# Implementing IBM Systems Director Active Energy Manager 4.1.1

Introduces energy management concepts and technologies

Discusses power and energy efficiency on IBM systems

Includes a step-by-step guide on implementing and using Active Energy Manager



David Watts
Phil Ainsworth
Srikanth Aithal
ByungJoon Choi
Mario Rodriguez

Redbooks

ibm.com/redbooks

International Technical Support Organization

# Implementing IBM Systems Director Active Energy Manager 4.1.1

October 2009



SG24-7780-00

**Note:** Before using this information and the product it supports, read the information in “Notices” on page ix.



**First Edition (October 2009)**

This edition applies to IBM Systems Director Active Energy Manager Version 4.1.1.

# Contents

# Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:
*IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.*

**The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:** INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.



 

# Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

AIX®
AS/400®
BladeCenter®
DB2®
DPI®
DS4000®
DS6000™
DS8000®
EnergyScale™
i5/OS®
IBM Systems Director Active Energy Manager™
IBM®
iDataPlex™
iSeries®
Maximo®
NetView®
Power Architecture®
Power Systems™
POWER6®
PowerExecutive™
POWER®
ProtecTIER®
Rational®
Redbooks®
Redpaper™
Redbooks (logo) ®
Smarter Planet™
System i®
System p®
System x®
System z10™
System z®
Tivoli®
WebSphere®
z/VM®

The following terms are trademarks of other companies:

Java, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

# Preface

Energy efficiency is a critical priority for IT managers because energy and power costs can be a significant portion of IT costs. Thus, understanding and investing in energy management is critical. With IBM® Systems Director Active Energy Manager™, an extension of IBM Systems Director, you can monitor and manage the power usage of systems. Originally designed to support IBM BladeCenter® and System x®, Active Energy Manager now supports the power management of additional IBM systems, including System z® mainframes and POWER6® processor-based systems, as well as storage devices using an intelligent Power Distribution Unit (PDU+).

Active Energy Manager helps determine the proper power allocation for each system in the data center. It can assist in determining how to allocate power to existing systems more efficiently so that additional systems can be accommodated without the need for additional power and cooling infrastructure. When power is constrained, the power management functions of Active Energy Manager allow power to be rationed on a system-by-system basis, enabling available processing power to match current workload closely.

This IBM Redbooks® publication can help system administrators effectively monitor and manage the power usage of systems in a data center. This book introduces energy management concepts and technologies, and then provides a step-by-step guide to implementing and using Active Energy Manager.

## The team who wrote this book

This book was produced by a team of specialists from around the world working at the International Technical Support Organization (ITSO), Raleigh Center.

**David Watts** is a Consulting IT Specialist at the IBM ITSO Center in Raleigh. He manages residencies and produces IBM Redbooks publications on hardware and software topics related to System x servers and associated client platforms. He has authored over 50 IBM Redbooks publications. He has a Bachelor of Engineering degree from the University of Queensland (Australia) and has worked for IBM for over 16 years. He is an IBM Certified IT Specialist.

**Phil Ainsworth** is an Senior IT Specialist with IBM Australia and works in the Systems and Technology Group as a presales technical specialist. He has 28 years of experience at IBM and has provided support for the Power i platform and its predecessors since 1980. He now provides presales technical support for Storage, specializing in N series. Since 1988 Phil has co-authored 10 IBM Redbooks publications about AS/400®, iSeries®, and System i® technologies, as well as the previous Active Energy Manager IBM Redbook publication. He is also a regular presenter at IBM conferences and other technical events.

**Srikanth Aithal** is an STG Lab services consultant in India. He has worked with IBM since 2007. His areas of expertise include IBM System x, BladeCenter systems, IBM Systems Director, Active Energy Manager, and VMware Infrastructure. He also focuses on Green IT related aspects and server consolidation projects.

**ByungJoon Choi** is an IT Specialist with IBM Korea and works in the Systems and Technology Group as a presales technical specialist. He has worked with IBM since 2007. His areas of expertise include IT simplification consulting, IBM System x, BladeCenter systems, and IBM Systems Director. Other areas of interest include various third-party virtualization technologies, such as VMware, Hyper-V, XEN, and Virtual Iron. Recently, he has focused on Green-related technologies, such as Active Energy Manager, Rear Door Heat eXchanger, iDataPlex™, and various sensors.

**Mario Rodriguez** is an IT Specialist in IBM Uruguay since 2001. He holds MCSE, AIX®, LPI and other Comptia certifications. His areas of expertise include SAN switches (Brocade, Cisco MDS), SAN Storage (DS3000, DS4000®, DS6000™, and DS8000®), Linux®, AIX, and VMware.

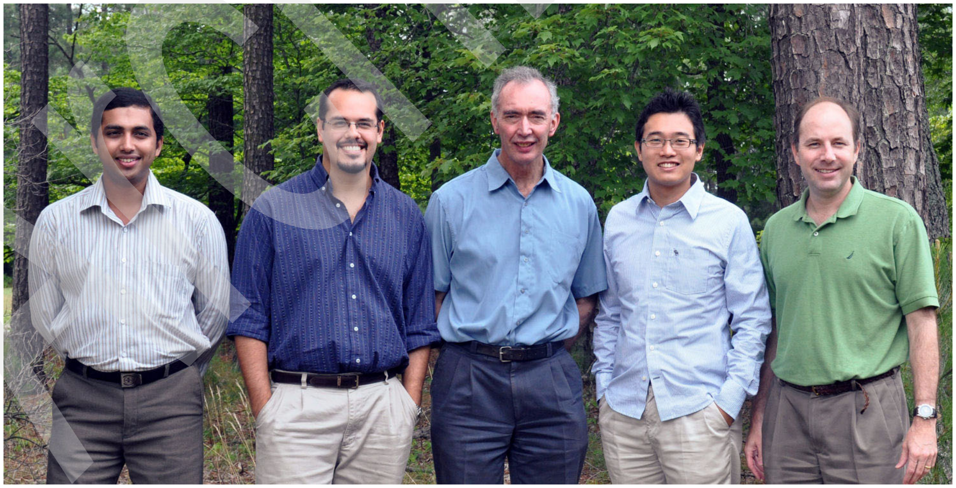

*The author team (l-r): Srikanth, Mario, Phil, ByungJoon, David*

A special thanks to Mary Wigham for her tireless efforts in helping us produce this book.

Also thanks to the following people for their contributions to this project:

Active Energy Manager development:

- Doug Beauchene
- Brad Behle
- Louis Behrens
- Andreas Bieswanger
- Ajay Dholakia
- Gordie Grout
- Jake Kugel
- Jacob Morlock
- Leigha Peterson
- Scott Piper
- Nathan Rabe
- Karthick Rajamani
- Matt Riedemann
- Michael Sullivan
- Mark Vanderweil
- Jeff Van Heuklon
- Janet Weber
- Mary Wigham

IBM Redbooks publications:

- Octavian Lascu
- Linda Robinson
- Margaret Ticknor
- Tamikia Barrow
- Dave Bennin

Other people from around the world:

- Don Roy, Raleigh, NC, U.S.
- KeeSeon Noh, Seoul Korea
- The UNIX® Software Group, Uruguay
- Thomas Domin, Bangalore India
- Sambath NA Parthasarathy, Bangalore India

Thanks to the authors of IBM Redpaper™ publication, *Going Green with IBM Systems Director Active Energy Manager 3.1*, REDP-4361:

- Phil Ainsworth
- Miguel Echenique
- Bob Padzieski
- Claudio Villalobos
- Paul Walters

## Become a published author

Join us for a two- to six-week residency program! Help write a book dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You will have the opportunity to team with IBM technical professionals, Business Partners, and Clients.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you will develop a network of contacts in IBM development labs, and increase your productivity and marketability.

Find out more about the residency program, browse the residency index, and apply online at:

**ibm.com**/redbooks/residencies.html

## Comments welcome

Your comments are important to us!

We want our books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks publications in one of the following ways:

- Use the online **Contact us** review Redbooks form found at:

  **ibm.com**/redbooks

- Send your comments in an e-mail to:

  redbooks@us.ibm.com

- Mail your comments to:

  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400

# 1

# Introducing IBM energy management

IBM Systems Director Active Energy Manager is a cornerstone of the IBM green infrastructure because it enables you to monitor and manage energy consumption in the data center. For Active Energy Manager to monitor and manage the energy consumption of IBM Power systems, it links into the IBM EnergyScale™ architecture which is implemented on IBM POWER6 processors.

EnergyScale provides a number of energy management capabilities of which Active Energy Manager can take advantage. Therefore, we describe the functions and capabilities of the EnergyScale architecture in detail in this chapter.

Active Energy Manager also supports energy management techniques used in IBM x86 architecture-based servers. We provide an overview of how these techniques are implemented on IBM servers, and the support provided by Active Energy Manager.

In this chapter, we also describe the components, techniques, equipment, and designs that comprise a *green* infrastructure and how this green infrastructure relates to the IBM Smarter Planet™ initiative.

## 1.1 The IBM approach to energy management

In terms of global issues, energy efficiency is one of the most important problems that we need to overcome. In the U.S. and elsewhere, computer data centers are at a tipping point. If left unchecked, the cost to power and cool servers in the future might well equal or even exceed the cost of acquisition.[1]

Data centers consume 10 to 100 times more energy per square foot than a typical office building and, therefore, represent a large potential for impact on electricity supply and distribution. At current usage rates, data centers in the U.S. are expected to double power requirements over the next 5 years.

So what are the drivers for energy management? To understand the drivers, you need to understand the exact power and cooling costs in a data center and break down the energy costs of all the data center components, not just the servers and associated equipment. For example, the energy consumed to cool a data center can be as high as 1.5 times the energy consumed by the center's computing resources.

Another important factor to consider in terms of energy planning and budgeting is the over allocation (safety margin) of the power supplied versus the power consumed. If you know how much power is actually consumed in the data center, rather than simply adding up the rated power consumption of all the center's devices, you can avoid over expenditure on power and cooling infrastructure. You can also take advantage of the infrastructure that is currently installed without having to make large disruptive changes.

In the future, companies will need to manage the energy efficiency of the data center environment by:

- Making more efficient use of existing hardware
- Saving energy in times of low utilization
- Improving performance without increasing the amount of infrastructure

Companies can manage energy efficiency by using tools that show *trending* of current energy consumption on individual or groups of systems to show growth *patterns* in power consumption and cooling *needs* and by actively controlling the amount of power that these systems consume.

IBM takes an integrated approach to energy management in the data center by tackling the problem at all levels:

- Data center design to optimize energy efficiency and growth options
- Hardware design for energy efficiency leadership

---

[1] See *Assessing trends over time in performance, costs, and energy use for servers* by Jonathan G. Koomey, et al. http://download.intel.com/pressroom/pdf/servertrendsrelease.pdf

- Software design to optimize processing efficiency and data management
- Virtualization and consolidation to boost utilization per watt
- Tools to measure and manage energy consumption in the data center

These are the pillars of energy efficiency in the data center and can form the framework for information technology infrastructure into the future. Let us look at each of these levels in more detail.

## 1.1.1 Data center design

IBM has adopted an innovative approach to data center design, as illustrated in Figure 1-1. Scalable modular design enables the speedy implementation of turnkey facilities from 500 to 2,500 square feet or of larger enterprise centers up to 20,000 square feet. A fully functional portable design can be deployed in around 3 months. IBM has also introduced the concept of *high density zones* within the data center that can increase server density without the high cost of a complete site retrofit.

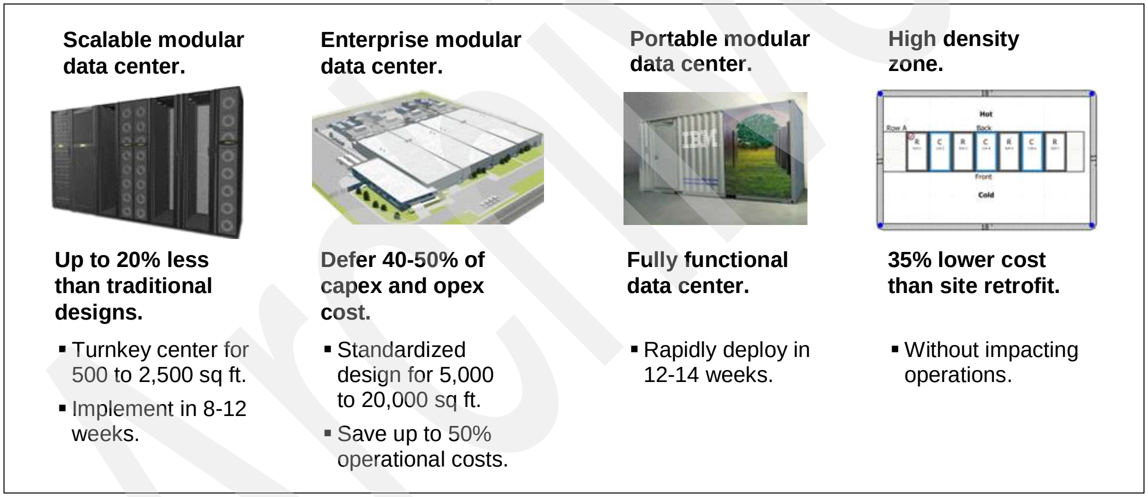


*Figure 1-1 IBM data center design leadership*

## 1.1.2 Hardware design

IBM server design is at the forefront of energy efficiency and technical excellence across all platforms, as illustrated in Figure 1-2.

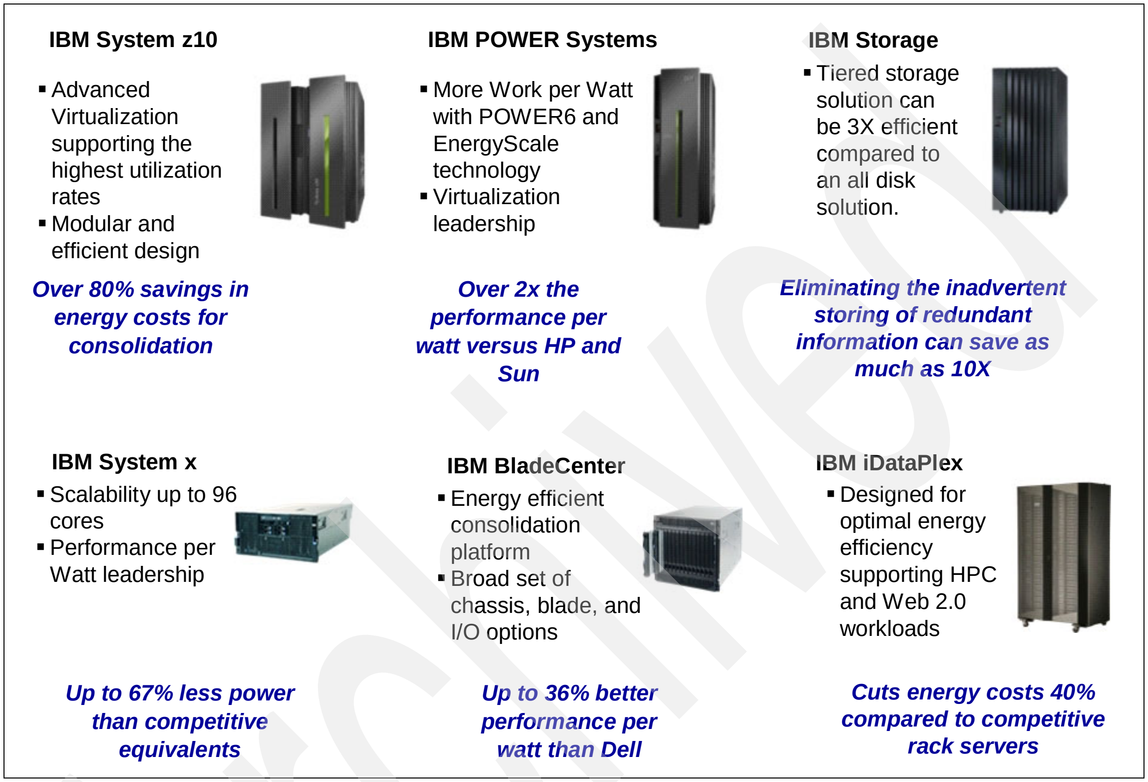


*Figure 1-2 IBM hardware design leadership*

The innovative hardware design includes the following systems:

- IBM System z10™ employs advanced virtualization techniques at the enterprise level to provide extremely high processor utilization levels.
- IBM POWER® processor-based systems provide virtualization leadership in the UNIX space, and employ the patented IBM EnergyScale architecture for advanced power monitoring and management.
- IBM System x86 rack and stand-alone servers are highly scalable, thereby providing a robust platform for server virtualization in addition to energy efficiency leadership.
- IBM BladeCenter is a highly energy efficient packaging infrastructure which supports a broad range of chassis, x86 and POWER blades, and I/O options.

- IBM iDataPlex is a recent innovation from IBM that is specifically designed for optimal energy efficiency, and optimized for high performance computing and Web 2.0 workloads.
- The broad range of storage products from IBM are integrated tightly with IBM software to provide energy efficient tiered storage, as well as advanced deduplication techniques to minimize storage redundancy.

## 1.1.3 Software design

Apart from the obvious data center and hardware components of an energy efficient IT infrastructure, IBM provides a suite of middleware that optimizes the software environment in the data center:

- IBM Tivoli® software to manage a tiered data backup and storage environment in addition to data center operations
- IBM WebSphere® software for application management and virtualization
- Combined hardware and software solutions, such as the industry recognized ProtecTIER® deduplication gateway
- IBM Rational® software for software development and collaboration

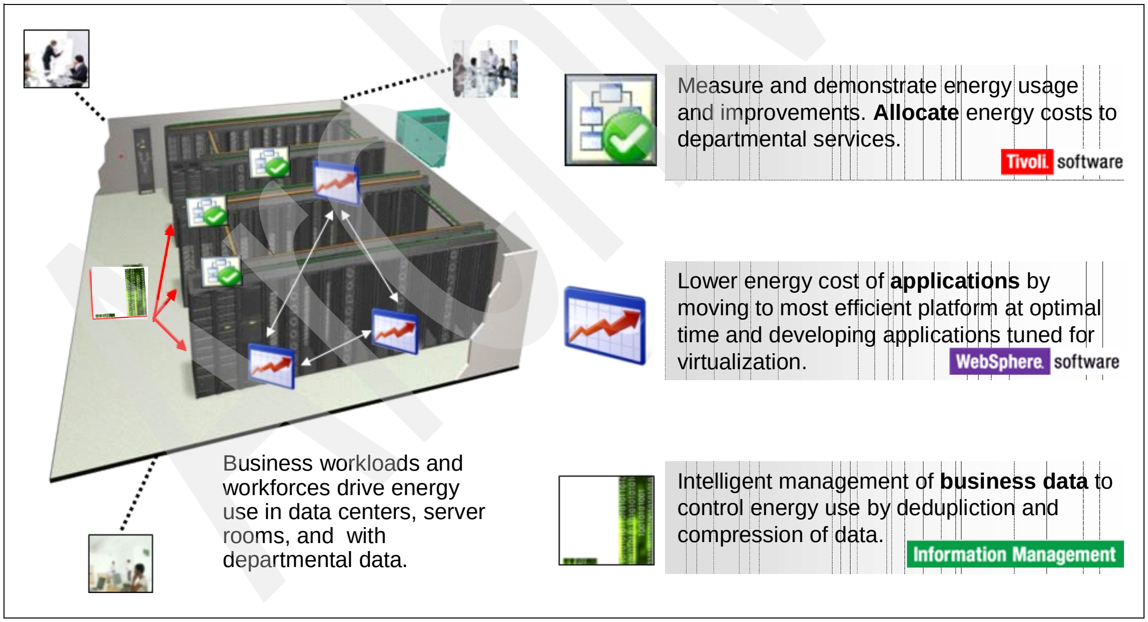


*Figure 1-3 IBM software design leadership*

## 1.1.4 Virtualization and consolidation

IBM hardware and software products support virtualization at the following levels (as shown in Figure 1-4):

- Server virtualization

  IBM builds advanced hardware virtualization into products, such as System z and IBM Power systems and also supports virtualization products from companies such as VMware. There is also the Virtualization Manager plug-in for IBM Systems Director.

- Storage virtualization

  IBM has one of the broadest range of disk, tape, and associated storage products in the industry and provides industry-leading tiered storage and deduplication products for unmatched storage virtualization.

- Client virtualization

  IBM has a long history of providing client virtualization solutions, as well as support for industry-leading solutions from other vendors such as VMware and Microsoft®.

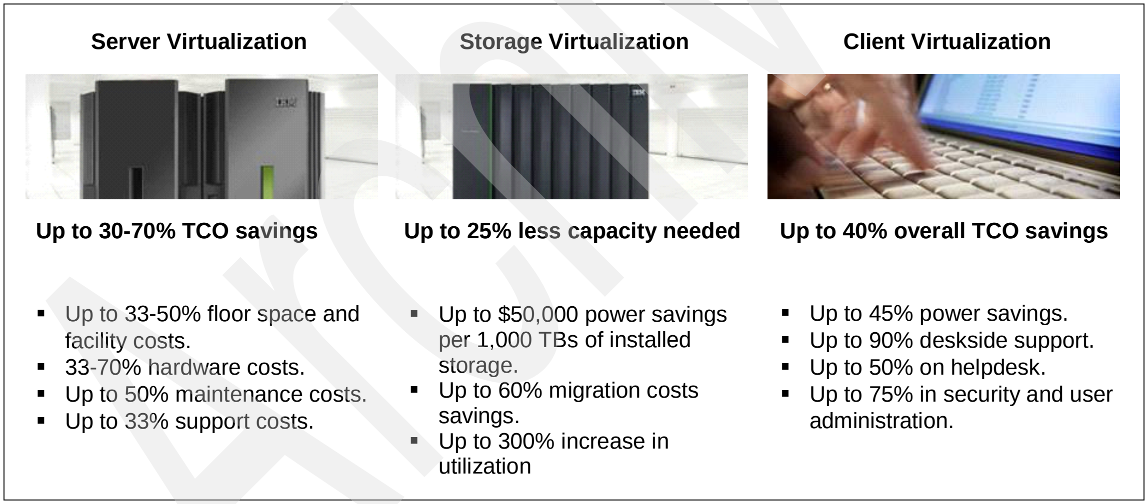


*Figure 1-4 IBM virtualization and consolidation leadership*

## 1.1.5 Energy management tools

You cannot manage energy usage in the data center unless you have the tools to monitor and actively manage power consumption. The IBM Systems Director Active Energy Manager product is a tool that you can use to:

- Monitor power consumption and thermal signature at a system and device level
- Manage the energy consumption of POWER6 processor-based systems by actively reducing the amount of power consumed during periods of reduced utilization
- Better plan and budget for energy consumption by setting a cap on the amount of power that supported systems can consume.

As shown in Figure 1-5, you can integrate Active Energy Manager into the Tivoli software environment to provide a consolidated approach to data center management and operations.

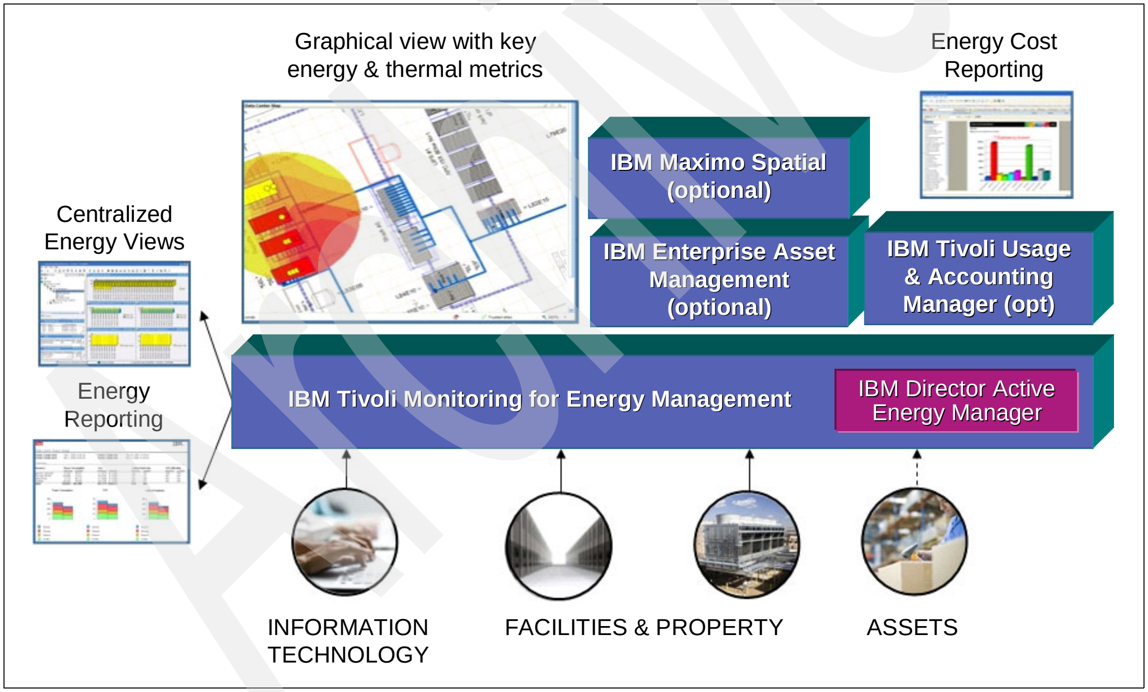


*Figure 1-5 Integrating Active Energy Manager to the Tivoli management and operations environment*

## 1.2 Active Energy Manager integration into Tivoli

Active Energy Manager is integrated into Tivoli through the IBM Tivoli Monitoring for Energy Management (ITMfEM) product. ITMfEM uses the IBM Tivoli Monitoring infrastructure, taking full advantage of its capabilities.

ITMfEM unifies energy, thermal, and other environmental metrics from IT equipment, data center infrastructure components, buildings and facilities, enterprise assets, and custom energy monitoring agents into a single management system that enables customers to have an integrated view of energy usage. The solution includes the integration of information from a range of leading power, cooling, and building management partners, allowing for a holistic view of energy consumption with historical awareness, trending, reporting, and rich visualization capabilities. ITMfEM provides a key link between the domains of energy management and service management to help customers reduce energy use while maintaining IT service levels.

IBM Systems Director and Active Energy Manager are included with ITMfEM. ITMfEM provides many of the Active Energy Manager capabilities available with IBM Systems Director, while extending the solution with additional capabilities as well as integrating energy information from additional partners. The resources that are monitored by Active Energy Manager, along with the energy and environmental metrics collected for those resources, flow directly into ITMfEM. Thus, you can take advantage of an existing Active Energy Manager installation or deploy a complete Tivoli energy management solution that includes Active Energy Manager.

You can install the ITMfEM Active Energy Manager component that is provided with ITMfEM, and the resources and data begin flowing between Active Energy Manager and ITMfEM. The resources that Active Energy Manager monitors become resources that ITMfEM monitors. ITMfEM then communicates with Active Energy Manager to obtain the energy and thermal metrics from those resources as needed. When IBM Systems Director discovers new resources and Active Energy Manager begins monitoring those resources, they flow automatically into ITMfEM.

Resources and metrics received from Active Energy Manager are visualized in the Tivoli Enterprise Portal workspaces. Energy and environmental metrics are used in situation definitions to alert operations and facilities when thresholds are reached or exceeded. These metrics are stored in the Tivoli Data Warehouse to provide historical views and reports. You can use these ITMfEM reports to analyze energy usage across the enterprise and to analyze energy metrics and IT metrics together to understand utilization, cost, and opportunities for efficiency gains. Included are over 20 reports, ready to use as is, that range from industry-standard Data Center Infrastructure Efficiency (DCIE) metrics, to optimization recommendations for server equipment, with an ability to easily create additional reports or dashboard views.

With the broad range of energy information that ITMfEM collects and the predefined thresholding that the solution provides, resulting threshold exceptions can generate events that are forwarded to the event management system, such as Tivoli's Omnibus, and incorporated into the overall service desk.

The data from ITMfEM also enables additional capabilities that take advantage of the energy information that is collected, as shown in Figure 1-6. For example, you can combine the energy data from ITMfEM with the Tivoli Maximo® solution to provide near real-time thermal heat maps of the data center. You can also combine the data with Tivoli Usage and Accounting Manager to provide financial accounting of energy usage. In addition, you can integrate the information with Tivoli Business Service Manager to provide rich energy management dashboards or to provide energy awareness at a business service level.

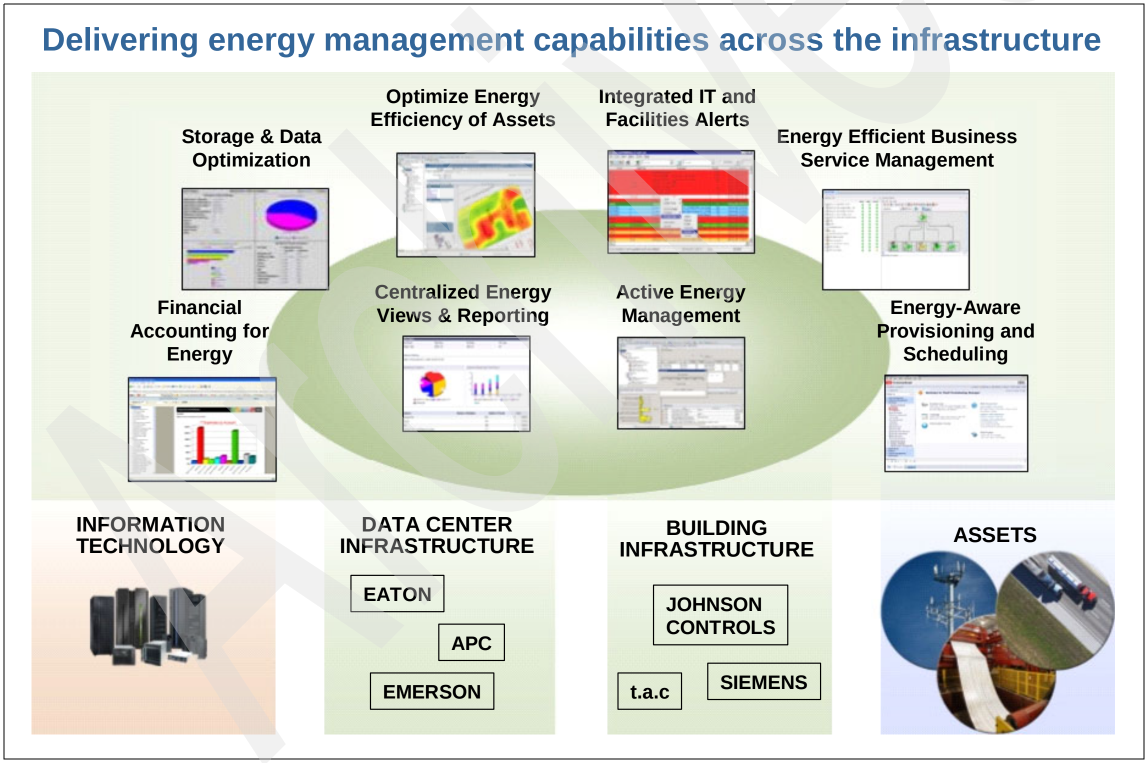


*Figure 1-6 Delivering energy management capabilities across the infrastructure*

For more information about these extended energy and environmental management solutions, consult the IBM white paper *Greening the data center with IBM Tivoli software*, which is available at:

http://www.ibm.com/software/tivoli/products/monitor-energy-management

## 1.3 IBM EnergyScale

Energy consumption is becoming an increasingly important issue in IT-based businesses. The energy that is required to power and cool computers can be a significant cost to the business, reducing profit margins and consuming resources. In addition, the cost of creating a power and cooling infrastructure can be prohibitive to business growth.

In response to these challenges, IBM EnergyScale technology was developed for IBM POWER6 processor-based systems. Although we discuss EnergyScale in terms of the POWER6 processor, this technology is applicable to, and will be enhanced on, future models of the IBM POWER processor line.

EnergyScale provides functions that help you to monitor and control IBM server power usage and cooling costs, which enables better facility planning, provides energy and cost savings, enables peak energy usage control, and increases system availability. Administrators can use EnergyScale capabilities to control the power consumption and performance of POWER6 processor-based systems to meet their particular data center needs.

It is important to note that the EnergyScale architecture is specific to the IBM POWER6 and later processors. Processors from other manufacturers implement power saving techniques in different ways compared with the POWER processors from IBM.

In this section, we describe the functions that EnergyScale provides along with usage examples, hardware and software requirements, and information about each of the functions that are available on IBM server offerings. All POWER6 processor-based systems support some, if not all, EnergyScale features. In turn, most POWER6 systems support IBM Systems Director Active Energy Manager, providing a single energy management interface to monitor and control these systems.

The family of POWER6 server models available include the following (as shown in Figure 1-7):

- IBM BladeCenter JS12 Express
- IBM BladeCenter JS22 Express
- IBM BladeCenter JS23 Express
- IBM BladeCenter JS43 Express
- IBM Power 520 Express
- IBM Power 550 Express
- IBM Power 560 Express
- IBM Power 570
- IBM Power 575
- IBM Power 595


*Figure 1-7 IBM POWER processor-based systems family*

The following sections describe each of the EnergyScale functions available on IBM POWER6 systems and their prerequisites. Note that all of the features described might not be available on a particular model.

## 1.3.1 Prerequisites

EnergyScale support has the following prerequisites:

- POWER6 processor or later

  POWER6 is the minimum level that supports EnergyScale functions.

- Service processor firmware level

  Support for EnergyScale functions is dependent on the level of firmware that is installed on the IBM Power system's service processor. Table 1-4 on page 30 lists the functional support for EM340 level of firmware. Support for previous levels of firmware is documented in *IBM EnergyScale for POWER6 Processor-Based Systems*, which is available at:

  http://ibm.com/systems/power/hardware/whitepapers/energyscale.html

- Operating system version

  EnergyScale requires the following levels of operating system support:

  – Compliant

    Operating systems that are EnergyScale *compliant* support the processor core nap function and tolerate power savings and power capping. These systems tolerate dynamic changes in processor speed, but are generally unaware of them. That is, they cannot recalculate processor utilization when power savings or power capping are in operation, and they report processor utilization as though the processor or processors are running unthrottled.

  – Ready

    Operating systems that are EnergyScale *ready* support the processor core nap, power savings, power capping, and dynamic power savings functions. They are also aware of dynamic changes in processor speed and are ready to support energy aware applications. These systems generally provide both processor time and scaled processor time in their system call interfaces.

  – Enabled

    Operating systems that are EnergyScale *enabled* are EnergyScale *ready* and also have performance and accounting tools that are aware of dynamic changes in processor speed. On these systems, energy awareness extends throughout the operating system's commands and utilities.

  For more information about processor utilization and accounting in a scaled environment, refer to 1.3.11, "EnergyScale impacts on processor utilization and accounting" on page 31.

Table 1-1 shows the minimum operating system level that is required to support EnergyScale functions and the compliance status of each operating system.

*Table 1-1 Operating system requirements*

| Operating system | Compliant | Ready | Enabled |
|---|---|---|---|
| AIX | 5.2 TL10<br>5.3 TL6 | 5.3 TL6 | 5.3 TL6 |
| SUSE Linux Enterprise Server | SLES 10 | SLES 10 | Not supported |
| Red Hat Enterprise Linux | RHEL 4.6 | RHEL 5.1 | Not supported |
| IBM i | V5R4M5 | V6R1M0 | Not supported |

## 1.3.2 Trending

In the context of the EnergyScale architecture, *trending* is the name given to the tracking of power consumption, thermal signature, and effective CPU percentage over time. EnergyScale defines the following types of trending:

- Power trending
- Thermal trending
- CPU trending

In this section, we provide a brief discussion of each type of trending. For more details about trending, refer to 6.7, “Viewing trend data” on page 221.

### Power trending

EnergyScale provides continuous collection of real-time server power consumption. Active Energy Manager can display or export this power usage data. Administrators can use such information to predict data center power consumption at various times of the day, week, or month. In addition, data can be aggregated across a number of systems to identify anomalies, manage power loads, and on certain models, enforce system-level power budgets.

Refer to the power capping examples that we describe in 7.3, “Setting a power cap” on page 306 to learn how to use power trending data to set a power budget (also referred to as a *power cap*).

**Note:** The collection of power data on some POWER6 systems, such as the IBM Power 560 and IBM Power 570 models, requires that an intelligent Power Distribution Unit (PDU+) provide power to the system. To learn how to set up a PDU and use IBM Active Energy Manager in this environment, refer to Chapter 8, “Using a PDU with Active Energy Manager” on page 367. Other POWER6 systems collect power data internally using built-in power meters and do not require additional hardware.

#### Thermal trending

Active Energy Manager can display a measured ambient temperature and a calculated exhaust temperature. This information can help identify data center hot-spots that need attention. Active Energy Manager can also monitor the exhaust temperature and use the data to trigger an action, such as reducing power to the system in the event of a dangerous rise in temperature.

#### CPU trending

When the power savings function is invoked, the processing power of the POWER6 processor is throttled back by reducing the clocking frequency and supplied voltage. The CPU is also throttled back when power capping needs to enforce a cap. CPU trending, which Active Energy Manager displays, shows the change in effective CPU speed relative to the rated speed (100%) over time.

### 1.3.3 Power savings

The power savings function is one of the key benefits of the EnergyScale architecture on the POWER6 processor. As previously mentioned, it enables the processing power of the POWER6 processor to be throttled back by reducing the clocking frequency and supplied voltage. This throttling can be implemented on demand or scheduled using a policy.

For a system with a single logical partition, power savings is quite straightforward. However, you need to experiment when implementing power savings on a logically partitioned system where there are multiple processors. The key point to note in this situation is that power savings is turned on at a *system* level, not at a processor or partition level. Thus, all processors on a POWER6 system are equally constrained in terms of their frequency (and, therefore, processing capacity) when power savings is turned on.

If all partitions in a system are using the same shared processor pool, there should be no difference in the overall effect of power savings between multiple partitions and a single partitioned system, because the CPU utilization is averaged across all processors. Therefore, all partitions are equally affected, depending on the type of power savings mode selected (static or dynamic).

However, in a situation where you have multiple processor pools (for example two or more dedicated processor partitions), each partition might be affected differently. For example, if you have a very busy dedicated processor partition and an idle one, in the case of static power savings (and to a lesser extent dynamic power savings - favor power savings), the system firmware throttles back the frequency of *all* processors on the system equally, even though some processors are busy and some are in an idle state. In this case, the busy partition might suffer a decrease in performance, even though the average utilization across all processors might be very low.

In this scenario, you might want to select the “power savings - favor performance” option (Dynamic Power Saving) because this option essentially allows the processors in all partitions to run unconstrained when required. It also saves power during periods of low workload demands. You might not save as much power as when using static power savings or dynamic power savings that favor power savings.

Note that the algorithms built in to the firmware that control power savings try to adjust for the fact that some processors in a multiple processor pool configuration are very busy and some are idle. Instead of simply calculating an average processor utilization across all pools, the algorithms adjust for the fact that some processors are very busy and some idle. Therefore, the *average* utilization calculated by the firmware can in fact be lower than the actual. This tends to delay the onset of throttling of the busy partitions.

Figure 1-8 shows the relationship between available processor capacity and the processor capacity requested by the workload when power savings is either off or not supported. In this case, the processor is running at 100% of its nominal frequency (for example 5.0 GHz) and all of the processor capacity is available to the workload (which varies over time as shown in the figure). Notice that the processor runs at its nominal speed (100%) no matter what the workload is.

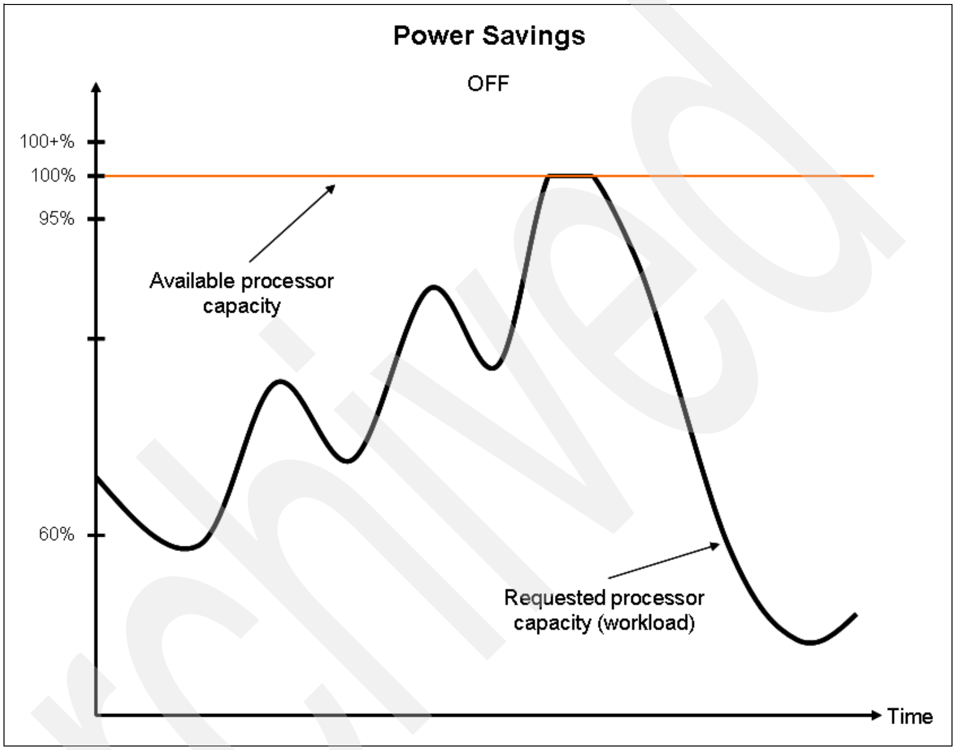


*Figure 1-8 Power savings: Power savings OFF*

In Figure 1-8, notice that the available processor capacity cannot rise above 100% even though there is a peak in the workload that can drive it higher.

There are two ways that you can activate power savings: statically or dynamically. We describe both methods in the following sections.

## Static power savings

*Static power savings* lowers the processor frequency and voltage on a system by a fixed amount, thereby reducing the power consumption of the system while still delivering predictable performance. This percentage is predetermined to be within a safe operating limit and is not user configurable. See Table 1-2 on page 22 for details about the actual percentage frequency drop and estimated processor power saved for each POWER6 processor-based system model and firmware release. On systems with supported firmware and supported DIMMs, memory is also allowed to enter a low power state when no memory accesses are occurring.

Active Energy Manager is the recommended user interface to enable or disable static power savings mode, although it can also be switched on and off through the Advanced Systems Management Interface (ASMI) and the Hardware Management Console (HMC) interfaces. ASMI is the graphical interface to the service processor on POWER6 machines, and the HMC controls IBM Power systems for the purpose of logical partitioning and related functions.

Static power savings can be enabled based on regular variations in workloads, such as predictable dips in utilization over night or over weekends. Also, static power savings can be used to reduce peak energy consumption, which can lower the total cost of power used. Note that when static power savings is enabled for workloads with low CPU utilization, workload performance is not affected.

The only time that a power savings-capable system does not support operating in power savings mode is during a system boot or reboot. Power savings can be enabled at any time; however, if power savings is enabled prior to a system boot, the processor voltage and frequency remains at the default (normal) levels until the service processor firmware reaches a standby or running state. Immediately before the firmware executes, the voltage and frequency enter power savings mode.

If a system reboot occurs while in power savings mode, the voltage and frequency are first raised back to normal levels. Following a successful system reboot, the voltage and frequency are dropped back to power savings mode levels. The power savings mode setting persists across system boots, service processor resets, and the loss of AC power, unless the power outage is long enough to drain the service processor's NVRAM battery, which retains the power savings mode setting. If the administrator has set an AEM policy for this system, AEM will put the system back into power savings mode, even if the NVRAM has lost this data.

Power savings is only supported on systems with 4.0 GHz or faster POWER6 processors. At lower frequencies the power saving realized by using the power savings function is minimal and, therefore, is not implemented. Note that some older versions of operating systems running on POWER6 processor-based systems might not report the correct information about utilization and accounting if the machine is placed into power savings mode. Refer to 1.3.11, "EnergyScale impacts on processor utilization and accounting" on page 31 for details.

Figure 1-9 shows a simplified view of the relationship between available processor capacity and the processor capacity that is requested by the workload when static power savings is turned on. The available processor capacity is reduced to a value that depends on the IBM Power system being used (refer to Table 1-2 on page 22). This reduced processor capacity in turn places a hard cap on the amount of work that the processor can perform, as shown by the requested processor capacity curve in Figure 1-9.

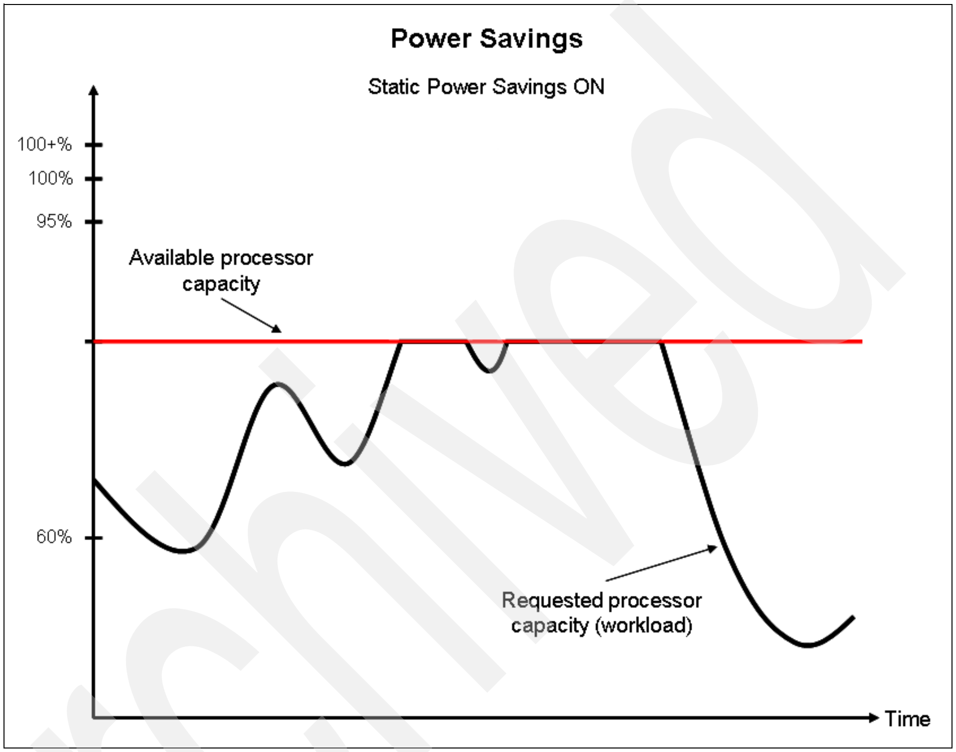


*Figure 1-9 Power savings: Static power savings ON*

In Figure 1-9, notice that the available processor capacity cannot rise above the imposed value even though there are peaks in the workload that could drive it higher.

## Dynamic power savings

Active Energy Manager includes the capability to implement an enhanced version of power savings called *dynamic power savings*. Active Energy Manager is the only interface that supports the activation of dynamic power savings.

Dynamic power savings varies the processor frequency and voltage based on the utilization of the system's POWER6 processors, where utilization is proportional to workload. The frequency and utilization of a processor are inversely proportional for most workloads, where utilization is defined as the percentage of the processor's capacity that is used to satisfy the workload demands. Therefore, for a constant workload, as the frequency of a processor decreases, its utilization increases.

Dynamic power savings takes advantage of this relationship between processor frequency and utilization to detect opportunities to save power, based on measured real-time system utilization. As system utilization decreases, the processor frequency can be turned down to match. Therefore, the POWER6 processor's utilization can still be maintained at a high level for a reduced workload using dynamic power savings. At lower workloads, power consumption (and therefore heat produced) are reduced without adversely affecting system performance.

To summarize, as the workload demands on the POWER6 processor go down, the processor's processing capacity is throttled back to match. Conversely, when the workload increases, the processor is throttled up to provide additional processing capacity to match the increased workload.

At higher workloads, the performance implications of activating dynamic power savings depend on whether the *favor system performance* or *favor system power savings* option is configured in Active Energy Manager. When dynamic power savings is enabled in either mode, the system firmware continuously monitors the performance and utilization of each of the system's POWER6 processor cores. Based on this data, the firmware dynamically adjusts the processor frequency and voltage, reacting to changes in workload within milliseconds to maintain acceptable performance, and also deliver power savings when the workload drops.

When a system is idle, the system firmware lowers the frequency and voltage to static power savings values. When fully utilized, the maximum frequency varies, depending on whether the favor system performance or favor system power savings option is selected:

- If the administrator selects the favor system power savings option and a system is fully utilized, the system firmware reduces the maximum processor frequency to 95% of the nominal value.
- If the favor system performance option is selected, the maximum processor frequency can climb to 100% of nominal and might in fact exceed nominal on certain platforms under ideal environmental conditions.

Refer to Table 1-3 on page 22 for more details about how dynamic power savings affects processor frequency on each supported platform.

### *Dynamic power savings that favor power savings*

Figure 1-10 shows a simplified view of the relationship between available processor capacity and the processor capacity requested by the workload when dynamic power savings (favor power savings) is turned on. The available processor capacity is reduced to a value which depends on the IBM Power system being used (refer to Table 1-2 on page 22). However, if the workload requires more processor capacity, the processor frequency increases automatically to allow for this up to 95% of nominal, at which point it is capped.

In Figure 1-10, notice that the available processor capacity cannot rise above 95% of nominal even though there is a peak in the workload that could drive it higher.

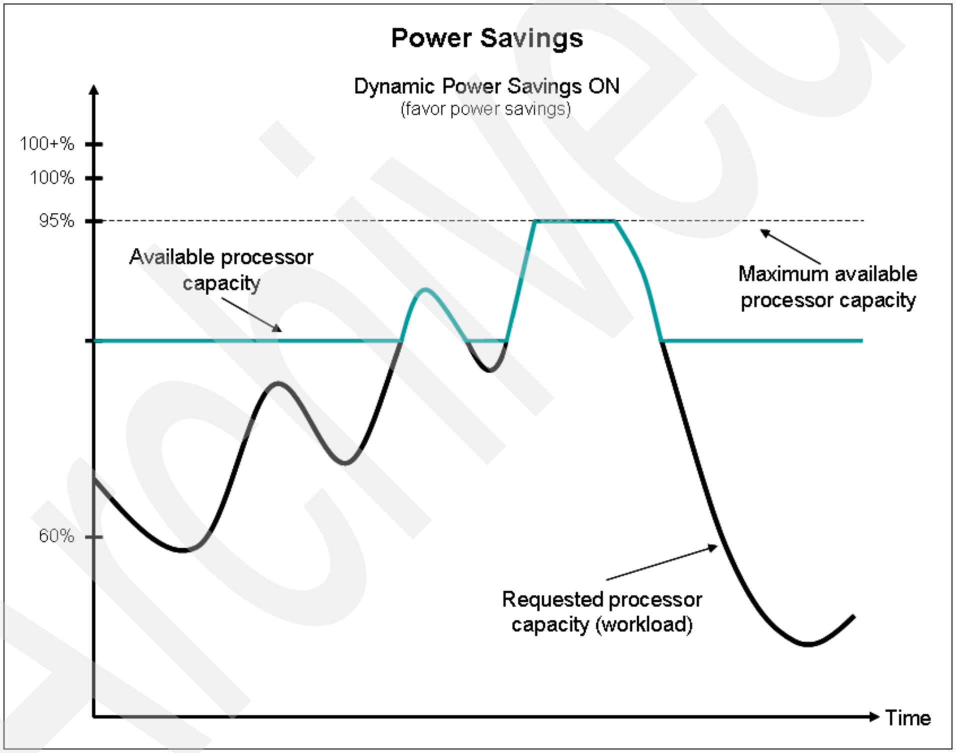


*Figure 1-10 Power savings: Dynamic power savings ON (favor power savings)*

### *Dynamic power savings that favor performance*

Figure 1-11 shows a simplified view of the relationship between available processor capacity and the processor capacity requested by the workload when dynamic power savings (favor performance) is turned on. For some IBM Power systems the available processor capacity can rise above 100% of nominal to a value which depends on the IBM Power system being used (refer to Table 1-2 on page 22). This increase in processing power results from the ability to overclock the processor to a frequency above 100%. Note that in this case, you might actually use *more* power than if dynamic power savings is turned off.

Dynamic power savings that favor performance does not constrain the amount of work performed but does save power when the workload requirement is low.

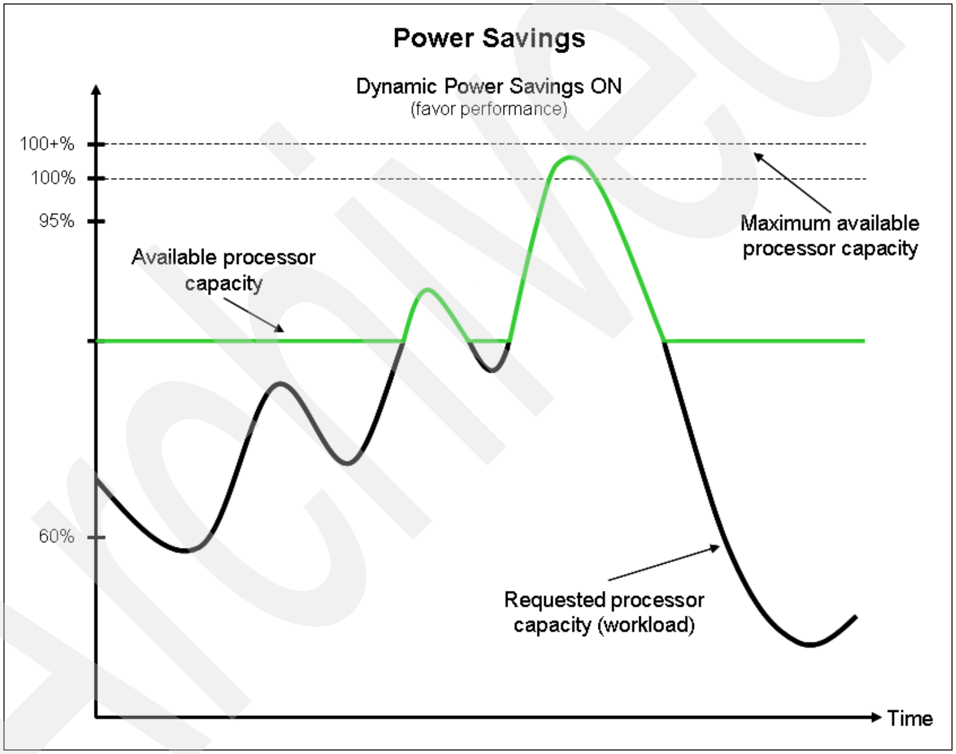


*Figure 1-11 Power savings: Dynamic power savings ON (favor performance)*

Note that dynamic power savings is mutually exclusive with static power savings mode. Therefore, only one of these modes can be enabled at a given time. Also, dynamic power savings is only available on systems with a supported firmware level, and on systems that are managed by Active Energy Manager. Refer to Table 1-4 on page 30 for details.

When either static power savings or dynamic power savings is enabled, power saving is achieved by lowering the processor frequency (clocking rate) and voltage. Table 1-2 shows the maximum frequency drop and corresponding power saving for each supported IBM Power system model.

*Table 1-2 Maximum frequency drop and estimated power saving*

| IBM Power system model | Maximum frequency drop | Estimated maximum power saving |
|---|---|---|
| 5.0 GHZ IBM Power 595 | 20% | 25 - 35% |
| 5.0 GHz IBM Power 570[a] | 14% | 20 - 30% |
| 5.0 GHz IBM Power 570[b] | 25% | 35 - 45% |
| 4.4 GHz IBM Power 570[a] | 8% | 5 - 15% |
| 4.4 GHz IBM Power 570[b] | 14% | 20 - 30% |
| All 4.2 GHz IBM Power systems[a] | 5% | 5 - 10% |
| All 4.2 GHz IBM Power systems[b] | 14% | 20 - 30% |
| JS12 blade server | 10% | 10 - 20% |
| All other systems with power savings support | 14% | 20 - 30% |

a. With GX dual port RIO-2 attachment
b. Without GX dual port RIO-2 attachment

As shown in Table 1-2 on page 22, the estimated maximum power saving is greater than the maximum frequency drop of the processor (which can be equated to a drop in processor capacity). This behavior is termed *super-linear*. The downside of super-linear behavior is that, if the processor capacity is increased by overclocking, the increase in power consumption is considerably greater than the increase in capacity (as measured by the increase in frequency).

Table 1-3 shows the maximum processor frequencies for IBM Power systems that support dynamic power savings.

*Table 1-3 Maximum frequencies for dynamic power savings*

| Power system model | Favor performance | Favor power savings |
|---|---|---|
| JS12, JS22 blade (<4.0 GHz) | 101.5% | 95% |
| JS12, JS22 blade (≥4.0 GHz) | 100% | 95% |
| IBM Power 520, 550 (≥4.0 GHz | 100% | 95% |

### 1.3.4 Power capping

With power capping you can specify a cap or limit on the power consumption for a system or group of systems which each system's firmware tries to enforce. If the cap is reached, the speed of the processor is throttled back and the voltage is reduced to try and keep the power used by the managed system under the cap.

Power capping is positioned as a planning and budgeting tool rather than a power saving technique as such. However, the saving of power is implied by better utilization of existing power capacity.

In most data centers and other installations, when a system is installed, a certain amount of power is allocated to it. Generally, the amount of power allocated is what is considered a safe value and typically has a large margin of reserved, extra power that is never used. This is called the *margined power*.

The main purpose of power capping is not to reduce power to a system (although it can do that) but rather to allow a data center administrator the ability to reallocate power from current systems to new systems by reducing the margined power that is assigned to the existing systems. In other words, power capping allows an operator the ability to add extra systems to a data center that previously had too much power allotted to its current systems. Power capping does this by guaranteeing that a system will not use more power than the power that is assigned to it, within certain defined limits.

Previously, the data center administrator had to plan the power consumption of the data center based on the Underwriters' Laboratories (UL) rating on the back of the servers being installed. This UL rating, commonly referred to as *label power* or *nameplate power*, represents the absolute worst case amount of power that the system can ever draw and is based on the capacity of the system's power supplies. It has to take into account a fully-configured system with the highest power usage components at the highest possible workload under the worst environmental conditions. This particular situation is almost always a highly unlikely scenario.

Power capping is similar to power savings in that both functions restrict the capacity of the processor in order to reduce power consumption. However, power savings is either on or off, whereas a power cap can be set between a minimum value (called a *minimum power cap* or *Pcap min*) and a maximum value (called a *maximum power cap* or *Pcap max*). These values are determined by an algorithm that takes into consideration all of the power-consuming devices within the managed system and an assumed workload.

Pcap min and Pcap max are calculated to ensure that the managed system always operates within safe power limits. If power capping needs to be enforced by the service processor firmware, it is the processor that has its power reduced, not other devices within the system such as disk drives or memory slots. When a power cap needs to be enforced, the service processor firmware turns down the processor's voltage and clocking rate to reduce power consumption to the required Pcap level.

Setting the power cap to Pcap min or Pcap max means that throttling starts to occur if power consumption reaches this value. For Pcap min, the amount of throttling depends on the propensity of the managed system to use power in excess of Pcap min. The degree of throttling will be somewhere between zero and the architected maximum. For Pcap max, the amount of throttling similarly depends on the propensity of the managed system to use power in excess of Pcap max, but throttling is far less likely to occur than if the power cap is set to Pcap min.

The BladeCenter management module displays information about allocated power and the power capping range for individual blade servers. The allocated maximum power cap for a blade server is not the worst case maximum amount of power that the blade server can consume (the label power), but is a typical maximum across various hardware configurations. This allocated maximum power cap is used by the management module to budget a typical maximum amount of power for a given blade server to determine if the blade server will fit within the BladeCenter domain's power budget and be allowed to power on. The maximum power in the power capping range for a blade server will be different than the allocated maximum power. The maximum power in the power capping range reflects the nameplate power for the blade server.

Figure 1-12 shows a simplified view of how hard and soft power capping are implemented.

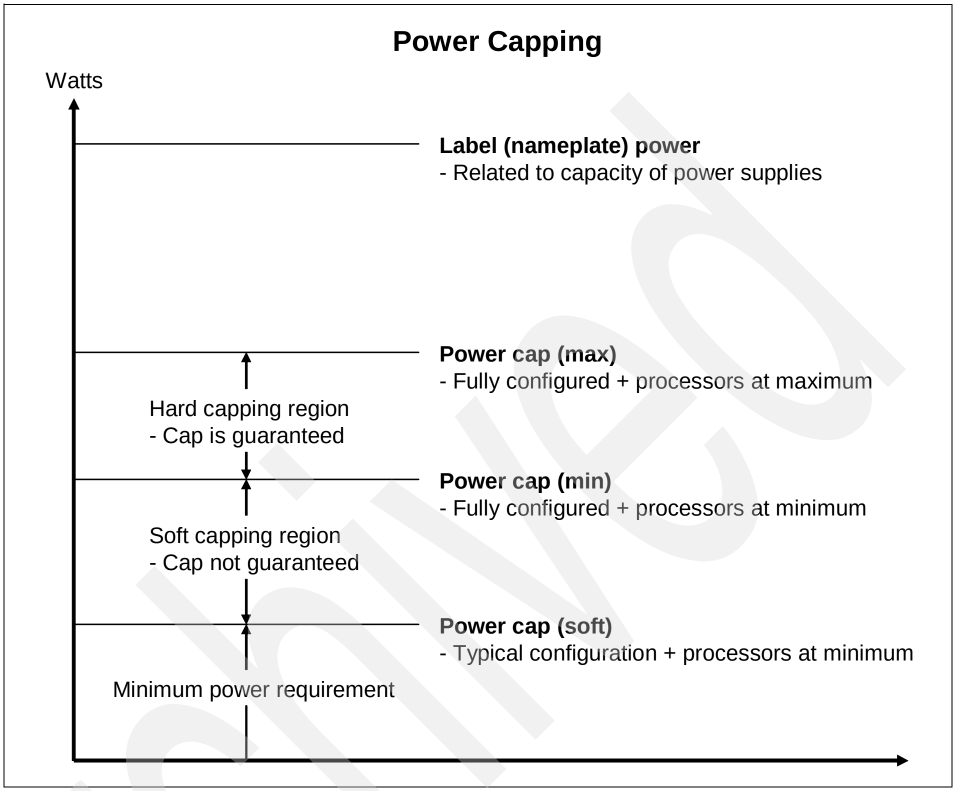


*Figure 1-12 Power capping implementation (simplified view)*

## Hard power capping

The system firmware calculates the Pcap max and the Pcap min according to an algorithm that considers the configured hardware as well as the potential workload on the processor. Notice that Pcap max is always less than the system's label (nameplate) power.

A power cap set between Pcap min and Pcap max is guaranteed not to exceed the power cap setting. To guarantee this, extreme system configuration and environmental conditions must be taken into account. Setting a power cap in this region allows for the recovery of some margined power, but in many cases no actual energy is saved because the allowable range is still well in excess of what the system typically consumes.

In cases where hard power capping needs to be enforced (the actual power consumption rises above the power cap), the system firmware reduces the processor frequency and voltage to bring the power consumption back under the power cap setting. This mechanism is the same mechanism that is used to implement power savings.

> **Note:** The hard power capping region is referred to as the *guaranteed* range in the Active Energy Manager graphical user interface.

### Soft power capping

*Soft power capping* extends the allowed power capping range down to lower values, beyond the hard power capping region that can be guaranteed for all configurations and environmental conditions. The extended range is called the soft power cap region, or *not guaranteed* range in the Active Energy Manager graphical user interface. You can think of the lowest value to which a soft power cap can be set as Pcap soft. Soft power capping is only supported on IBM POWER6 systems, not x86 architecture servers.

By setting a power cap in the soft power capping region, the system can save power by throttling back the processors with a corresponding lower level of performance. As for hard power capping, the technique used is the same as for power savings, that is by lowering the processor frequency and voltage. The performance impacts of a particular soft power cap setting can be determined using the power trending and CPU speed information displayed using Active Energy Manager.

> **Note:** A soft power cap is not guaranteed because the system firmware might not be able to reduce power consumption down to the power cap setting. In this case, an alert is logged in the IBM Systems Director event log.

A soft power cap setting will hold in most cases, depending on external factors such as room temperature. However, even if the soft power cap value is exceeded, the power consumed will never exceed Pcap Min, which is guaranteed.

If the power management goal is to try and meet a particular power consumption limit that is more representative of the actual power used, then soft power capping might be the preferred option. However, because it is not guaranteed, you cannot rely on it for power planning and budgeting purposes. Soft power capping works on a “best effort” basis.

**Note:** Failure to enforce a soft power cap below the minimum-guaranteed range is not an error and will not result in the system firmware entering safe mode. The system continues to operate normally with all EnergyScale features at the minimum-supported frequency until the soft power cap is disabled or raised.

### Safe mode

Certain failures can cause the system firmware to enter a *safe mode*. In safe mode all power and thermal information is longer reported to Active Energy Manager. An error is logged when safe mode is entered. When safe mode is entered and a power cap is set, the firmware drops the system to the power savings mode voltage and frequency to guarantee that the power cap set is still held. To re-enable full EnergyScale functionality, a firmware update, FRU replacement, or complete recycle of the system must be completed. If the problem that originally caused the system to enter safe mode is still present, the system re-enters safe mode and generates an additional error event.

## 1.3.5 Processor core nap

The IBM POWER6 processor can use a low power mode called *processor core nap* that halts execution when there is no work for the processor core to perform. The hypervisor firmware running on the POWER6 processor can use processor core nap as a general purpose idle state. This capability is always available, and is not controlled by AEM.

There are a maximum of two threads available on IBM Power processors. The second thread is made available when Simultaneous Multithreading (SMT) is turned on.

If the operating system detects that a processor thread is idle, it yields control of the thread to the hypervisor. The hypervisor immediately puts the thread into a nap state. If the operating system yields control of the second thread, and the processor core belongs to a dedicated processor partition, the second thread also enters nap mode. If the processor core is in a shared processor pool (used by multiple partitions) and if there is no micro-partition to dispatch, the hypervisor also puts the second thread into nap mode.

When both hardware threads running on a given processor core enter nap mode, the whole core enters nap mode, which allows the hardware to clock off most of the circuits inside the processor core. Reducing active power consumption by turning off the clocks allows the temperature to fall, which further reduces static power leakage of the circuits, resulting in additional power savings.

Capacity-on-demand processor cores are kept in core nap until they are activated and return to core nap whenever they are deactivated.

### 1.3.6 Energy optimized fan control

On supported POWER6 systems, the firmware can adjust fan speed dynamically based on power consumption, altitude, ambient temperature, and power savings mode. Systems are designed to operate in worst case environments, that is in hot ambient temperatures, at high altitudes, and at high utilization. Typically, not all of these conditions are experienced by the system simultaneously.

When no power savings setting is enabled, fan speed is based on ambient temperature and assumes a high-altitude environment. When a power savings setting is activated (either static power savings or dynamic power savings), fan speed varies based on power consumption, ambient temperature, and altitude (if specified). Note that altitude is not sensed automatically but can be set manually for a system in Active Energy Manager (“Setting the altitude for a resource” on page 211). If no value for the altitude is set, the system assumes a default value of 350 meters above sea level.

### 1.3.7 Processor folding

While processor core nap provides substantial energy savings when processors become idle, additional savings can be realized if processors remain idle by intent rather than by default. *Processor folding* is a consolidation technique that adjusts dynamically, over the short-term, the number of processors that are available for dispatch to match the number of processors that are demanded by the workload. As the workload increases, the number of processors made available increases, whereas as the workload decreases, the number of processors made available decreases. Active Energy Manager does not control this feature.

Processor folding increases energy savings during periods of low to moderate workload because unavailable processors remain in low power idle states longer than they otherwise would. Because the idle condition is intentional, the hypervisor can also exploit special purpose idle states that can reduce power consumption even further than with nap mode alone but without the stringent latency requirements of the nap function. Processor folding achieves power savings similar to what could be achieved by intelligent, utilization-based logical partitioning configuration changes, but it does so with much greater efficiency and fidelity, and without impacting the configuration or processor utilization of the partition.

Processor folding can be enabled only with folding-aware service processor firmware (minimum level EM340) and operating system levels as follows:

- AIX 5.3 technology level 9
- AIX 6.1 technology level 2
- IBM i 6.1 with PTF MF45515
- Linux is not supported

### 1.3.8 EnergyScale for I/O

EnergyScale for I/O support is available on all POWER6 processor-based systems and the expansion units that they support. Note that EnergyScale for I/O applies to hot pluggable PCI slots only. Power controls for other types of I/O features and built-in (embedded) PCI adapters are not visible to the firmware, so they cannot be powered off independently of their enclosure power. Active Energy Manager does not control this feature.

IBM Power systems turn off pluggable PCI adapter slots that are not being used automatically. This automatic shut off saves approximately 14 watts per slot. A PCI adapter slot is considered not being used when the slot is empty, when the slot is not assigned to a partition, or when the partition to which the slot is assigned is not turned on.

A PCI slot is turned off immediately by system firmware when it is removed dynamically from the partition to which it was assigned and when the partition to which it is assigned is turned off. It is also powered off immediately by the system firmware if the adapter is physically removed from the slot, although such *hot* removal of PCI adapters is not supported. The user needs to explicitly power off the slot using the I/O adapter concurrent maintenance utilities prior to removing the installed adapter. Furthermore, system firmware scans all pluggable PCI slots automatically at regular intervals looking for slots that meet the criteria for not being in use and then turns off these slots. Thus, when the system is powered on, any slots powered on but not in use are powered off.

The power saving achieved by EnergyScale for I/O is factored into the calculation of the minimum and maximum allowable power cap values.

### 1.3.9 EnergyScale user interfaces

The primary user interface for EnergyScale functions on a POWER6 processor-based system is Active Energy Manager running on IBM Systems Director. For customers who do not have IBM Systems Director and Active Energy Manager, the activation of static power savings mode (only) is also supported from the Advanced System Management Interface (ASMI) or a controlling Hardware Management Console (HMC). Static power savings mode is the only EnergyScale feature supported using ASMI and HMC.

Note that Active Energy Manager is the preferred user interface for setting power savings because it provides far more flexibility and function than ASMI or HMC. For information about how to work with the HMC and ASMI interfaces to turn on static power savings, refer to *Going Green with IBM Systems Director Active Energy Manager 3.1*, REDP-4361.

### 1.3.10 EnergyScale functions by IBM Power system model

Table 1-4 shows the EnergyScale functions for each supported IBM Power system model. Note that the supported EnergyScale functions vary according to the installed level of firmware on the system's service processor (FSP).

> **Note:** Table 1-4 is correct for firmware level EM340.

*Table 1-4 EnergyScale functions by IBM Power system model (firmware level EM340)*

| EnergyScale function | IBM Power blades: JS12 JS22 JS23 JS43 | IBM Power servers: 560/570 < 4 GHz | IBM Power servers: 570 ≥ 4 GHz | IBM Power servers: 520/550 < 4 GHz | IBM Power servers: 520/550 ≥ 4 GHz | IBM Power servers: 575 < 4 GHz | IBM Power servers: 575/595 ≥ 4 GHz |
|---|---|---|---|---|---|---|---|
| Power trending | Yes | No[3] | No[3] | Yes | Yes | Yes | Yes |
| Thermal and CPU trending | Yes[6] | No[3] | No[3] | Yes | Yes | Yes | Yes |
| Static power savings[4, 5] | Yes | No | Yes | No | Yes | No | Yes |
| Dynamic power savings[1] | Yes | No | No | No | Yes | No | No |
| Hard power capping[4] | Yes | No | No | Yes | Yes | No | No |
| Soft power capping[1] | Yes | No | No | No | Yes | No | No |
| Enhanced fan control | No[7] | No | No | No | Yes | No | No |
| Processor nap | Yes | Yes | Yes | Yes | Yes | Yes | Yes |

| EnergyScale function | IBM Power blades: JS12 JS22 JS23 JS43 | IBM Power servers: 560/570 < 4 GHz | IBM Power servers: 570 ≥ 4 GHz | IBM Power servers: 520/550 < 4 GHz | IBM Power servers: 520/550 ≥ 4 GHz | IBM Power servers: 575 < 4 GHz | IBM Power servers: 575/595 ≥ 4 GHz |
|---|---|---|---|---|---|---|---|
| Processor folding[2] | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| EnergyScale for I/O | No[7] | Yes | Yes | Yes | Yes | Yes | Yes |

1 Active Energy Manager is required to enable this function.
2 Refer to 1.3.7, “Processor folding” on page 28 for minimum operating system requirements.
3 An iPDU can be used to provide this support.
4 Can be enabled by Active Energy Manager 3.1 or 4.1.
5 Can be enabled using ASMI or HMC.
6 Chassis level only.
7 Fans and I/O are controlled at a BladeCenter chassis level, not blade level.

## 1.3.11 EnergyScale impacts on processor utilization and accounting

Historically, CPU elapsed time on a server was used to measure both processor utilization (the percentage of processor clock cycles that are being used to perform work) and accounting (processing time for a given job). Every tick of processor time represented the same amount of processing capacity and work performed because the processor operated at constant speed.

IBM POWER6 processor-based systems with EnergyScale technology introduce variable processor speed due to the implementation of the power savings and power capping functions. Because the speed (clocking frequency) of the processor can change, the processing capacity also changes correspondingly. A tick of processor time at reduced frequency (power) performs the same amount of work as a tick of processor time at full power. However, if there are fewer ticks, less work is performed. Therefore, elapsed processor time alone cannot be relied on to measure both processing capacity and work performed.

To address the challenge of measuring both processing capacity and work performed, the IBM POWER6 processor implements a new processor timekeeping facility that can be used to adjust for the relationship between elapsed time and processing capacity in a variable processor frequency (speed) environment.

Prior to the development of POWER6 processors, processor utilization on IBM POWER processors was measured by the *processor utilization resource register* (PURR). Every tick of PURR represented the same processing capacity because there was no capability to change the clocking rate. Therefore, PURR based processor time could be used for both processor utilization and accounting.

The POWER6 processor also includes a *scaled PURR* (SPURR) to measure processor utilization in a variable speed processor environment.

For example:

- When the POWER6 processor is operating at full speed, the PURR and SPURR tick in lockstep.
- When the POWER6 processor is operating at reduced speed, the SPURR ticks slower than the PURR because each processor cycle is longer.
- When the POWER6 processor is operating at half speed, the SPURR ticks at half the rate of PURR.
- When the POWER6 processor is operating in excess of full speed, the SPURR ticks faster than the PURR.

Scaled processor time measures processing capacity relative to the reduced speed of the processor. If the processor is indeed running at full speed, then processor time and scaled processor time are one and the same. The ratio of scaled processor time to processor time is the relative processor speed.

To summarize, processor utilization using SPURR is relative to the reduced processor processing capacity. Therefore, a SPURR enabled operating system accurately accounts for reduced processing capacity when calculating processor time (ticks). With PURR, the operating system does not adjust for the reduced processing capacity and therefore reports more processing time (processor ticks) than actual processor time. In terms of processor utilization, PURR could report 100% processor utilization when the processor is in power savings mode whereas a SPURR enabled operating system would always report < 100% utilization.

While the addition of scaled processor time provides the information necessary to accurately measure processor utilization in a variable processor speed environment, but it does not address the ambiguity of processor time in the historical context. Simply stated, whether a processor time value in an existing environment should represent PURR based processor time or SPURR based processor time is open to interpretation. In some contexts, such as when processor utilization data is to be used to determine how many more jobs should be started to consume idle capacity, PURR based processor time is appropriate. In other contexts, such as when process accounting data is used for billing purposes, SPURR based processor time might be more suitable.

There is no single best answer to handle all cases, and in fact, the issue has not been uniformly addressed by all operating systems:

- In IBM i and Linux, processor time remains PURR based. The result is that IBM i and Linux performance tools that use processor time remain accurate, but they do not take into account processor speed variations. Versions of IBM i and Linux that are EnergyScale-ready provide additional SPURR based scaled processor time in some APIs and performance tools that also provide PURR based processor time.
- In AIX, processor time is SPURR based rather than PURR based, which means that accounting and performance tools are accurate irrespective of processor speed variations. AIX also continues to make PURR based information available through performance tools and APIs.

## 1.4 Introducing x86 energy management techniques

In this section, we provide a brief introduction to the energy management techniques used in x86 architecture processors—System x, BladeCenter, and iDataPlex servers—and how Active Energy Manager can use these techniques.

x86 architecture processors provide two fundamental mechanisms for controlling power consumption:

- P-states or *processor performance states*
- T-states or *processor throttling states*

These mechanisms are the two primary controls used to reduce power consumption on x86 machines. They both employ the same general principle which is that a processor consumes less power when running slower. However, they differ in how they slow the processor down.

- P-states

  The P-states mechanism is used on x86 architecture servers to implement the equivalent of power savings on IBM POWER6 processors. P-states is turned on in the x86 server’s BIOS and is thereafter controlled by the installed operating system. When enabled in the BIOS, the operating system has the option of doing power savings or not. The user decides what happens by setting an operating system-level power policy. If not enabled in the BIOS, the operating system cannot do power savings.

Figure 1-13 shows an example of the BIOS CPU options that can be set on a typical x86 server. To enable P-states (and therefore power savings) on an x86 server, change the Processor Performance States option from **Disabled** to **Enabled**.

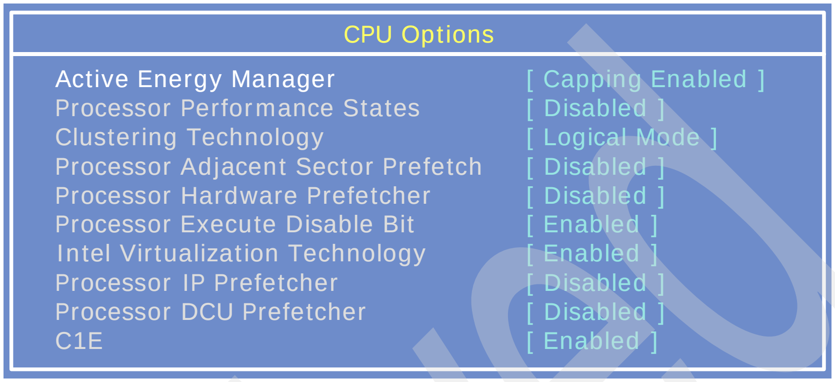


*Figure 1-13 Enabling P-states for power savings on an x86 server*

The P-states mechanism involves running the processor clock at a lower frequency. This results in less power consumption because power consumption is proportional to clock speed. P-states has the additional advantage that a processor running at a lower frequency can operate at a lower voltage which further reduces power requirements. The combination of the two provides a super-linear reduction in power consumption, with a very good power-performance trade-off.

These processor performance states can be changed very quickly in response to processor demand while software continues to execute. This technique allows the operating system to provide automatic scaling of the processor's power consumption in response to varying workloads, with no required user intervention and no perceivable effect on system performance.

> **Note:** Active Energy Manager cannot control power consumption using P-states. However, Active Energy Manager can only report whether or not the operating system is controlling power saving.

The P-states mechanism is an example of a Dynamic Voltage/Frequency Scaling (DVFS) mechanism where both frequency and voltage are reduced. Intel®'s implementation of DVFS is called Demand Based Switching (DBS) whereas AMD calls their implementation PowerNow!.

- T-states

  The T-states mechanism involves running the processor clock at the nominal frequency, but periodically stopping, or gating, the clock. Therefore, this technique is often referred to as *clock gating*. The processor still runs at its nominal speed, but some clock cycles are not made available for executing instructions. The result is an effective decrease in clock speed with a corresponding reduction in power consumption. Clock gating provides a very effective method of reducing power (and therefore temperature), but the resulting decrease in performance is greater than when using P-states.

  Clock gating is generally used in scenarios where processor temperature is too high. In this case, the goal is simply to reduce processor power (and therefore temperature) as quickly as possible to avoiding damaging the processor, without concern about the resulting impact on performance.

  The Active Energy Manager power capping function exploits clock gating in order to enforce a power cap setting on x86 architecture servers. Active Energy Manager only supports *hard* power capping in an x86 environment.

  As shown in Figure 1-13 on page 34, to enable T-states (and, therefore, power capping) on an x86 server, change the Active Energy Manager option from Capping Disabled to Capping Enabled.

  The T-states mechanism is an example of a Dynamic Frequency Scaling (DFS) mechanism where only the frequency is reduced while the voltage is unchanged.

P-states and T-states are discussed in the document *Advanced Configuration and Power Interface Specification*, which available at:

`http://www.acpi.info/DOWNLOADS/ACPIspec30b.pdf`

2

# Planning for Active Energy Manager

In this chapter, we introduce IBM Systems Director and the Active Energy Manager plug-in. We also discuss the terminology and concepts of Active Energy Manager and the new features of Active Energy Manager Version 4.1.1.

We explain the hardware and software requirements, the devices that can be monitored and managed by Active Energy Manager, how to download Active Energy Manager in preparation for installation, and the licensing requirements of Active Energy Manager.

We discuss the following topics in this chapter:


- 2.1, "Introduction to IBM Systems Director" on page 38
- 2.2, "Active Energy Manager product overview" on page 41
- 2.3, "Requirements" on page 53
- 2.4, "Downloading Active Energy Manager" on page 63
- 2.5, "Licensing" on page 64
- 2.6, "Backup and recovery" on page 68


**Tip:** For additional planning information, see the Active Energy Manager 4.1.1 Information Center:

`http://publib.boulder.ibm.com/infocenter/director/v6r1x/index.jsp?topic=/aem_410/frb0_main.html`

 

## 2.1 Introduction to IBM Systems Director

IBM Systems Director is an integrated suite of tools that improves and facilitates the management of physical and virtual systems across an heterogeneous environment including Intel Linux, Power Linux, AIX, Windows®, i5/OS® and System z, running on IBM Hardware, as well as on x86-based non-IBM Hardware.

From a single Web browser interface, you can:

- Monitor the hardware status from different system environments

  IBM Systems Director generates an event if, for example, server redundancy is lost when one of the components fails (power supply, hard disk, CPU, and so forth) or even if a component is about to fail (as indicated by Predictive Failure Analysis[1]).

- Collect inventories from each managed server

  IBM Systems Director can generate a detailed report of each hardware component (FRU, firmware versions, and so forth) and software applications installed on the systems.

- Automate processes

  IBM Systems Director can automate the processes required to enable a proactive environment, upgrading hardware components to a certain firmware level, installing software updates, or reducing energy costs bringing down power consumption through Active Energy Manager.

- Monitor and manage the power usage of system

  IBM Director through Active Energy Manager is able to control power in the data centers. This feature will be described later in this publication.

- Manage virtual environments

  IBM Systems Director help to analyze, visualize and manage different virtual environments including HMC, IVM, Microsoft Virtual Server, VIO, VMware ESX, VMware ESXi, VMware VirtualCenter, Xen, and z/VM®.

- Deploy operating systems images

  Through Tivoli Provisioning Manager for OS Deployments, IBM Director performs unattended operating system installations and creates and deploys operating system images.

---

[1] Predictive Failure Analysis (PFA) is a proprietary IBM technology that can predict the failure of supported hardware components using advanced heuristic techniques and periodic self-diagnostics. IBM includes PFA on selected models for power supplies, processors, memory, hard drives, and fans.

- Monitor the usage of hardware and software resources

  IBM Director enables the use of thresholds to generate an event when resources are reaching a predefined maximum limit (that is CPU, memory, and hard disk usage) or when the status of a process, application, or service changes.

- Simplify server connections and adding failover capabilities

  With BladeCenter Open Fabric Manager, IBM Systems Director provides I/O virtualization of Ethernet MAC and Fibre Channel WW Names to reduce configuration complexity. In addition, it adds capabilities to monitor blade servers for failure events and to take automatic actions to fail over from a faulty blade to a cold standby blade.

Figure 2-1 on page 40 shows the IBM Systems Director server topology. IBM Systems Director architecture has a central IBM Systems Director Server component and agents installed on the managed end-points either the Common Agent and smaller-footprint Platform Agent. Each agent provides a different footprint size, level of performance, and set of management functions. IBM Systems Director can also discover and manage some systems on which neither of these operating-system agents is installed, but the level of management is limited.

IBM Systems Director Server can also be configured to forward alerts to higher-level enterprise managers, including:

- CA Unicenter NSM
- HP OpenView
- Tivoli NetView®
- Tivoli Enterprise
- Microsoft Systems Center Operations Manager
- Microsoft Systems Management Server

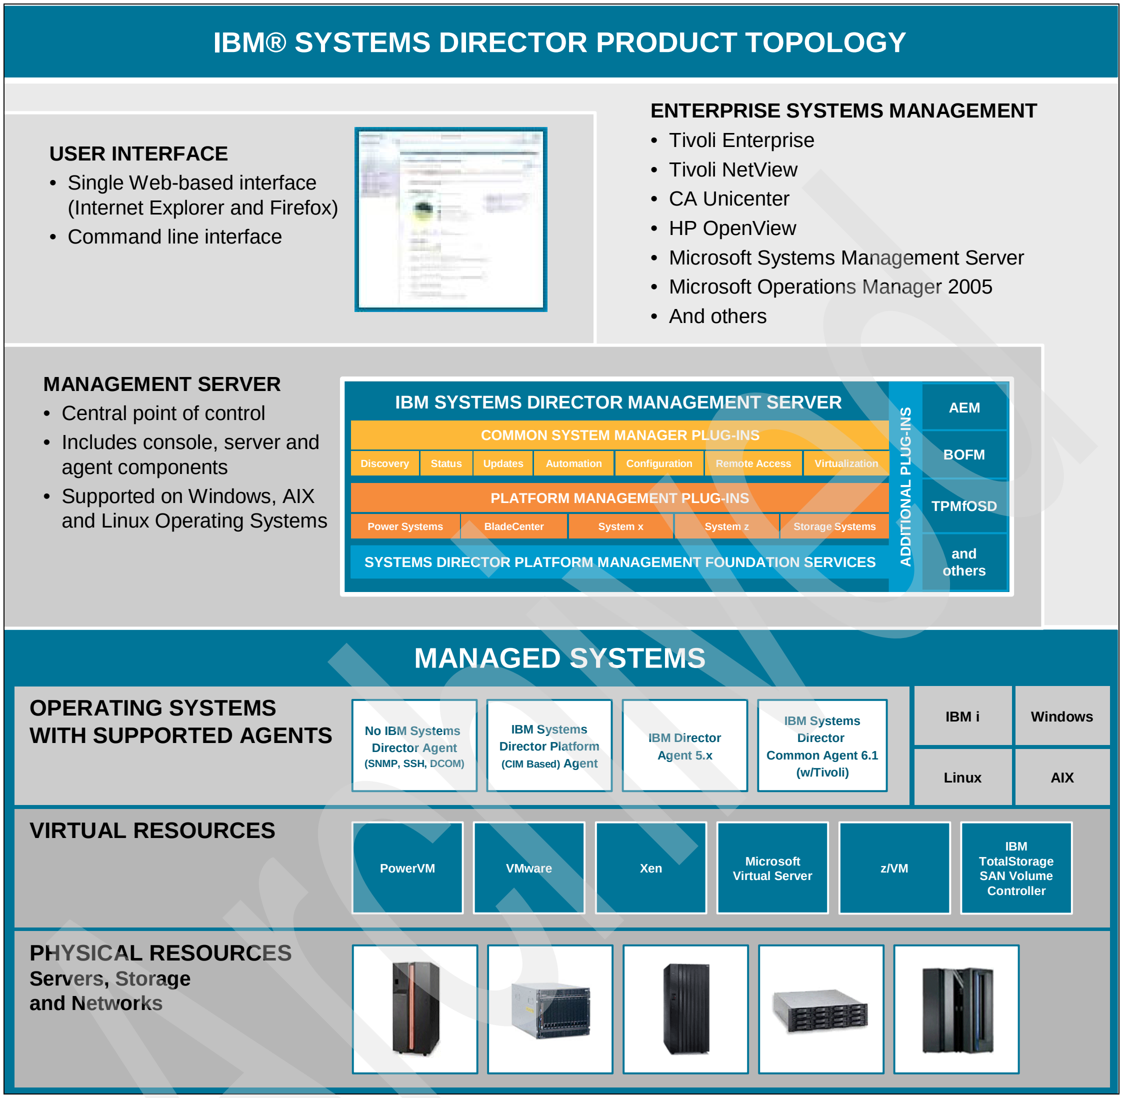


*Figure 2-1 IBM Systems Director 6.1 Topology*

## 2.2 Active Energy Manager product overview

IBM Systems Director Active Energy Manager is an IBM Systems Director plug-in. It provides power management functions on top of the base IBM Systems Director product.

Active Energy Manager monitors and manages the power and cooling needs of IBM servers and BladeCenter systems. Non-IBM systems can also be monitored using metering products, such as intelligent power distribution units (iPDUs or PDU+), sensors, and integration with infrastructure software.

Some tasks that you can perform on Active Energy Manager resources include:

- Monitoring power consumption data
- Collecting power consumption data
- Managing power, including
  - Setting power savings options
  - Setting power caps
  - Automating power-related tasks
- Configuring metering devices, such as PDUs and sensors
- Exporting data
- Viewing events
- Calculating energy cost
- Setting thresholds
- Creating and setting power policies
- Monitoring of power and cooling equipment that affect the IT resources

The first step in making a data center run more efficiently is to understand the power and cooling characteristics of the individual pieces of equipment using real-time monitoring in Active Energy Manager. Using Active Energy Manager, you can take steps to save energy costs.

For example, a data center using the Liebert SiteScan Web application from Emerson Network Power to manage power and cooling equipment can collaborate with Active Energy Manager to collect and view data, resulting in an integrated view of where power is used in the data center and where additional cooling might be required.

Energy management functions are also integrated with IBM Systems Director functions. For example, you use the same interface to set energy-related thresholds as other thresholds that are set in IBM Systems Director. Also, when viewing system properties in the IBM Systems Director resource navigator, you can view Active Energy Manager properties. You can even display a thumbnail view of an energy trending graph within IBM Systems Director to call attention to the most critical systems. In addition, most Active Energy Manager tasks are scriptable using the systems management command-line interface.

Active Energy Manager is part of a larger energy-management implementation that includes hardware and firmware components. For more information, see IBM Project Big Green at:

http://ibm.com/press/us/en/presskit/21440.wss

## 2.2.1 Architecture

This section shows how the different parts of the IBM Systems Director and Active Energy Manager environments relate to each other and describes the various terms that we use throughout this chapter.

Figure 2-2 shows a high-level view of the IBM Systems Director and Active Energy Manager environments and their object management domains.

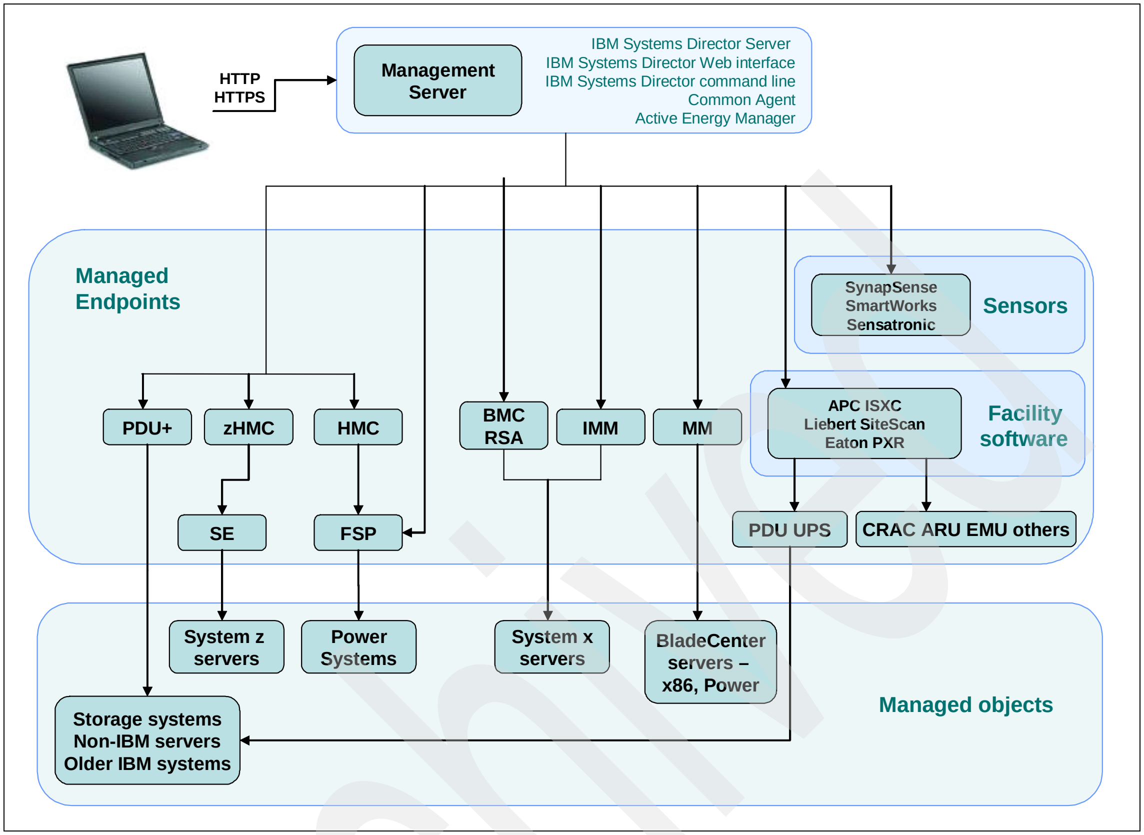


*Figure 2-2 IBM Systems Director and Active Energy Manager object management domains*

Figure 2-2 illustrates the following components:

- *Managed objects* are objects that can be managed by an Active Energy Manager server through the Active Energy Manager endpoints. It comprises Storage, older IBM servers that do not support Active Energy Manager, and non-IBM servers.
- Active Energy Manager integrates *facility software* that are management applications for collecting energy data from all the equipment in a data center such as uninterruptible power supplies, PDUs, Computer Room Air Conditioning (CRACs), Air Removal Units (ARUs), Environmental Monitoring Units (EMUs), and so forth.
- *Sensors* are metering devices that monitor values such as the power use, temperature, humidity, and dew point of other resources.

- Power Systems™ (System p® and System i with POWER processors) can be managed by Active Energy Manager. Active Energy Manager can manage the power of a Power system in one of the following ways:
  - Directly, by communicating with the system's integrated Flexible Service Processor (FSP).

    A POWER6-based System i, System p, or Power 5xx system always has an integrated FSP that provides management capabilities. You normally communicate directly with the FSP to power manage a POWER6 system only when it is not managed by an HMC.

    If the POWER6-based system is being managed by IVM, Active Energy Manager communicates directly to the FSP. AEM does not communicate directly with IVM.
  - Indirectly, by communicating with the system's managing HMC.

    A POWER6-based System i, System p, or Power 5xx system is usually managed by an external HMC that provides additional management capabilities compared with an FSP and that can manage multiple POWER6 systems concurrently. In this case, Active Energy Manager communicates with the FSP on the power managed system through its controlling HMC.
  - Indirectly, by communicating with a PDU+ that has a Power Systems server attached.

    Note, however, that depending on the servers, the PDU+ provides a lower level of power management function than if the Active Energy Manager server were to be connected directly or indirectly to the FSP or HMC.
- Power distribution units: The PDU+ is intended to provide limited Active Energy Manager functionality to attached existing devices, storage, I/O drawers, and other devices that cannot otherwise be managed by Active Energy Manager.
- Active Energy Manager also supports x86 architecture server (System x servers, BladeCenter servers, and iDataPlex servers). The Active Energy Manager interfaces to these systems is as follows:
  - The Active Energy Manager interface to a BladeCenter server is the advanced management module (AMM) installed in the BladeCenter chassis.
  - The Active Energy Manager interface to an IBM System x rack or tower server is either an Remote Supervisor Adapter (RSA II), Baseboard Management Controller (BMC) or Integrated Management Module (IMM).

  The AMM, RSA, and BMC provide an energy management interface to an Active Energy Manager server in a similar way to an FSP in a POWER6-based System p, System i, or Power server.

- The Active Energy Manager server can monitor a System z system by communicating with its managing zHMC. System z is always managed by an external zHMC that provides additional management capabilities to the Support Element (SE). The zHMC can manage multiple System z systems as well as other System z models.

  The Support Element of a System z system provides similar management functions to an FSP, AMM, RSA, or BMC, with the following differences:

  – AEM cannot connect directly to the System z Support Element. It must connect through a zHMC.

  – The Support Element is a separate but physically integrated computer that is tightly coupled with the System z Central Processor Complex (CPC) from a management perspective but is not physically integrated into the frame of the System z such as the FSP, AMM, RSA, and BMC.

## 2.2.2 Terminology

In this section, we define some of the terms that Active Energy Manager uses:

- Label power

  Most items that use electricity have Underwriters' Laboratories (UL) power rating.[2] For servers, this rating is printed on a label on the back of the server and represents the most power that system can ever draw. This label power or nameplate power is based on the capacity of the system's power supplies. It takes into account a fully-configured system with the highest power usage components running the highest possible workload.

- Input power

  Input power defines the power that a system or device is actually using. This is almost always less than the calculated maximum defined by the label power. Knowing the input power of all the equipment in the data center is important when planning for new equipment. Some servers and power distribution units can monitor and report on their input power usage. If that is not available, an electrician can usually provide a device that can measure and record the current flowing into the system. Then, using the formula of 1 watt = 1 volt × 1 amp x PF (power factor), input power can be calculated.

[2] Underwriter's Laboratories ratings are primarily for products made for the US market. However all IBM products that are for a world-wide market have a UL label on them. The exceptions are for systems and devices that are made exclusively for a non-US market. These systems and devices might not have a UL label but will still indicate a maximum power consumption value which has the same meaning.

- Output power

  Output power is a value that is determined from several system sensors. It represents the amount of power system components are consuming from the internal power supplies. It is less than the input power due to losses when the power is converted from AC input to DC output or from other circuitry that filters out input power fluctuations. It can vary based on changes in demand for power by individual components.

- IBM Systems Director plug-in

  A piece of software that extends the functionality of IBM Systems Director. Some of the IBM Systems Director plug-ins are Active Energy Manager, Virtualization Manager, BladeCenter Open Fabric Manager, and Tivoli Provisioning Manager for OS Deployment.

- Intelligent power distribution unit (IBM products have the name PDU+)

  An electrical device that controls power distribution, provides circuit protection, and monitors the power and temperature of the environment. On systems that support power trending, the power data is collected internally from the system itself and does not require any additional hardware to obtain the power usage data. A PDU+ can also be used to collect power usage data from non-server equipment, like storage units, I/O drawers, and computer room air conditioning (CRAC) units.

- Group policy

  A power policy associated with a group of resources, for power capping settings only. A group power capping policy specifies an overall power cap that the systems in the group collectively cannot exceed and can be applied to any number of groups.

- Power policy

  A set of rules that are enforced at every poll and that influence the power management behavior of a resource or group of resources. A system power policy is either a power cap or power savings setting that can be defined and applied to any number of individual systems.

  Both system power policies and group power policies are continually enforced by Active Energy Manager on the systems or groups to which the policies are applied. A power policy is an Active Energy Manager construct. No values or settings are defined on the target resources and groups of resources.

- Metering device

  An object that monitors values such as power use, temperature, humidity, and dew point of other resources. When you associate a resource with a metering device, such as a PDU outlet group or a SynapSense sensor, you are effectively monitoring the power or environment of the resource through the metering device. The metered values are shown as properties of the resource with “externally metered” added to the label. For example, input power metered for a resource through a PDU outlet group is shown as a property of the resource with the label “Average input power (externally metered)”.

- Metered device

  The resource associated with a *metering* device is called a *metered* device. A metered device can be a power-managed resource, a server which does not have built-in power metering support, or any other resource which is defined using the Active Energy Manager Configure Metering Device task.

- Outlet group

  A group of outlets in a PDU device that have power reported as a single logical use of power. An outlet group might have only one outlet in it. Therefore, the outlet group represents the power for a single resource plugged into that outlet.

- Sensor node

  A wireless device that contains, and to which are attached, one or more sensors, such as temperature and power sensors.

- Power Savings mode

  Power Savings mode provides a way to save power by dropping the voltage and frequency a fixed percentage. This percentage is predetermined to be within a safe operating limit and is not user configurable. Under current implementation this is a frequency drop. Power Savings mode can be enabled and disabled through the HMC and FSP interfaces. However, Active Energy Manager is the recommended user interface to enable or disable Power Savings mode.

  One possible use for Power Savings is to enable it when workloads are minimal, such as at night, and then disable it in the morning. When Power Savings is used to reduce the peak energy consumption, it can lower the cost of all power used. At low CPU utilization, the use of Power Savings increases processor utilization such that the workload notices no performance impact.

  The IBM Systems Director scheduler could be used to automate the enabling and disabling of Power Saver mode based on projected workloads. You can also run scripts to enable and disable it based on CPU utilization.

- Power capping

  Power capping enforces a user specified limit on power usage. The user must set and enable a power cap from the Active Energy Manager user interface.

  In most data centers, when a machine is installed, a certain amount of power is allocated to it. Generally, the amount is what is considered to be a “safe” value, often the label power for the system. This often means that there is a large amount of reserved, extra power that is never used. This is called the *margined power*. The main purpose of the power cap is not to save power but rather to allow a data center operator the ability to reallocate power from current systems to new systems by reducing the margin assumed for the existing machines.

Thus the basic assumption of power capping allows an operator the ability to add extra machines to a data center which previously had all the data center power allotted to its current systems.

- Power trending

  Power trending shows the consumption of power by a supported power managed object over time. It is possible to use this data not only to track the actual power consumption of monitored devices, but also to determine the maximum value over time. The data can be presented either graphically or in tabular form.

- CPU trending

  CPU trending determine the actual CPU speed of processors for which either the power saver or power cap function is active. The data can be presented either graphically or in tabular form.

- Thermal trending

  Thermal trending monitors the temperatures of a supported power managed object in real time. This data is used to help avoid situations where overheating can cause damage to computing assets and to study how the thermal signature of various monitored devices varies with power consumption. The data can be presented either graphically or in tabular form.

- IBM POWER6 EnergyScale

  The IBM POWER6 processor-based systems offer the IBM EnergyScale technology, which provides several power management functions such as collecting power usage data, dropping processor voltage and frequency by a predetermined percentage, enforcing a specified power usage limit, reducing power consumption by turning off the clock for the core, and powering off automatically pluggable PCI adapter slots when they are not used. We discuss EnergyScale in 1.3, "IBM EnergyScale" on page 10.

## 2.2.3 New features and enhancements in Active Energy Manager 4.1 and Active Energy Manager 4.1.1

Active Energy Manager V4.1 and V4.1.1 provide a number of new features and enhancements, which we describe in this section.

### New in Active Energy Manager V4.1.1

New features and enhancements in Active Energy Manager V4.1.1 over V4.1 include:

- Configuring metering devices

  The Configure Metering Device page allows you to associate metering devices with multiple resources, thus enabling you to configure more complete and accurate relationships within Active Energy Manager. For example, associating a temperature sensor with multiple servers because those servers are physical near the sensor.

- Additional facilities manager integration

  – Eaton Power Xpert Reporting (PXR) 1.0.x

    Eaton's Power Xpert and Foreseer are popular facility management applications that manage all the power equipment in a data center. Active Energy Manager collects energy data and events from Power Xpert Reporting System for uninterruptible power supplies, PDUs, ePDUs, Remote Power Panels, and Rack Power Modules.

  – APC InfraStruXure Central (ISXC) 5.1

    APC InfraStruXure Central is a popular facility management application that manages the power and cooling equipment in a data center. Active Energy Manager collects energy data from ISXC for uninterruptible power supplies, Rack Managers, PDUs, Rack PDUs, CRACs, Air Removal Units, NetBotz sensor pods, Environmental Monitoring Units, and Automatic Transfer Switches.

- Additional sensor support

  Active Energy Manager can now monitor temperature, humidity, dew point, and power readings from Smart Works Smart-Watt meters and Smart-SenseTH sensors. Active Energy Manager can now monitor temperature, humidity, and dew point readings from iButton and Sensatronics sensors.

- Additional hardware support

  For complete hardware support information, see 2.3, "Requirements" on page 53.

- Performance and memory utilization

  Active Energy Manager V4.1.1 includes many improvements to allow increased overall performance and better memory usage on the IBM Systems Director server, allowing Active Energy Manager to monitor and manage power for more resources.

- Support for the Remote Supervisor Adapter II

  Active Energy Manager V4.1 on IBM Systems Director 6.1 did not support the Remote Supervisor Adapter II on System x servers. However this is corrected with V4.1.1 and the RSA II is now supported as a service processor.

### New in Active Energy Manager V4.1

New features and enhancements in Active Energy Manager V4.1 over V3.x include:

- New Web interface

  Active Energy Manager is now integrated into the Web-based interface of IBM Systems Director 6.1. The result is tighter integration of Active Energy Manager and IBM Systems Director, thus eliminating the need to install the Active Energy Manager console interface separately.

- Power policies

  A system power policy is either a power cap or power savings setting that you can define and apply to any number of individual systems or groups of systems. A group power capping policy specifies an overall power cap that the systems in the group collectively cannot exceed and can be applied to any number of groups. These policies are enforced continually by Active Energy Manager on the systems or groups to which the policies are applied.

- Additional CLI support

  Full support for Active Energy Manager systems management command-line interface (**`smcli`**) commands is now added.

- Altitude setting

  On the latest IBM Power systems, specifying the altitude for an Active Energy Manager resource allows Active Energy Manager to adjust power usage and cooling needs accordingly. Active Energy Manager can make this adjustment for altitude on the latest IBM Power systems only.

- Threshold support
- Emerson-Liebert SiteScan Web integration
- Soft power capping
- Dynamic power savings

## 2.2.4 Performance and scalability considerations

You can configure settings to improve the performance and scalability of Active Energy Manager. The performance of Active Energy Manager, as well as the maximum number of resources that can be managed, depends on several factors:

- IBM Systems Director Server system performance (memory, disk, and processing power)
- Network speed
- Types of managed resources

In addition to these considerations, you can configure the following values to improve the performance and scalability of Active Energy Manager:

- Default metering interval

  The default setting in Active Energy Manager is to collect data from resources once every 5 minutes. The metering interval can be increased or decreased so that data is collected from a given resource at a different interval.

  Consider the following when setting metering intervals:

  – How often does data need to be gathered in order to meet specific goals? If only basic energy data is required, metering every 5 or 10 minutes is probably sufficient.

  – The more frequently resources are metered, the more data that is collected, and the more disk space that is needed to store this data.

  – Depending on the environment in which Active Energy Manager is running (CPU speed, disk speed, network bandwidth, number of resources being metered, and metering interval), the Active Energy Manager server might not be able to keep up metering all of the resources at the set metering intervals. In this case, Active Energy Manager meters the resources as fast as it can. Increasing the metering intervals can decrease the CPU, disk, and network load on the IBM Systems Director server.

  – System z servers contain a cache that holds 1 hour of input power data. The default metering interval can be set as high as 60 minutes without having any gaps in the input power data. For System z servers, the ambient and exhaust temperatures are recorded only once per metering interval.

  – On BladeCenter servers, the minimum metering interval is 10 minutes. The Active Energy Manager metering interval can be set to something less than 10 minutes, but new data is available only for BladeCenter components every 10 minutes.

- Default data refresh interval

  The default setting in Active Energy Manager is to gather new data from the server automatically and to display this data every 1 minute. This default value works well in most cases. However, for viewing table or chart data for a very large number of resources over a long period of time, you might need to change the value to make viewing the data easier.

  For example, when viewing the aggregate data for 1000 resources over a 24-hour period, it might take the Active Energy Manager 30 seconds or more to gather this data from the database and aggregate it. In this case, a 1 minute data refresh interval is not recommended. For situations such as this, setting the Active Energy Manager data refresh interval to a higher value might be the best option.

- Data retention setting

  The default setting in Active Energy Manager is to keep the energy data for 365 days, then delete it. When managing a large number of resources, a large amount of data is gathered and saved, often in the range of GB. Decreasing this setting so that data is saved for fewer days decreases the amount of disk space that is needed.

- Configure the number of IBM Systems Director Server connections allowed on Linux systems

  On Linux, there is a limit on the number of resources to which it can send data, at least on the local network. You set this parameter in the `/proc/sys/net/ipv4/neigh/default/gc_thresh3` file, which controls the growth of the neighbor table.

  You can verify whether the IBM Systems Director Server system is hitting this limit by checking the `/var/log/messages` file for the following message:

  ```
  Neighbor table overflow
  ```

  If these messages are in the log file, then the limit is being reached and Active Energy Manager cannot request data from some resources. This situation might result in failures when adding resources to IBM Systems Director.

  To increase this limit, simply increase the number in the `gc_thresh3` file so that it is larger than the number of resources that are being managed. The default value in this file for some Linux systems is 1024. One option is to increase the value to 8192. The side-effect of increasing this value is that the neighbor hash table increases from the default of 256 KB to 2 MB.

- Compression of historical data

  To save disk space and improve query performance of large intervals of data, Active Energy Manager automatically compresses metered data that is more than 7 days old. Each night at midnight, data is compressed to 1 hour averages for both power and environmental values.

### 2.2.5 Further planning information

The following Web sites can provide additional planning information:

- Active Energy Manager home page:

  http://www.ibm.com/systems/management/director/plugins/actengmgr/
- IBM Systems Director home page:

  http://www.ibm.com/systems/management/director/
- *Implementing IBM Systems Director 6.1*, SG24-7694:

  http://www.redbooks.ibm.com/abstracts/sg247694.html
- Active Energy Manager V4.1.1 Information Center:

  http://publib.boulder.ibm.com/infocenter/director/v6r1x/index.jsp?topic=/aem_410/frb0_main.html
- Active Energy Manager V4.1.1 limitations, known problems and workarounds:

  http://publib.boulder.ibm.com/infocenter/director/v6r1x/index.jsp?topic=/aem_410/frb0_r_tbs_solving_problems.html
- IBM Fix Central:

  http://www.ibm.com/support/fixcentral/
- Migrating from Active Energy Manager 3.1:

  http://publib.boulder.ibm.com/infocenter/director/v6r1x/index.jsp?topic=/aem_410/frb0_t_migrating.html
- IBM EnergyScale for POWER6 Processor-Based Systems:

  ftp://ftp.software.ibm.com/common/ssi/sa/wh/n/pow03002usen/POW03002USEN.PDF

## 2.3 Requirements

This section provides information about Active Energy Manager product requirements.

### 2.3.1 Hardware requirements for the management server

The following guidelines for minimum hardware requirements were defined after laboratory testing of Active Energy Manager optimized for DB2® database under simulated stress conditions. Because it is not possible to simulate all configurations under all conditions, these must be accepted as guide, rather than absolute fact.

The following guidelines apply to 32-bit Windows, 32-bit Linux, and AIX environments. In 64-bit environments, more memory will be required to support the application, but other requirements will be similar.

> **Note:** Active Energy Manager V4.1.1. requires IBM Systems Director 6.1.1.

The following requirements are in addition to any requirements for other functions of IBM Systems Director 6.1:

- Installation requirements

  Active Energy Manager will install on any system that currently supports the installation of IBM Systems Director 6.1.1. See the IBM Director 6.1 Information Center for the list of supported systems and products:

  http://publib.boulder.ibm.com/infocenter/director/v6r1x/index.jsp?topic=/director.plan_6.1/fqm0_r_os_supported_by_ibm_director_61.html

  If large numbers of power-managed objects are involved, larger storage capacity is required, which might affect the location where IBM Systems Director and Active Energy Manager are installed.

- Storage requirements

  By default, Active Energy Manager stores 365 days of power-related information in its database. After data is collected for a 7 day period, Active Energy Manager compresses historical data older than 1 week into 1 hour average intervals.

  If you are planning to manage more than 1000 resources, it is recommended that DB2, IBM Systems Director, and Active Energy Manager be on a logical disk that is striped (RAID) over multiple drives to allow I/O spreading. The space needed depends on the number of objects that are monitored.

  Note that a single blade center can have over 20 objects, including system blades, power supplies, and network interfaces. In addition, a single HMC can have multiple systems under its control, and a single power distribution unit can have multiple components plugged in to it.

  The space needed also depends on the metering rate. Active Energy Manager is shipped with a default rate of 5 minutes, but you can adjust this value.

  Assuming the default of 5-minute metering interval, the 31-day space is approximately 5 GB for 1000 blade managed resources and PDUs. As with processing requirements, storage is exercised most heavily by the regular metering of power information and by the processing of queries to refresh console screens. Thus, using multiple consoles simultaneously might increase the need to have multiple disks in a striped array to support the data.

In general, using one or two consoles will allow smooth performance if the data is spread across one disk per thousand managed objects.

- Processor requirements

  Active Energy Manager has the following minimum processor requirements:

  - System x, BladeCenter, and iDataPlex (x86):
    - 500 or fewer resources: Single processor 3 GHz or faster
    - Up to 1000 resources: Single processor with Hyper-Threading, faster than 3 GHz
    - Up to 2000 resources: Dual-core or Quad-core processor
  - Power
    - An rPerf rating of 3 or higher should be reserved to run Active Energy Manager for moderate numbers (500 or less) of managed resources with DB2 database.
    - An rPerf rating of at least 4.5 is recommended for environments up to 1000 managed resources with DB2 database.
    - When scaling to larger numbers of managed objects, add 1.5 to the rPerf requirement for every 1000 managed resources with DB2 database.

  Processors are exercised during the normal metering function, while refreshing the console to show trend information and when making power management decisions. Of these, the first two are the most significant, which means that initiating multiple consoles with simultaneous queries of the database can affect the amount of processing capacity that is needed to maintain smooth performance.

- Memory requirements

  As with processor and disk, the amount of memory needed for Active Energy Manager depends on the number of managed resources and the frequency with which objects are metered and power histories are queried. If other IBM Systems Director functions are not stressed, a moderate configuration of 500 objects can be monitored and managed within the 2 GB of memory that is recommended for IBM Systems Director. If, however, both IBM Systems Director and Active Energy Manager functions are exercised on the same system, it is likely that more memory will be needed to allow an optimal amount of Active Energy Manager data to remain memory-resident. In configurations supporting well over 1000 managed resources, with a limit of 2000, it is recommended to have 4 GB of memory in the system.

- Database requirements

  Active Energy Manager supports 1000 managed resources of a single type and up to 2000 managed resources of multiple types combined with DB2 database. DB2 is the recommended database for more than 500 managed resources with Active Energy Manager.

### 2.3.2 Hardware requirements for managed systems

Active Energy Manager can monitor and manage power consumption for rack servers, BladeCenter chassis and blade servers, System z servers, and iDataPlex servers.

Table 2-1 lists the supported servers and the minimum firmware levels needed. It also lists whether each system supports the overall functions of Active Energy Manager:

- Power monitoring
- Power capping
- Power savings

> **Note:** For the latest information about supported Managed Systems, see the Active Energy Manager Information Center:
>
> `http://publib.boulder.ibm.com/infocenter/director/v6r1x/index.jsp?topic=/aem_410/frb0_r_HW_reqs_managed_systems.html`

*Table 2-1 Active Energy Managed system support for IBM Servers*

| Server | Firmware level[a] | | Power monitoring | Power capping | Power savings |
|---|---|---|---|---|---|
| **System x Servers** | **BIOS** | **BMC** | | | |
| System x3350 (4192) | K3E123A | K3BT10A | Yes | No | Yes[b] |
| System x3400 M2 (7836, 7837) | Latest level | Latest level | Yes | Yes | Yes[b] |
| System x3500 M2 (7839) | Latest level | Latest level | Yes | Yes | Yes[b] |
| System x3550 (7978, 2805) | GFE127A | GFBT39A | Yes | Yes | Yes[b] |
| System x3550 M2 (7946, 4198) | Latest level | Latest level | Yes | Yes | Yes[b] |
| System x3650 (7979, 4388) | GGE127A | GGBT38A | Yes | Yes | Yes[b] |
| System x3655 (7943) | C9E110A | C9BT09A | Yes | Yes | Yes[b] |
| System x3650 M2 (7947, 4199) | Latest level | Latest level | Yes | Yes | Yes[b] |
| System x3655 (7985) | C2E128A | C2BT31A | Yes | Yes | Yes[b] |

| Server | Firmware level[a] | | Power monitoring | Power capping | Power savings |
|---|---|---|---|---|---|
| System x3755 (7163) | C8E109A | C8BT11A | Yes | Yes | Yes[b] |
| System x3755 (8877) | ZYE129A | ZYBT39A | Yes | No | Yes[b] |
| System x3850 (8864) | ZSE121B | ZSBT20A | Yes | No | No |
| System x3950 (8878) | ZSE121B | ZSBT20A | Yes | No | No |
| System x3950 E (8879) | ZSE121B | ZSBT20A | Yes | No | No |
| System x3850 M2 (7141, 7144, 7233, 7234) | A3E111T | A3BT18A | Yes | Yes[c] | Yes[b] |
| System x3950 M2 (7141, 7144, 7233, 7234) | A3E111T | A3BT18A | Yes | Yes[c] | Yes[b] |
| iDataPlex dx360 M2 (7321, 7323) | Latest level | Latest level | Yes | Yes | Yes[b] |
| BladeCenter HS12 (8014, 8028) | N1E125A | N1BT07L | Yes | Yes | Yes[b] |
| BladeCenter HS20 (8843, 7981) | BWE128A | BWBT35A | Yes | Yes | No |
| BladeCenter HS21 (8853) | BCE121A | BCBT36A | Yes | Yes | Yes[b] |
| BladeCenter HS21 XM (7995) | MJE112A | MJBT13A | Yes | Yes | Yes[b] |
| BladeCenter HS22 (7870, 1936) | Latest level | Latest level | Yes | Yes | Yes[b] |
| BladeCenter HC10 (7996) | DOE122A | DOBT20A | Yes | No | Yes[b] |
| BladeCenter JS21 (8844) | 01EA330_031_031 | | Yes | Yes | Yes |
| BladeCenter JS22 (7998) | 01EA320_030_030 | | Yes | Yes | Yes |
| BladeCenter JS23 (7778) | Latest level | Latest level | Yes | Yes | Yes |
| BladeCenter JS43 (7778-23X w/ FC 8446) | Latest level | Latest level | Yes | Yes | Yes |
| BladeCenter LS20 (8850) | BKE126A | BKBT28A | Yes | Yes | Yes[b] |
| BladeCenter LS21 (7971) | BAE126A | BABT36A | Yes | Yes | Yes[b] |
| BladeCenter LS22 (7901) | L8E123A | L8BT11A | Yes | Yes | Yes[b] |
| BladeCenter LS41 (7972) | BAE126A | BABT36A | Yes | Yes | Yes[b] |
| BladeCenter LS42 (7902) | L8E123A | L8BT11A | Yes | Yes | Yes[b] |
| BladeCenter QS21 (0792) | QB-1.9.1-3 | BNBT17A | Yes | No | No |
| BladeCenter QS22 (0793) | QD-1.26.0-A | BLBT11A | Yes | No | No |
| **IBM Power Systems** | **HMC** | **FSP** | | | |
| IBM Power 520(8203-E4A) | R320.0 | EM320 | Yes | Yes[d] | Yes[e] |
| IBM Power 520 (9407-M15) | R320.0 | EM320 | Yes | Yes[d] | Yes[e] |

| Server | Firmware level[a] | | Power monitoring | Power capping | Power savings |
|---|---|---|---|---|---|
| IBM Power 520 (9408-M25) | R320.0 | EM320 | Yes | Yes[d] | Yes[e] |
| IBM Power 550 (8204-E8A) | R320.0 | EM320 | Yes | Yes[d] | Yes[e] |
| IBM Power 550 (9409-M50) | R320.0 | EM320 | Yes | Yes[d] | Yes[e] |
| IBM Power 560 (9116-561) | R320.0 | EM320 | No | No | Yes[f] |
| IBM Power 570 (9406-MMA) | R320.0 | EM320 | No | No | Yes[f] |
| IBM Power 570 (9117-MMA) | R320.0 | EM320 | No | No | Yes[f] |
| IBM Power 575 (9125-F2A Feature code 6316 air-cooled) | R330.0 | EM330 | Yes | No | No |
| IBM Power 575 (9125-F2A Feature code 7298 water-cooled) | R330.0 | EM330 | Yes | No | Yes |
| IBM Power 595 (9119-FHA) | R330.0 | EM330 | Yes | No | Yes |
| **System z mainframes** | | | | | |
| System z10 BC (2098) | | | Yes | No | No |
| System z10 EC (2097) | | | Yes | No | No |
| System z10 E12 (2097) | | | Yes | No | No |
| System z10 E26 (2097) | | | Yes | No | No |
| System z10 E40 (2097) | | | Yes | No | No |
| System z10 E56 (2097) | | | Yes | No | No |
| System z10 E64 (2097) | | | Yes | No | No |
| **Other servers** | | | | | |
| Other servers and devices | | | Yes[g] | | |

a. The supported firmware levels listed for each server/machine type are those that have been tested with Active Energy Manager. The listed firmware levels and later levels are supported.
b. On Intel and AMD processors, power saving is enabled in the BIOS and controlled by the operating system. For power savings to work on x86 architecture servers, the operating system must also support it. Active Energy Manager can detect whether power savings mode is on or off on supported systems, but cannot control its operation.
c. Supports power capping in single-node configurations only. Multi-node configurations are not supported.
d. Supports soft power capping.
e. Supports dynamic power savings.
f. Supports static power savings only
g. Power monitoring is supported when the server or device is associated with a metering device.

## 2.3.3 Hardware requirements for BladeCenter chassis

Table 2-2 lists the minimum firmware levels for Advanced Management Modules in each BladeCenter chassis.

*Table 2-2 Active Energy Manager managed system support for BladeCenter chassis*

| **BladeCenter chassis** | **Machine type** | **Advanced Management Module firmware level**[a] |
|---|---|---|
| BladeCenter E | 8677 (AC powered) | BPET28G |
| BladeCenter H | 7989 (AC powered) | BPET28G |
| BladeCenter H | 8852 (AC powered) | BPET28G |
| BladeCenter HT | 8740 (DC powered) | BPET28G |
| BladeCenter HT | 8750 (AC powered) | BPET28G |
| BladeCenter S | 8886 (AC powered) | BPET34B |
| BladeCenter T | 8720 (DC powered) | BBET28G |
| BladeCenter T | 8730 (AC powered) | BBET28G |

a. The supported firmware levels listed for each chassis/machine type are those that have been tested with Active Energy Manager. The listed firmware levels and later levels are supported.

## 2.3.4 Metering devices

Active Energy Manager can monitor many metering devices and hardware from external vendors. Specifically, it can monitor IBM PDU+s, and non-IBM PDUs, as well as Emerson-Liebert power units, SynapSoft environments from SynapSense, Smart Works monitored power units, Eaton Power Xpert Reporting System monitored power units, and many power and environmental devices monitored by APC ISXC, Sensatronics sensors, and iButton and 1-wire sensors.

Active Energy Manager can monitor the following metering devices:

- IBM PDU+s
  - IBM DPI® C13 PDU+ (IBM part numbers 39M2816, 39M2818)
  - IBM DPI C13 3-phase PDU+ (IBM part numbers 39M2817, 39M2819, 43V6045, 43V5994, 44V3897)
  - v IBM Ultra Density Enterprise PDU C19 PDU+ (part number: 71762MX | model: 43V5967)
  - IBM Ultra Density Enterprise PDU C19 3 phase 60A PDU+ (part number: 71763MU | model: 43V5968)

- SynapSense sensor nodes

  Active Energy Manager can monitor SynapSense sensor nodes that are connected to a SynapSense SNMP agent that is discovered by Active Energy Manager. Each sensor node can contain one or more current, power, humidity, temperature, or battery voltage sensors. SynapSense sensor networks at the following versions are supported:

  – SynapSoft version 3.0.4 or later
  – SynapSoft version 4.x

  > **Note:** SynapSoft environments at version 2.x and 3.0.3 and earlier might also work; however they are not supported officially in Active Energy Manager and can prevent SynapSense SNMP agents from being discovered. In these cases, upgrade to version 3.0.4 or later.

- Non-IBM PDUs

  Active Energy Manager can monitor the following Smart Works devices:

  – Smart-Watt
  – Smart-Sense sensors
  – Smart-SenseTH sensors

  Active Energy Manager can monitor the following Sensatronics hubs:

  – E16
  – U16
  – E4
  – U4
  – EM1
  – Senturion

  > **Note:** All Sensatronics sensors supported by the these hubs are supported.

  Active Energy Manager can monitor the following iButton and 1-wire hubs:

  – HA7
  – LinkHubE

Active Energy Manager can monitor these PDUs:

- Eaton Monitored Powerware PDUs
  - PW102SW0U150
  - PW102SW0U151
  - PW103SW0U152
  - PW103SW0U153
  - PW105SW0U154
  - PW306SW0U155
  - PW306SW0U156
- Eaton Switched Powerware PDUs
  - PW105MI0U096
  - PW105MI0U097
  - PW105MI0U098
  - PW105MI0U099
  - PW306MI0U113
  - PW309MI0U114
  - PW309MI0U115
  - PW110MI0U116
  - PW110MI0U117
  - PW110MI0U118
- Eaton Remote Power Panels
  - Models with a Power Xpert Gateway Series 1000 card
- Emerson Network Power power units (using Liebert SiteScan Web)
  - Breaker Cabinet Power Monitoring (FPC)
  - Power Monitoring Panel (Ext. Protocol) (PM2)
  - Power Monitoring Panel (PMP)
  - Static Transfer Switch PDU Dual Output (STS)
  - Static Transfer Switch PDU Dual Output (STS-2)
  - Static Transfer Switch PDU (EDS)
  - Voltage-Current-Frequency Monitor Panel (VCF)
  - Voltage-Current Monitoring Panel (VCM)
- Raritan
  - Dominion PX PDUs

- Uninterruptible power supplies
  - Eaton models with a ConnectUPS-X, ConnectUPS-BD or ConnectUPS-E Web/SNMP UPS Connectivity Device
  - Emerson Liebert uninterruptible power supply devices (using Liebert SiteScan Web)
    - Multi Module UPS – S600/610 Extended Prot. – MM4
    - Multi Module Series – MMS
    - System Control Cabinet – S600600/610 Ext. Prot. – SC4
    - System Control Cabinet – SCC
    - Single Module Series – SMS
    - Single Module UPS – S600600/610 Ext. Prot. – SM4
    - Multi-Module SICE 7200 and HiPulse UPS – SMM
    - Systems Cabinet SICE 7200 UPS – SSC
    - Single Module SICE 7200 and HiPulse UPS – SSM
    - UPStation S3 – US3
    - Single Module Series AP301/302 – SM3
    - Single Module UPS – NPower—IMP
    - Single Module UPS – NX—PNX
    - HiNet – PHN
    - NFinity – PNF
    - GXT – PGX
    - PSI – PPS
- Computer room air conditioning (CRAC) units

  Emerson Liebert hardware monitored and controlled by Emerson Liebert SiteScan Web Version 3.0
- Eaton Power Xpert Reporting System monitored power units

  Active Energy Manager can monitor PDUs and uninterruptible power supplies that are defined in an Eaton Power Xpert Reporting Hierarchy database server that has been discovered by Active Energy Manager.
- APC InfraStruXure Central metering devices

  Active Energy Manager can monitor APC CRACs, ARUs, PDUs, Rack PDUs, uninterruptible power supplies, Rack Managers, NetBotz sensor pods, energy management units, and automatic transfer switches that are being managed by an InfraStruXure Central server that has been discovered by Active Energy Manager. These metering devices can contain one or more current, power, humidity, temperature, or dew point sensors.

### 2.3.5 Supported operating systems

Active Energy Manager is supported for use on the AIX, Linux, and Windows operating systems supported by IBM Systems Director 6.1.1 management servers. See the IBM Systems Director 6.1.1 Information Center for details:

http://publib.boulder.ibm.com/infocenter/director/v6r1x/index.jsp?topic=/director.plan_6.1/fqm0_r_os_supported_by_ibm_director_61.html

Active Energy Manager 4.1.1 requires IBM Director 6.1.1. Active Energy Manager 4.1 requires IBM Director 6.1.

## 2.4 Downloading Active Energy Manager

To download Active Energy Manager from the IBM Web site, complete the following steps:

1. Go to the Active Energy Manager home page:

   http://www.ibm.com/systems/management/director/plugins/actengmgr/

2. In the navigation area to the left, click **Downloads**.
3. On the Downloads page, select **IBM Systems Director Active Energy Manager v4.1.1**.
4. Click Submit.

   > **Note:** You are prompted for your IBM sign in if you have not already signed in.

5. Select the files that you want to download. Table 2-3 lists the download package files and their associated installation files.

*Table 2-3 Active Energy Manager download packages and installation files*

| Operating system | Download package | Installation package file name |
|---|---|---|
| AIX | SysDir_AEM4_1_1_AIX.tar.gz | IBMSystemsDirector-AEM-Setup.sh |
| Linux for System p | SysDir_AEM4_1_1_Linux_Power.tar.gz | IBMSystemsDirector-AEM-Setup.sh |
| Linux for System x | SysDir_AEM4_1_1_Linux_x86.tar.gz | IBMSystemsDirector-AEM-Setup.sh |
| Linux for System z | SysDir_AEM4_1_1_Linux_System_z.tar.gz | IBMSystemsDirector-AEM-Setup.sh |
| Windows | SysDir_AEM4_1_1_Windows.zip | IBMSystemsDirector-AEM-Setup.exe |

6. Copy the downloaded file to a local drive on each IBM Systems Director management server on which you want to install Active Energy Manager. Be sure that you copy the correct file based on the operating system that is running on the IBM Systems Director management server.

## 2.5 Licensing

Active Energy Manager performs power *monitoring* functions and, with the addition of a license, it also performs power *management* functions on supported systems.

The no-charge power monitoring functions are as follows:

- Power trending
- Thermal trending
- CPU trending

The licensed power management functions are as follows:

- Power savings
- Power capping

> **Note:** Because the power savings and power capping functions are only applicable to x86 and Power systems, you do not need an Active Energy Manager license for any other IBM platforms (such as System z) because they only support the no-charge monitoring functions.

Some key points to licensing:

- Active Energy Manager licenses are *in addition* to any IBM Systems Director licenses required.
- When we refer to *Active Energy Manager licensing*, we are describing licenses for the power-managed systems. The Active Energy Manager Server itself does not require a license; however, you need to install a license key on the Active Energy Manager Server, depending on how many systems you want to run the power savings or power capping functions. If all you want to do is monitor your systems, you do not need any Active Energy Manager licenses.
- As with IBM Systems Director 6.1.1, there is no Active Energy Manager 4.1.1 client software product as such. The client code that Active Energy Manager Server communicates with is built into the service processor of the IBM power managed systems.

When you download, install, and begin using Active Energy Manager, you are granted a 60-day evaluation license which enables use of the optional management function. When the evaluation license expires, you can continue monitoring energy, but you must purchase a license to continue using the optional management function.

The 60-day evaluation period begins the first time you begin using Active Energy Manager. The number of days left on the evaluation license displays in the License section at the bottom of the Active Energy Manager summary page. The day the evaluation license expires also displays, as well as information about obtaining a license.

The Active Energy Manager license is packaged on a CD-ROM with authorization key and installation program. After you install the license, the optional management functions are enabled and will function just as they did during the evaluation period, with your configurations and settings remaining intact.

The license CD-ROM only contains the license installer and assumes you already have IBM Systems Director and Active Energy Manager installed as described in Chapter 3, “Installing Active Energy Manager” on page 71.

## 2.5.1 Obtaining an Active Energy Manager license

To obtain the correct type and number of Active Energy Manager licenses that you need, follow these steps:

1. Establish which IBM systems require an Active Energy Manager license.
2. Categorize each identified IBM system as to its size.
3. Order Active Energy Manager licenses and software subscription.

We cover each of these steps in the following sections.

### Establish which IBM systems require an Active Energy Manager license

As previously mentioned, the only Active Energy Manager functions that you need a license for are:

- Power savings
- Power capping

Not every IBM System x and Power system supports the power savings or power capping functions. Systems that support power savings and power capping are listed in 2.3.2, “Hardware requirements for managed systems” on page 56 which is accurate at the time of writing. Note that you need to add other systems to this list over time.

Compare the list of IBM systems that you want to enable for power savings or power capping with Table 2-1 on page 56. Note the details of those systems that actually support the power savings or power capping functions.

After you establish which IBM systems support either the power savings or power capping functions, you then need to categorize each server in terms of its size.

### Categorize each identified IBM system as to its size

The chargeable management functions have a system size-specific price structure that is based upon the type of server that is managed. Therefore, after you have determined which of your systems actually require an Active Energy Manager license, you need to categorize each of these systems by size.

Here is the definition of the server categories:

- Small server
  - IBM System x models
  - IBM BladeCenter x86 and Power Architecture® technology-based blades
  - IBM Power Systems, System p models in Processor Groups C5, D5, and E5
  - IBM Power Systems, System i models in Processor Groups P05, P10, and P20
- Medium server
  - IBM Power Systems, System p models in Processor Group F5
  - IBM Power Systems, System i models in Processor Groups P30 and P40
- Large server
  - IBM Power Systems, System p models in Processor Groups G5 and H5
  - IBM Power Systems, System i models in Processor Groups P50 and P60
  - IBM System z models

These models are currently available in the associated processor groups as at the time of writing but might change over time. Check the Active Energy Manager announcement letters for updates to this list.

## Order Active Energy Manager licenses and software subscription

After you have determined the number and size of the IBM systems for which you want to order Active Energy Manager licenses, you need to obtain the program and feature numbers for the licenses and software subscription so that they can be ordered.

Software subscription provides support and program upgrades for Active Energy Manager.

> **Note:** Although one year of software subscription for Active Energy Manager and associated programs is included with the license for the product, you must place a separate order for software subscription registration at the time of initial order or supply of the product.

After one year, you must renew the software subscription. Acquiring program upgrades for eligible offerings under software subscription is only available under a current and active software subscription contract.

Look under *Ordering information* in the following announcement letters to determine the program and feature numbers for both the Active Energy Manager licenses and software subscription:

- x86 architecture systems

  http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?infotype=dd&subtype=ca&&htmlfid=897/ENUS208-332

  To place an order for Active Energy Manager licenses and software subscription for x86 architecture systems, see

  http://www.ibm.com/systems/management/director/plugins/actengmgr/

- Power systems

  http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?infotype=dd&subtype=ca&&htmlfid=897/ENUS208-284#h2-ordinfx

  To place an order for Active Energy Manager licenses and software subscription for Power systems, contact your IBM representative.

- Linux on System z systems

  http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?infotype=dd&subtype=ca&&htmlfid=897/ENUS208-344#h2-ordinfx

  To place an order for Active Energy Manager licenses for Linux on System z systems and software subscription, contact your IBM representative.

# 2.6 Backup and recovery

To preserve IBM Director Server data in case of any kind of failure, a backup is recommended. In this section, we describe the commands that you can use to backup and restore the IBM Systems Director management server environment.

## 2.6.1 The smsave command

The **smsave** command backs up all data that is associated with IBM Systems Director Server, including any file systems data (called the *master data*) and database data. The command also saves all AEM metering collection data. The command syntax is as follows:

```
smsave [-dbUserName user_name] [-dbUserPwd password] [-targetDir
directory] [-dbTargetDir directory] [-noPrompt {true | false}]
```

In this command:

- `-dbUserName` and `-dbUserPwd` are used to specify access to the database. This user must have privileges to back up the database.
- The `-targetDir` and `-dbTargetDir` options determine where the data is stored. The data is saved to `<install_root>\backups\time_stamp` and the database data set is stored with the master set unless otherwise specified.

> **Note:** The `-dbUserName` and `-dbUserPwd` options do not apply to the Apache Derby database.

If your management server is running on AIX ensure that the Bash shell is installed. The database data is stored in a format specific to the type. You cannot restore data from one database type to another (that is, from an Apache Derby database to a DB2 database).

If you have a remote database, the **smsave** command produces two data sets, one on the remote database server and another on the IBM Systems Director Server. These data sets are mated sets. You must maintain and restore these two data sets together. The IP address of the management server and the database server can change from the original installation without any adverse effects.

> **Note:** You must stop the IBM Systems Director Server before using the **smsave** command to backup data.

## 2.6.2 The smrestore command

The **smrestore** command restores a backup created by **smsave**. The command syntax is as follows:

```
smrestore [-dbUserName user_name] [-dbUserPwd password]
{-sourceDirdirectory} {-dbsourceDir directory} {-timestamp timestamp}
[-noPrompt {true | false}]
```

This command uses the following options:

- `-dbUserName` and `-dbUserPwd` specify access to the database. This user must have privileges to back up the database.
- `-sourceDir` and `-dbSourceDir` point to the data previously backed up with **smsave**.
- If you have more than one backup in the source directory you must specify the `-timestamp` option. The time stamp is in the format of *yyyymmddhhss*, where *yyyy* is the year, `mm` is the month, *dd* is the day, `hh` is the hour, and *ss* is the seconds.

3

# Installing Active Energy Manager

This chapter provides step-by-step instructions for installing IBM Systems Director Active Energy Manager 4.1.1 on Windows, Linux, and AIX platforms. In it, we discuss the following topics:

## 3.1 Updating IBM Systems Director

To install IBM Active Energy Manager 4.1.1, you must ensure IBM Systems Director is at Version 6.1.1 or later. In this section, we describe how to update IBM Systems Director from V6.1 to V6.1.1 using Update Manager.

You can update IBM Systems Director with Update Manager using either of the following two methods:

- Connect to the internet using HTTP (port 80) and download the updates.

  The IBM Systems Director server must have direct access to the Internet or must be able to connect through a proxy. If your proxy server requires user authentication, you can specify that authentication information in the **Settings** menu (shown in Figure 3-1) to provide an Internet connection settings. In this menu, you can also set the size limit on the storage location, use a NIM Server for AIX, and other settings.

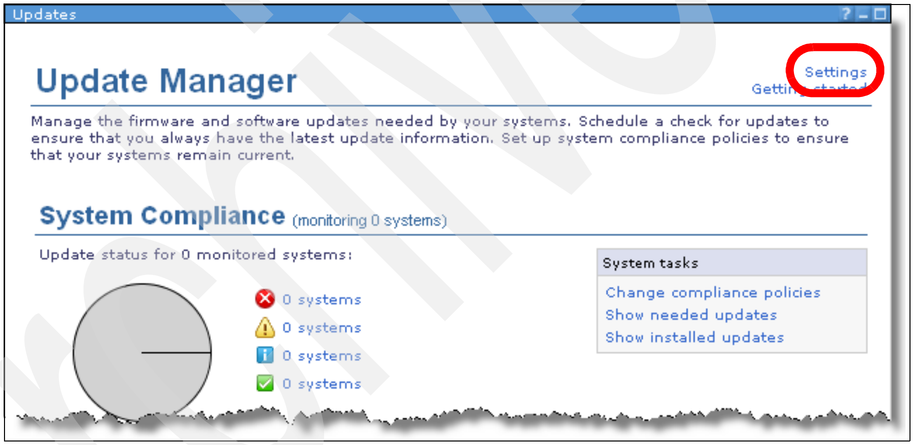


*Figure 3-1 Update Manager settings*

- Download the updates from the IBM Web site and import them using IBM Systems Director Console.

If your IBM Systems Director server has access to the Internet, we recommend that you use the first method.

This section assumes that IBM Systems Director V6.1 is installed. If you need details about how to install IBM Systems Director V6.1 refer to *Implementing IBM Systems Director 6.1*, SG24-7694, which is available at:

http://www.redbooks.ibm.com/abstracts/sg247694.html

**Notes:**

- You need to restart IBM Systems Director server after the update process is complete.
- We recommend you first remove IBM Systems Director updates from the repository base before you apply the new IBM Systems Director updates. You can do it through the Web browser or through the command-line interface. The command to empty the repository base is:

```
install_root/bin/smcli cleanupd -mFv -P “Platform=’Director’”
```

For information about updating IBM Systems Director, go to:

http://publib.boulder.ibm.com/infocenter/director/v6r1x/index.jsp?topic=/director.updates_6.1/fqm0_t_um_updating_director_features.html

The updates that you need to apply should be similar to those shown in Figure 3-2.

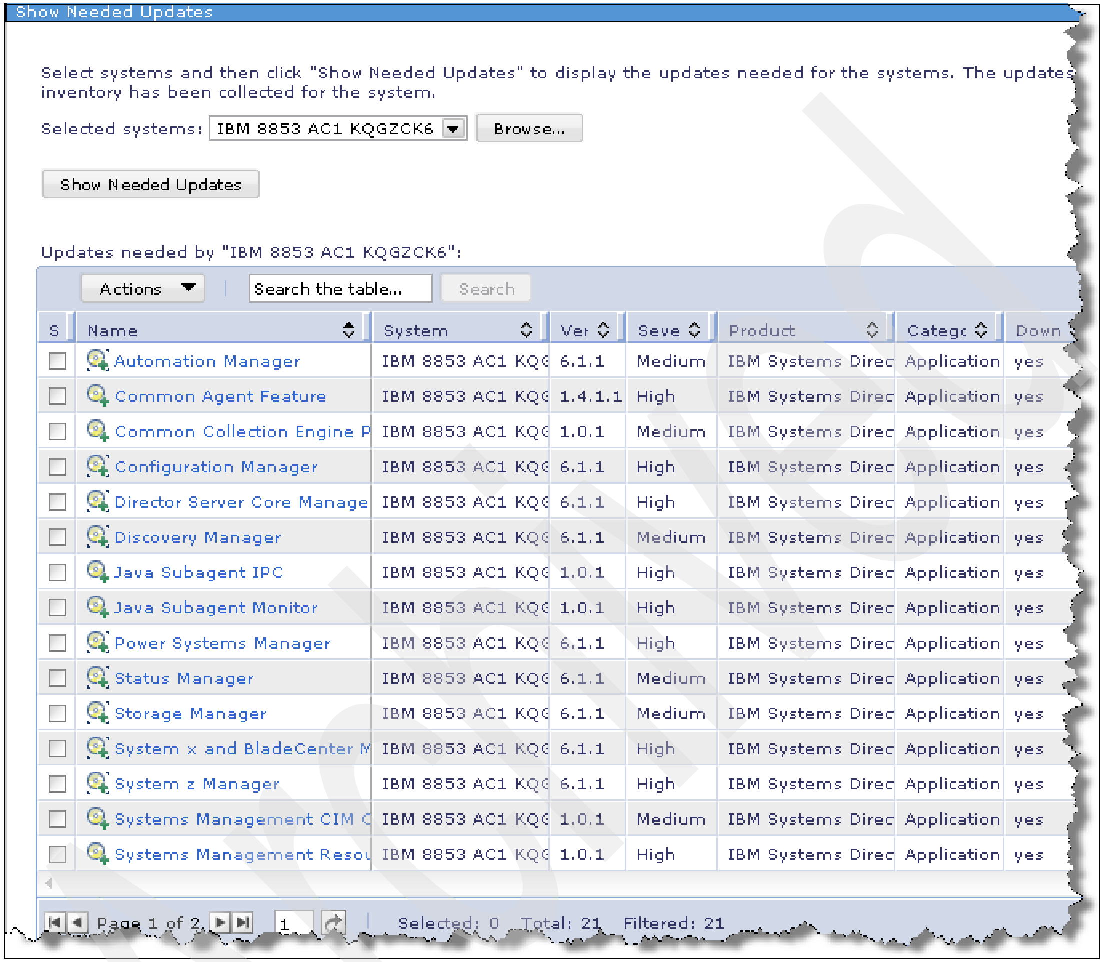


*Figure 3-2 Selecting the updates to apply*

After applying the updates, verify in the Welcome page of IBM Systems Director Console that IBM Systems Director has the value of `6.1.1` as shown in Figure 3-3.


*Figure 3-3 IBM Systems Director version*

## 3.2 Installing Active Energy Manager

Active Energy Manager must be installed on the IBM Systems Director server. For a list of supported hardware and software see 2.3.1, “Hardware requirements for the management server” on page 53. If you need details about how to install IBM Systems Director V6.1, refer to *Implementing IBM Systems Director 6.1*, SG24-7694. In this section, we discuss how to install Active Energy Manager on Windows, AIX, and Linux.

## 3.2.1 Installing Active Energy Manager on Windows

You can install Active Energy Manager on Windows using either the graphical user interface or the command line.

### Installation using the graphical user interface

To install Active Energy Manager on Windows using the graphical user interface, follow these steps:

1. Using an account with either local or domain Administrator authority, log on to the operating system.
2. Extract the contents of the downloaded installer package to a folder.
3. Run the Active Energy Manager installation file to start the installation wizard. In the window that opens (Figure 3-4), select the language for the installation, and click **OK**.

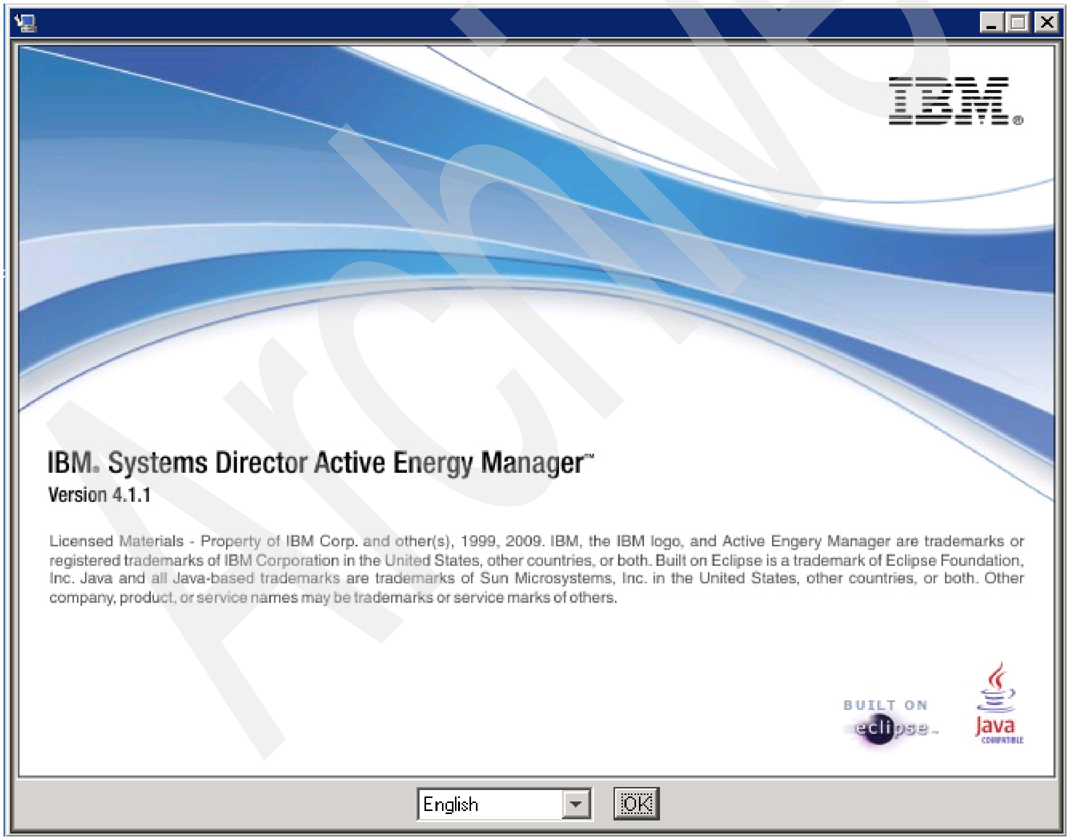


*Figure 3-4 Active Energy Manager installation wizard*

4. On the next window, click **Next**.
5. Then, select the "I Accept the terms in the license agreement" option, and select "Yes, restart IBM Systems Director automatically" if you want the services to start automatically (recommended). Click **Install**.
6. When the installation is complete, click **Done**.

### Unattended installation using the command line

To install Active Energy Manager using a response file and the command line, follow these steps:

1. Using an account with either local or domain Administrator authority, log on to the operating system.
2. Extract the contents of the downloaded installer package to a temporary directory.
3. In an ASCII text editor, open the `installer.properties` file.
4. Edit the following lines to enable the installation to run silently:

   ```
   INSTALLER_UI=Swing
   LICENSE_ACCEPTED=false
   START_SERVER=false
   ```

   becomes

   ```
   INSTALLER_UI=silent
   LICENSE_ACCEPTED=true
   START_SERVER=false
   ```

   > **Note:** To have the IBM Systems Director server restart automatically at the conclusion of the installation, edit the file so that `START_SERVER=true`.

5. Open a command prompt, and then change to the directory that contains the Active Energy Manager installation file. From the command prompt, enter the following command:

   ```
   IBMSystemsDirector-AEM-Setup.exe
   ```

6. The Active Energy Manager silent installation starts.

## 3.2.2 Installing Active Energy Manager on AIX

To install Active Energy Manager on IBM Systems Director management servers running AIX, using either the graphical user interface or the command line. To install it using the GUI, you need to run the installation from an graphical environment such as CDE, or you can export the display if you are accessing by remote access.

### Installation using the graphical user interface

To install Active Energy Manager on a system running AIX using a GUI, follow these steps:

1. Log on to the operating system as root.
2. Extract the contents of the downloaded installer package to a temporary directory:

   ```
   gzip -cd SysDir_AEM4_1_1_AIX.tar.gz | tar -xvf -
   ```

   Note that the trailing hyphen (-) is required.
3. Change the permissions and run the installation file with the following commands:

   ```
   chmod +x IBMSystemsDirector-AEM-Setup.sh
   ./IBMSystemsDirector-AEM-Setup.sh.
   ```

4. The Active Energy Manager installation wizard opens as shown in Figure 3-4 on page 76. From this point, the procedure is the same as the procedure in Windows. Follow the instructions described in 3.2.1, “Installing Active Energy Manager on Windows” on page 76 to complete the installation.

### Unattended installation using the command line

To install Active Energy Manager using a response file from the command line, follow these steps:

1. Log on to the operating system as root.
2. Extract the contents of the downloaded installer package to a temporary directory.
3. Read and acknowledge the software agreements in the `/license` directory.
4. Change the permissions of the installation file with the following commands:

   ```
   chmod +x IBMSystemsDirector-AEM-Setup.sh
   ```

5. In an ASCII text editor, open the `installer.properties file`.
6. Edit the following lines to enable the installation to run silently:

   ```
   INSTALLER_UI=Swing
   LICENSE_ACCEPTED=false
   START_SERVER=false
   ```

   becomes

   ```
   INSTALLER_UI=silent
   LICENSE_ACCEPTED=true
   START_SERVER=false
   ```

**Note:** To have the IBM Systems Director Server restart automatically at the conclusion of the installation, edit the file so that `START_SERVER=true`.

7. Open a command prompt, then change to the directory that contains the Active Energy Manager installation package.
8. Run the AIX installer file `IBMSystemsDirector-AEM-Setup.sh`.
9. The Active Energy Manager silent installation starts.

**Note:** The installation is complete when a file named `installLog.txt` appears in the `/opt/ibm/director/ActiveEnergyManager directory`.

## 3.2.3 Installing Active Energy Manager on Linux

To install Active Energy Manager on IBM Systems Director management servers running Linux, you can use either the graphical user interface or the command line.

### Installation using the graphical user interface

To install Active Energy Manager on a system running Linux using a GUI, follow these steps:

1. Log on to the operating system as root.
2. Extract the contents of the downloaded installer package to a temporary directory:

   ```
   gzip -cd SysDir_AEM4_1_1.tar.gz | tar -xvf -
   ```

   Note that the trailing hyphen (-) is required.
3. Run the installation file:

   ```
   chmod +x IBMSystemsDirector-AEM-Setup.sh
   ./IBMSystemsDirector-AEM-Setup.sh.
   ```

4. The Active Energy Manager installation wizard opens as shown in Figure 3-4 on page 76. Follow the instructions described in 3.2.1, “Installing Active Energy Manager on Windows” on page 76 to complete the installation.

### Unattended installation using the command line

To install Active Energy Manager on a system running Linux using a response file and a command prompt, follow these steps:

1. Log on to the operating system as root.
2. Extract the contents of the downloaded installer package to a temporary directory.
3. Read and acknowledge the software agreements in the `/license` directory.
4. Change the permissions of the installation file with the following command:

   ```
   chmod +x IBMSystemsDirector-AEM-Setup.sh
   ```

5. In an ASCII text editor, open the `installer.properties file`. Edit the following lines to enable the installation to run silently:

   ```
   INSTALLER_UI=Swing
   LICENSE_ACCEPTED=false
   START_SERVER=false
   ```

   becomes

   ```
   INSTALLER_UI=silent
   LICENSE_ACCEPTED=true
   START_SERVER=false
   ```

   > **Note:** To have the IBM Systems Director Server restart automatically at the conclusion of the installation, edit the file so that `START_SERVER=true`.

6. Open a command prompt, then change to the directory that contains the Active Energy Manager installation file for the appropriate operating system.
7. Run the Linux installer file for the appropriate operating system.
8. The Active Energy Manager silent installation starts.

> **Note:** The installation is complete when a file named `installLog.txt` appears in the `/opt/ibm/director/ActiveEnergyManager` directory.

## 3.3 Installing the Active Energy Manager license

As described in 2.5, “Licensing” on page 64, when you install Active Energy Manager from the Web download, you can perform all energy monitoring functions. In addition, you have a 60-day trial license to perform all energy management functions. After 60-days, you need to purchase and install a license to continue using the management functions.

The license that you purchase is in the form of a CD-ROM that contains software that you need to install and to activate the license.

> **Note:** You cannot install Active Energy Manager using this CD-ROM. Active Energy Manager is available *only* as a download from the Web as described in 2.4, “Downloading Active Energy Manager” on page 63. The CD-ROM contains only the license. To install the license, you must first download Active Energy Manager and install both IBM Systems Director and Active Energy Manager.

To install the Active Energy Manager license, follow these steps:

1. Insert the license CD-ROM in an optical drive that is accessible by the operating system on which IBM Systems Director 6.1.1 and Active Energy Manager 4.1.1 are installed. The CD-ROM contains the files and directories shown in Figure 3-5.


*Figure 3-5 Files on the license CD-ROM*

2. For Windows servers, we recommend that you use the graphical user interface to install the license as follows:

   a. Run the installer executable file from the Windows directory on the CD-ROM. An Introduction window opens, as shown in Figure 3-6.

   b. Read the text and click **Next**. Then, follow the prompts to complete the license installation.

   c. Continue to step 4 on page 83 to confirm that the license is installed successfully.

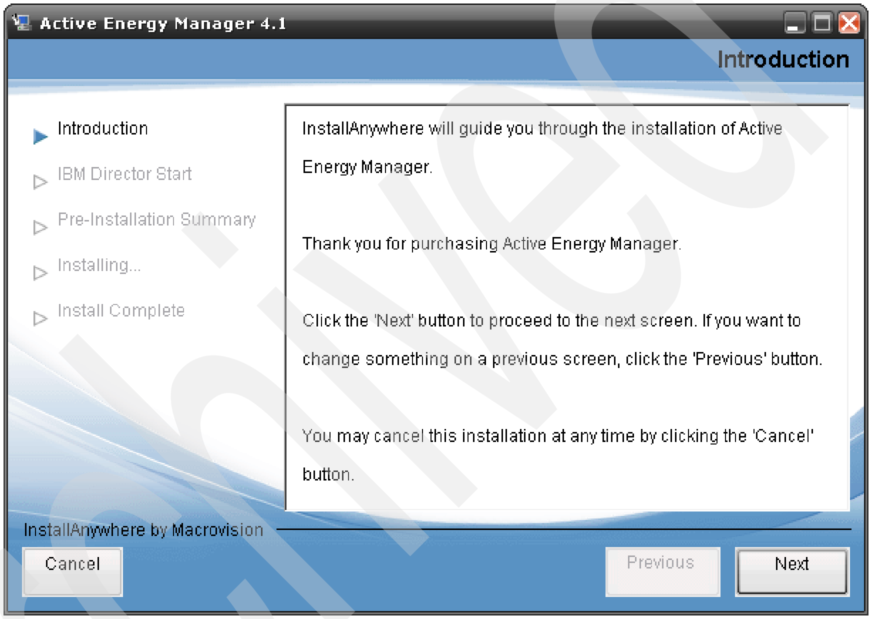


*Figure 3-6 Installing the license using the GUI*

3. For AIX and Linux servers, we recommend that you use the command-line interface to install the license as follows:

   a. Copy the directory from the CD-ROM that contains the installer for your operating system (for example `Linux_x86`) and paste the directory to a temporary directory on the IBM Systems Director server.

   b. Copy the `installer.properties` file from the root of the CD-ROM (see Figure 3-5 on page 81) and paste the file into the directory that you copied to the server in step a (in our example, `Linux_x86`). The `installer.properties` file should now be with the other installation files for that operating system, not the directory one level up.

c. As described in the `ReleaseNotes.txt` file on the CD-ROM, edit the `installer.properties` file that you copied onto the IBM Systems Director server in the previous step. Change the following line:

```
INSTALLER_UI=Swing
```

to:

```
INSTALLER_UI=silent
```

This configures the installer to run in silent mode.

d. Add the following line to the `installer.properties` file:

```
LICENSE_ACCEPTED=true
```

This step is required to correct an issue with the installer. It is required for all point releases of Active Energy Manager 4.x.

e. Run the installer executable file on the IBM Systems Director server.

f. Verify that the license installation is complete by browsing the file `director/ActiveEnergyManager/installLog.txt`. This file is updated with the current date when the installation completes without errors.

g. You can verify that the permanent license is installed by looking in the IBM Systems Director keys directory. You should see a file named `AEM41.lic`.

4. Start an IBM Systems Director Active Energy Manager browser session and sign in.

5. On the Active Energy Manager home page, scroll to the License section at the bottom of the Web page. You should see the display shown in Figure 3-7.


*Figure 3-7 License installed*

# 3.4 Uninstalling Active Energy Manager

This section explains the procedures to uninstall Active Energy Manager on AIX, Windows, and Linux.

## 3.4.1 Uninstalling Active Energy Manager on systems running Windows

Use these instructions to uninstall Active Energy Manager on IBM Systems Director management servers running Windows.

> **Note:** If you want the uninstaller to remove Active Energy Manager data from the IBM Systems Director database, ensure that both Active Energy Manager and IBM Systems Director are running before you continue. The Active Energy Manager group is removed only if IBM Systems Director is active when you begin the uninstallation process.

As with the installation, you can use either the GUI or the command-line interface to uninstall Active Energy Manager.

### Uninstall using the graphical user interface

To uninstall Active Energy Manager on Windows using the GUI, follow these steps:

1. Using an account with either local or domain Administrator authority, log on to the operating system.
2. Click **Start** → **Settings** → **Control Panel**. The Control Panel window opens.
3. Double-click **Add/Remove Programs**. The Add/Remove Programs window opens.
4. Click IBM Systems Director Active Energy Manager, and then click **Change/Remove**.
5. The Active Energy Manager uninstallation wizard begins. Follow the instructions and prompts to remove the necessary files.

### Uninstall using the command line

To uninstall Active Energy Manager on Windows using the command line, follow these steps:

1. Using an account with either local or domain Administrator authority, log on to the operating system.
2. In the `ActiveEnergyManager/uninstall` directory, edit the `installer.properties` file as follows:

   ```
   INSTALLER_UI=Swing
   DATABASE_CLEANUP=false
   ```

   becomes

   ```
   INSTALLER_UI=silent
   DATABASE_CLEANUP=false
   ```

   > **Note:** Set `DATABASE_CLEANUP=true` to have database cleanup performed; Note that this option work only if IBM Systems Director Server is running at the time of the uninstall procedure. If IBM Systems Director Server is inactive at the time of the uninstall, then the database is not deleted, regardless of the value of this parameter.

3. Launch the uninstaller.

## 3.4.2 Uninstalling Active Energy Manager on systems running AIX

To uninstall Active Energy Manager on IBM Systems Director management servers that are running AIX, you can using either the graphical user interface or the command line.

> **Note:** If you want the uninstaller to remove Active Energy Manager data from the IBM Systems Director database, ensure that both Active Energy Manager and IBM Systems Director are running before you continue. The Active Energy Manager group is removed only if IBM Systems Director is active when you begin the uninstallation process.

### Uninstall using the graphical user interface

To uninstall Active Energy Manager on AIX using the GUI, follow these steps:

1. Log on to the operating system as root.
2. From the command prompt, launch the uninstaller by entering the following command:

   ```
   /opt/ibm/director/ActiveEnergyManager/uninstall/uninstall
   ```

### Uninstall using the command line

To uninstall Active Energy Manager on AIX using the command line, follow these steps:

1. Log on to the operating system as root.
2. In the ActiveEnergyManager/uninstall directory, edit the `installer.properties` file as follows:

```
INSTALLER_UI=Swing
DATABASE_CLEANUP=false
```

   becomes

```
INSTALLER_UI=silent
DATABASE_CLEANUP=false
```

> **Note:** Set `DATABASE_CLEANUP=true` to have database cleanup performed.

3. Launch the uninstaller.

## 3.4.3 Uninstalling Active Energy Manager on systems running Linux

To uninstall Active Energy Manager on IBM Systems Director management servers that are running Linux, you can using either the graphical user interface or the command line.

> **Note:** If you want the uninstaller to remove Active Energy Manager data from the IBM Systems Director Database, ensure that both Active Energy Manager and IBM Systems Director are running before you continue. The Active Energy Manager group is removed only if IBM Systems Director is active when you begin the uninstallation.

### Uninstall using the graphical user interface

To uninstall Active Energy Manager on Linux using the GUI, follow these steps:

1. Log on to the operating system as root.
2. From the command prompt, launch the uninstaller by typing the following command and pressing Enter:

```
/opt/ibm/director/ActiveEnergyManager/uninstall/uninstall
```

### Uninstall using the command line

To uninstall Active Energy Manager on Linux using the command line, follow these steps:

1. Log on to the operating system as root.
2. In the `ActiveEnergyManager/uninstall` directory, edit the `installer.properties` file such that:

   ```
   INSTALLER_UI=Swing
   DATABASE_CLEANUP=false
   ```

   becomes

   ```
   INSTALLER_UI=silent
   DATABASE_CLEANUP=false
   ```

3. Launch the uninstaller.

**Note:** Set `DATABASE_CLEANUP=true` to have database cleanup performed.

## 3.5 Migrating Active Energy Manager

You can migrate data and settings from Active Energy Manager 3.1.x or later to Active Energy Manager 4.1 using the IBM Systems Director Migration Tool. Migration to later versions, such as version 4.1.1, requires that you first migrate to Active Energy Manager 4.1, then upgrade to the later version.

The IBM Systems Director Migration Tool is a command line utility that includes the following main functions:

- **`smexport`**
- **`smimport`**

Both **`smexport`** and **`smimport`** are invoked from the command line. You export data from Active Energy Manager 3.1.x by running **`smexport`**, and you import data into Active Energy Manager 4.1 by running **`smimport`**.

> **Important:** To download and use the IBM Systems Director Migration Tool, see the *Upgrading and Migrating IBM Systems Director* topic in the IBM Systems Director V6.1 documentation, which is available at:
>
> http://publib.boulder.ibm.com/infocenter/director/v6r1x/index.jsp?topic=/director.upgrade_6.1/fqm0_t_upgrading_and_migrating.html

### 3.5.1 Migration considerations and notes

You can migrate only data and settings from Active Energy Manager 3.1 or later. You cannot migrate from PowerExecutive™ (the former name of Active Energy Manager prior to Version 3.1).

Consider the following settings when using the IBM Systems Director Migration Tool to migrate Active Energy Manager 3.1.x data and settings:

- Global default settings

  The following global default settings are migrated:

  - Default polling interval in minutes (3.1.x) → Default metering interval (4.1)

    > **Note:** The 4.1 default value is 5 minutes. If the 3.1.x value was set to 4 minutes or less, the value migrated to 4.1 is set to 5 minutes. Otherwise, the 3.1.x value is migrated.

  - Number of days to keep trend data (3.1.x) → Retain data (4.1)
  - Refresh rate in minutes (3.1.x) → Default data refresh interval (4.1)

- Price per kilowatt-hour (3.1.x) → Default energy price and Default currency type (4.1)

  > **Note:** This setting is migrated only if the currency type on the 4.1 system is set to *USD*.

- Cooling rate factor (3.1.x) → Default cooling rate multiplier (4.1)
- Default voltage (3.1.x) → Voltage (4.1)

► Resources

Resource considerations related to migration include:

- Migration of Active Energy Manager 3.1.x resources might require network connectivity to those resources during import. Resources that cannot be contacted cannot be migrated. Resources not migrated will not have associated settings, events, or polling data migrated.
- If a resource is not migrated, a warning message displays during import.
- A resource that is plugged into a PDU+ outlet but is not an Active Energy Manager resource is not associated with the PDU+ outlet. Instead, Active Energy Manager migration creates another resource of type *Metered Device* that represents the first, has the same name, and is an Active Energy Manager resource. This resource is associated with the PDU+ (Power Unit) outlet during migration.
- Each BladeCenter resource is migrated, as well as the data for it, such as event and power use data at the BladeCenter level. However, some or all modules within the BladeCenter are not migrated, because inventory is required to discover them, and inventory is not run during migration. Therefore, Active Energy Manager data for BladeCenter slots cannot be migrated.

► Resource settings

The following settings for individual Active Energy Manager resources are migrated:

- Polling interval in minutes (3.1.x) → Metering interval (4.1)

  > **Note:** The 4.1 default value is 5 minutes. If the 3.1.x value was set to 4 minutes or less, the value migrated to 4.1 is set to 5 minutes. Otherwise, the 3.1.x value is migrated.

- Time (3.1.x) → Last time metered (4.1)
- Polling active (3.1.x) → Metering active (4.1)

  - Relationships to other resources, such as PDU+ outlets associated with resources plugged into them.
  - Text descriptions for location and description, in some cases.

- Trend data

  Trend data for all resources discovered or created during migration is migrated. Historical data compression of migrated trend data is not done during migration. In the 4.1 release, to save disk space and improve query performance of large intervals of data, Active Energy Manager automatically compresses metered data that is more than 7 days old. Each night at midnight, data is compressed to 1 hour averages for both power and environmental values.

- Power management settings and policies

  Power management settings and policies considerations related to migration include:

  - Migration of power management settings, such as policies, occur only if an evaluation or permanent Active Energy Manager license is active when the import is performed.
  - Policies that set power cap based on historical data are not migrated.
  - Policies that set both power cap and power saver are migrated but are split into two policies, one ending with `-- POWER CAP PORTION` and the other ending with `-- POWER SAVER PORTION`.
  - Tasks are created in Active Energy Manager 4.1 for all migrated policies. However, scheduling of these tasks is not migrated.
  - Policies are not applied to any resources by migration. You must apply policies manually to resources post-migration.

- Events

  Event considerations related to migration include:

  - Migrated events are displayed in English only.
  - Events not associated with specific Active Energy Manager 3.x resources are not migrated. This might include some SNMP trap events and events that are associated with nodes that could not be discovered during migration.

- Storage requirements

  Migration of Active Energy Manager objects and data requires that extra disk space be available on both the export and import systems before you start. Before exporting this data, look up the size of the Active Energy Manager database. This value is displayed in different places depending on which version of Active Energy Manager you are running:

  – Active Energy Manager 3.1: Click **Edit** → **Manage Trend Data**.

  – Active Energy Manager 3.1.1: Click **Edit** → **Preferences**, then select the Server tab.

  There are two locations in the file system that must have free space in order for Active Energy Manager migration to work:

  – The directory that will hold the data that is exported by the migration tool. Export will create the directory, and import will read from it.

  – The temporary working directory that Active Energy Manager migration uses during migration to accomplish its work. By default, this is either the Windows temporary directory, or `/tmp` on Linux. You can select a different location, however, by creating a file called `aemmigration.properties` in the <*IBM Director install path*>`/data` directory, containing one entry, as follows:

    ```
    com.ibm.aem.migration.TemporaryDirectory=fully-qualified
    directory name
    ```

    You can select a different location for each system (if needed) for which you want to specify a temporary location. If `none` is specified on a system, the system temporary directory is used. The directory specified on the source system must exist. The directory specified on the target system will be created if it does not already exist. If you use a backslash in the path, it must be escaped with another backslash character. Use a double-backslash instead of a single-backslash in this file.

  After you know the database size and these locations, calculate the needed disk space:

  – For export from 3.1.x:

    - Approximately 50% of the Active Energy Manager database size must be available in the temporary working directory used by Active Energy Manager migration.
    - Approximately 50% of the database size must be available in the directory to which the migration data will be exported.

– For import to 4.1:

  - Approximately 50% of the database size must be free in the directory into which the exported migration data is copied.
  - Approximately 100% of the Active Energy Manager database size must be available in the temporary working directory used by Active Energy Manager migration.
  - At least 100%, but recommended to be much more, of the database size must be available in <IBM Director install path>/data, to hold the migrated data in Active Energy Manager 4.1 format.

## 3.6 Updating to Active Energy Manager 4.1.1 from 4.1

This section shows a step-by-step procedure to update Active Energy Manager 4.1.1 in the scenario where Active Energy Manager 4.1 is already installed.

As explained in 3.1, "Updating IBM Systems Director" on page 72, you can update Active Energy Manager using two methods. In this case, we import the updates instead of updating online.

For more information about updating IBM Systems Director, refer to:

http://publib.boulder.ibm.com/infocenter/director/v6r1x/index.jsp?topic=/director.updates_6.1/fqm0_t_um_updating_director_features.html

> **Important:** IBM Systems Director 6.1.1 and Active Energy Manager 4.1 are required to run this procedure.

### 3.6.1 Using the command-line interface to update from 4.1 to 4.1.1

To update Active Energy Manager using the command-line interface, follow these steps:

1. Using a Web browser, go to the following Web address:

   http://ibm.com/support/fixcentral/

   Then, follow these steps:

   a. Select **IBM Systems Director** in the Product Group drop-down menu.
   b. Select **IBM Systems Director** from the Product drop-down menu.
   c. Select **6.1** from the Installed version drop-down menu.
   d. Select **ALL** from the Platform drop-down menu.
   e. Click **Continue**.
   f. Select **com.ibm.aem_4.1.1**.

   g. Click **Continue**.
   h. Select **I agree** after you review the terms and conditions of the agreement.
2. Run inventory, to ensure that Active Energy Manager 4.1.1 is in the software inventory:

```
smcli collectinv -i <systemName> -p "All Software Inventory"
```

3. Import the fix *com.ibm.aem_4.1.1*:

```
smcli importupd -v -t c:\temp\com.ibm.aem_4.1.1
```

4. Verify that the import was successful by running **smcli lsupd**. The output should include the following lines:

```
com.ibm.aem.common_4.1.1
com.ibm.aem.console_4.1.1
com.ibm.aem.discovery_4.1.1
com.ibm.aem.nonplugin_4.1.1
com.ibm.aem.server_4.1.1
com.ibm.aem_4.1.1
```

5. Run the following command, and pipe it into an output file because the output can be long:

```
smcli lsrelationships appliesTo > output.txt
```

6. Browse the output file, and look for `com.ibm.aem`. If you find it, then the task completed successfully. If you do not see this line, then wait a couple of minutes and try the command again. If the line is still not there, run the inventory as described in step 2.
7. Install the updates:

```
smcli installupd -i <systemName> -u com.ibm.aem_4.1.1 -v
```

8. Restart IBM Systems Director Service.
9. Verify that Active Energy Manager is at release 4.1.1 in the Welcome panel.

### 3.6.2 Using the GUI to update from 4.1 to 4.1.1

To update Active Energy Manager using the GUI, follow these steps:

1. In a Web browser, go to:

   http://www-933.ibm.com/support/fixcentral/

   Then, follow these steps:

   a. Select **IBM Systems Director** in the Product Group drop-down menu.
   b. Select **IBM Systems Director** from the Product drop-down menu.
   c. Select **6.1** from the Installed version drop-down menu.
   d. Select **ALL** from the Platform drop-down menu, and click **Continue**.

e. Select **com.ibm.aem_4.1.1**, and click **Continue**.
f. Click **I agree** after reviewing the terms and conditions.

2. In the IBM systems Director Web interface, select **Inventory** → **View and Collect Inventory** as shown in Figure 3-8.

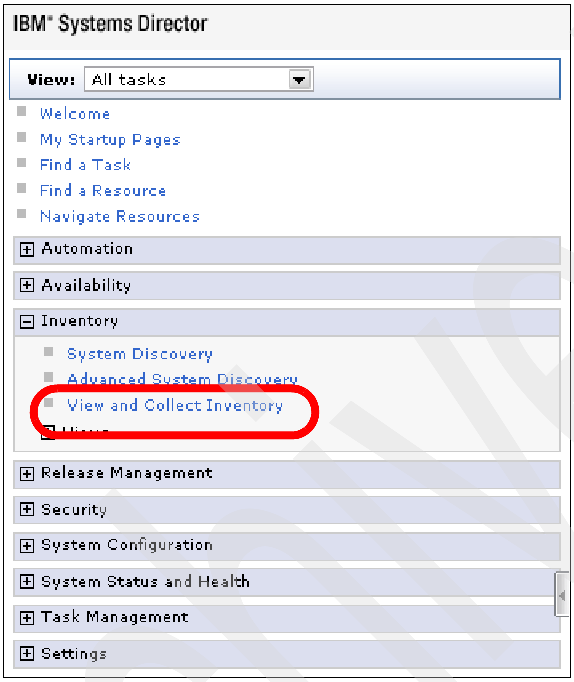


*Figure 3-8 View and Collect Inventory*

3. In the View and Collect Inventory window, shown in Figure 3-9, select IBM Systems Director server by clicking **Browse** under Target Systems. Then click **Collect Inventory**.

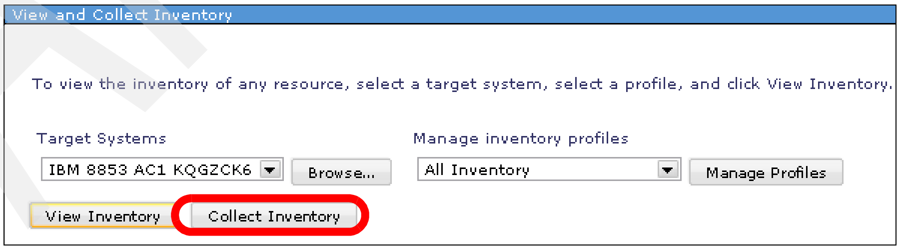


*Figure 3-9 Collect Inventory*

4. In the Run - Collect Inventory window, shown in Figure 3-10, click **OK** to run the inventory collection immediately.

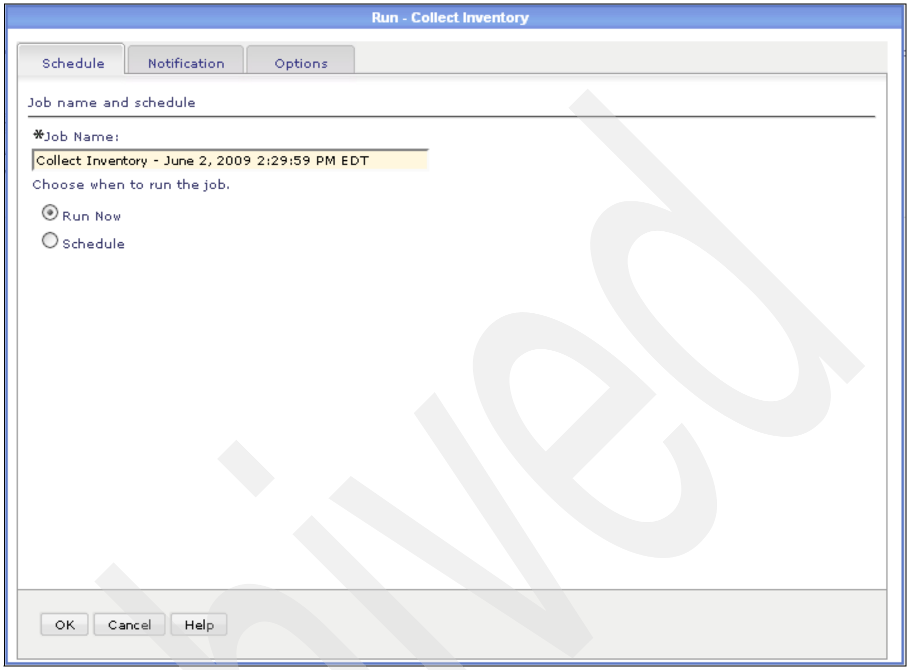


*Figure 3-10 Run the job immediately*

5. To verify that IBM Systems Director successfully collected the inventory, select **Display Properties**, as shown in Figure 3-11.

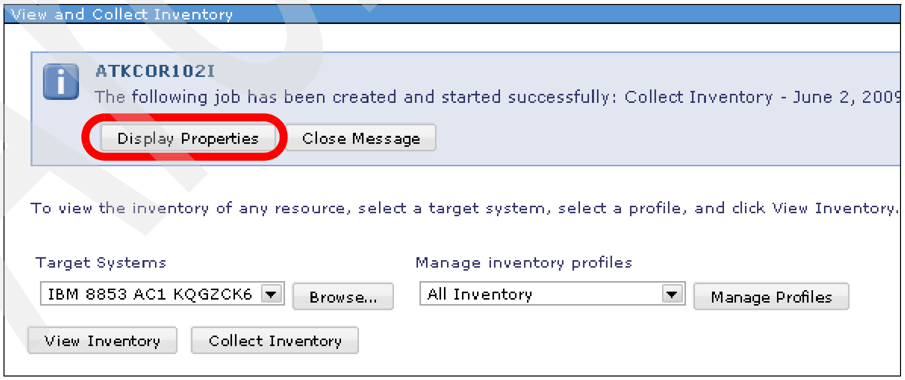


*Figure 3-11 Verify the collect inventory by selecting Display Properties*

6. Now, you can import the update. In the IBM Systems Director Web Interface, click **Release Management** → **Updates** as shown in Figure 3-12.

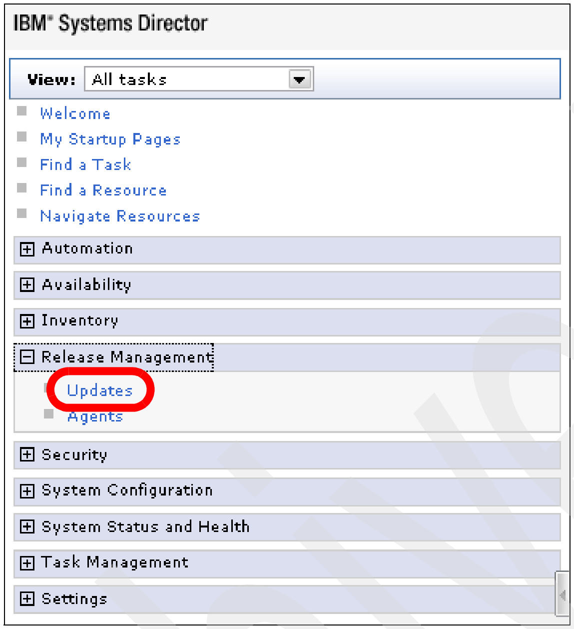


*Figure 3-12 Release Management menu*

7. In the Update Manager window, click **Import Updates** under Command tasks, as shown in Figure 3-13.

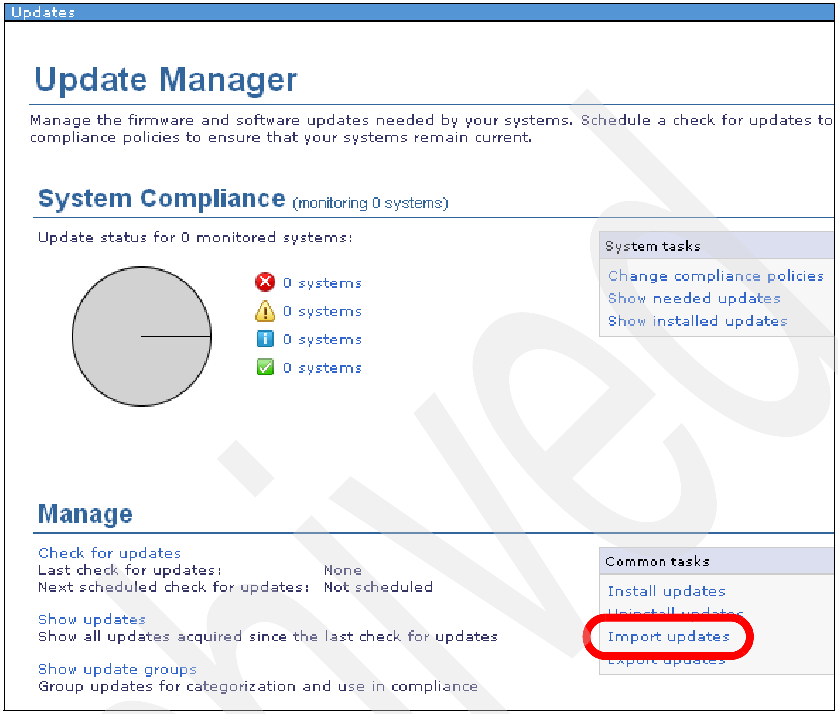


*Figure 3-13 Importing the updates through the Update Manager window*

8. Specify the path to the directory in which the updates to import are located. Figure 3-14 shows the update path for a Linux installation.

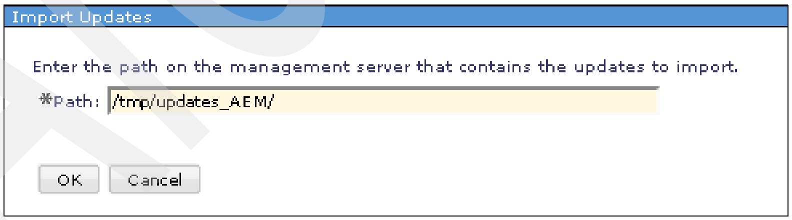


*Figure 3-14 Specifying the path to the updates*

9. In the Run - Import Updates window, shown in Figure 3-15, click **OK** to run the importing update process immediately.

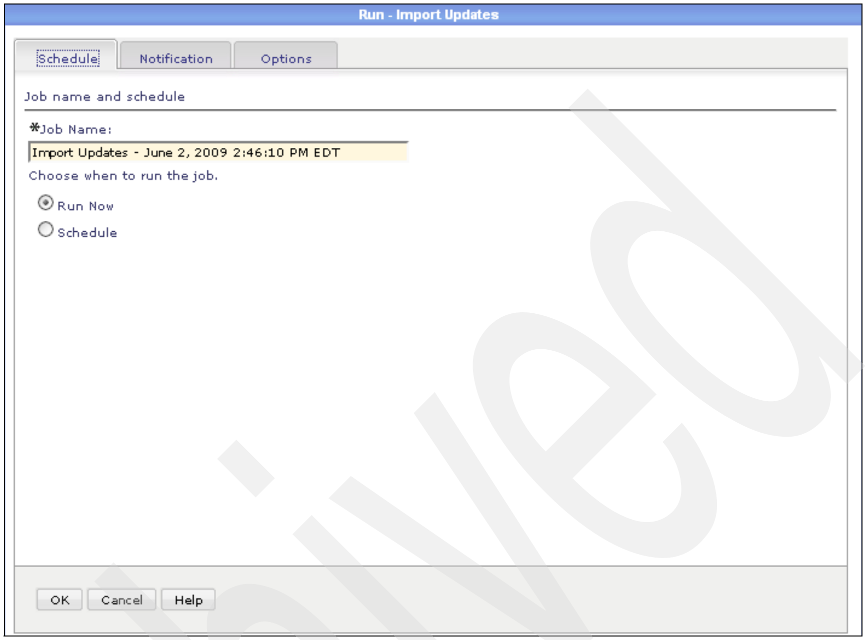


*Figure 3-15 Run the job immediately*

10. In the Active and Scheduled Jobs window on the General tab, when the status bar shows Complete at 100%, as shown in Figure 3-16, confirm that it the update shows no errors (Complete with errors). If errors occur, go to the Logs tab to see the errors.

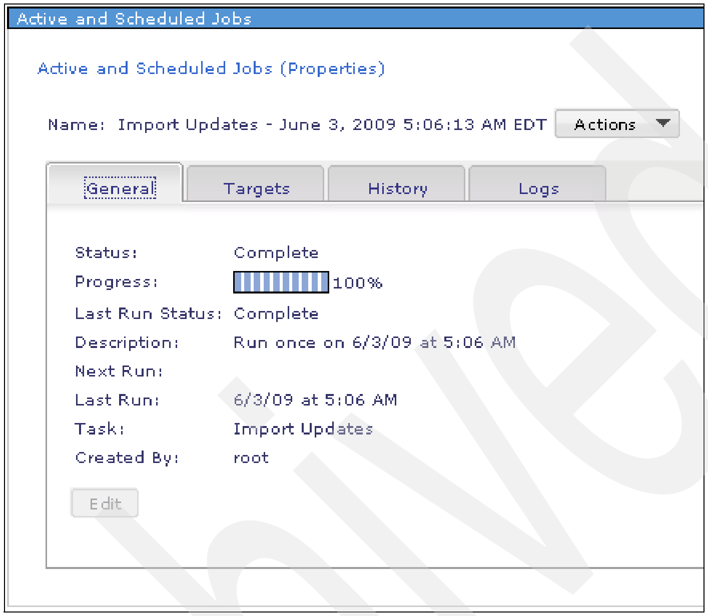


*Figure 3-16 Status 100% Complete*

11. In the Update Manager window, select **Install updates** as shown in Figure 3-17.

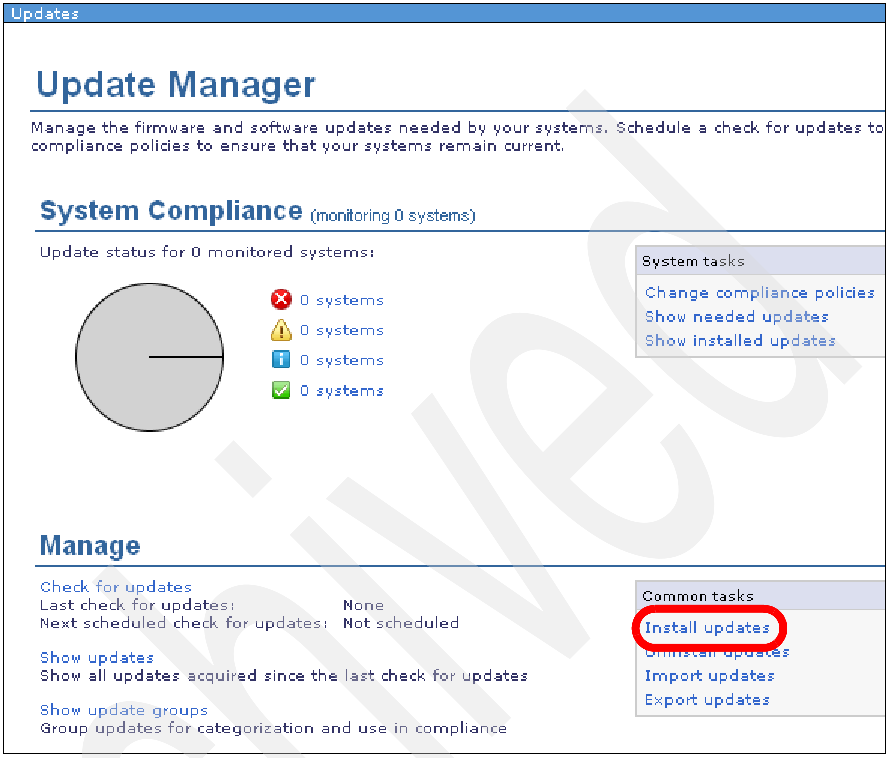


*Figure 3-17 Installing the updates through the Update Manager window*

12. In the Welcome section of the installation wizard, click **OK** (Figure 3-18).

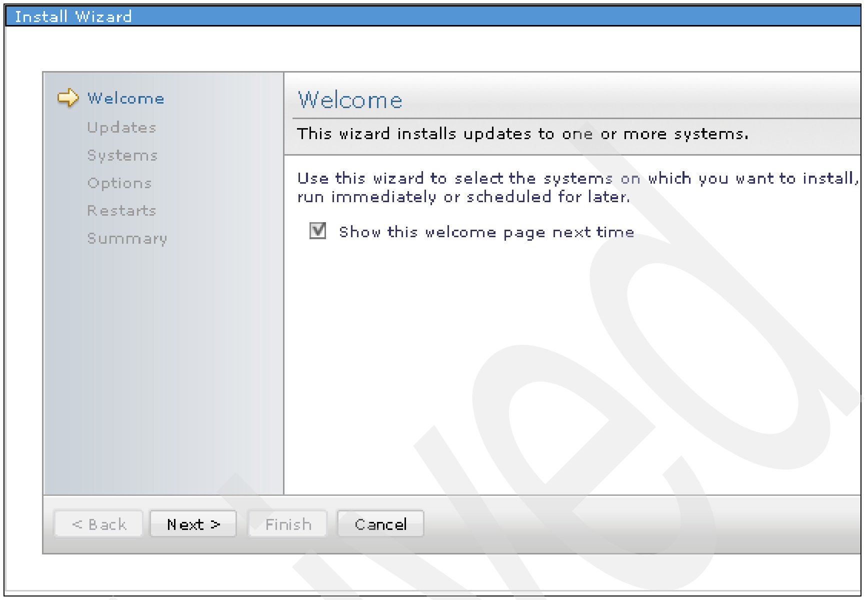


*Figure 3-18 Welcome page of the installation wizard*

13.In the Updates window, shown in Figure 3-19, the wizard shows which updates are available to install. Select **Active Energy Manager Version 4.1.1** and click **Add**.

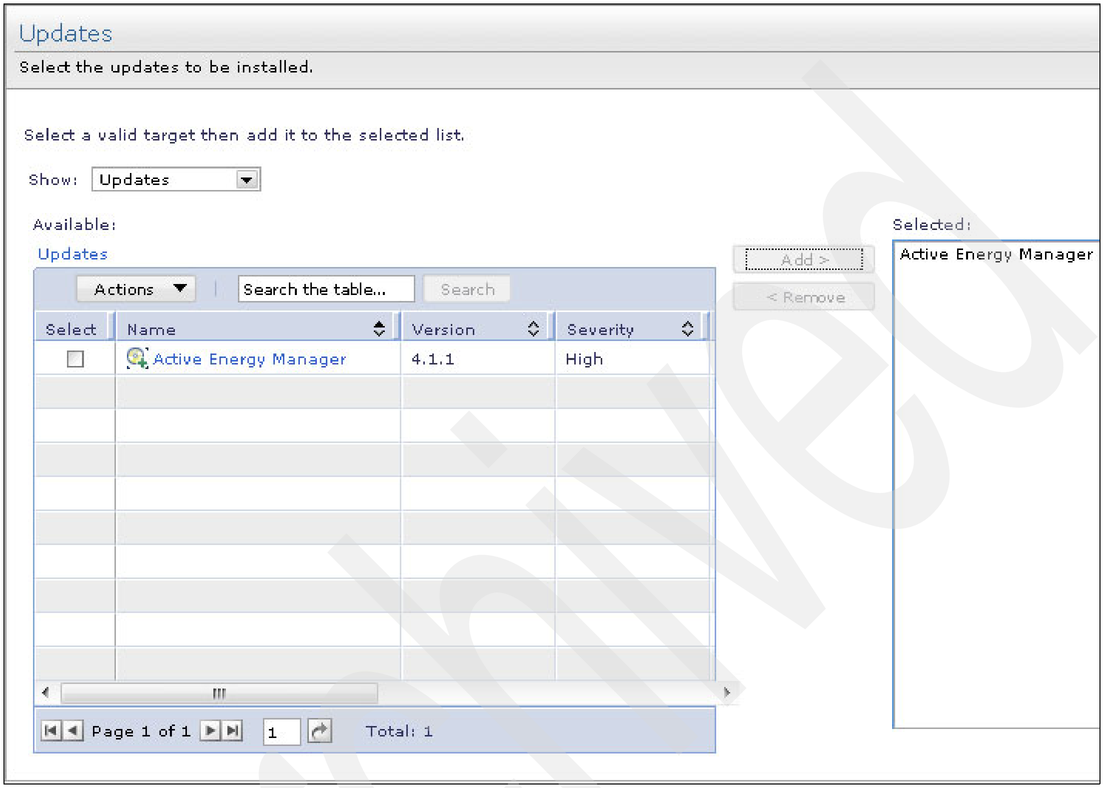


*Figure 3-19 Available updates to apply*

14. Select the IBM Systems Director server on which the updates are needed. Then click **Add** (Figure 3-20).

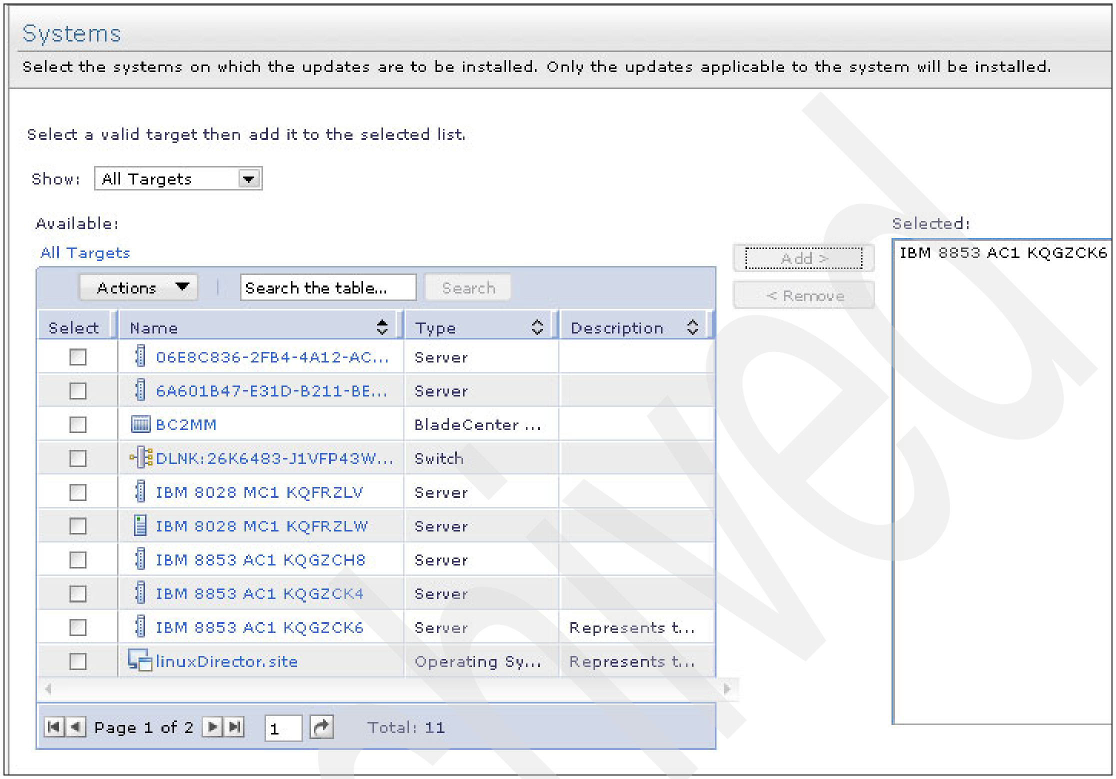


*Figure 3-20 Selecting the system in which apply the updates*

15. Select **Automatically install missing updates requirements** and click **OK** to install missing updates automatically, as shown in Figure 3-21.

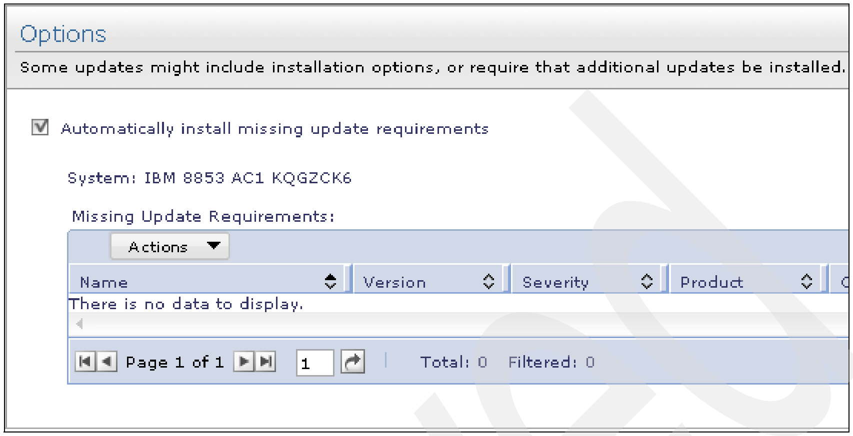


*Figure 3-21 Options section of the wizard*

16. Next, select **Automatically restart as needed during installation** and click **OK** to restart the management server automatically, as shown in Figure 3-22.

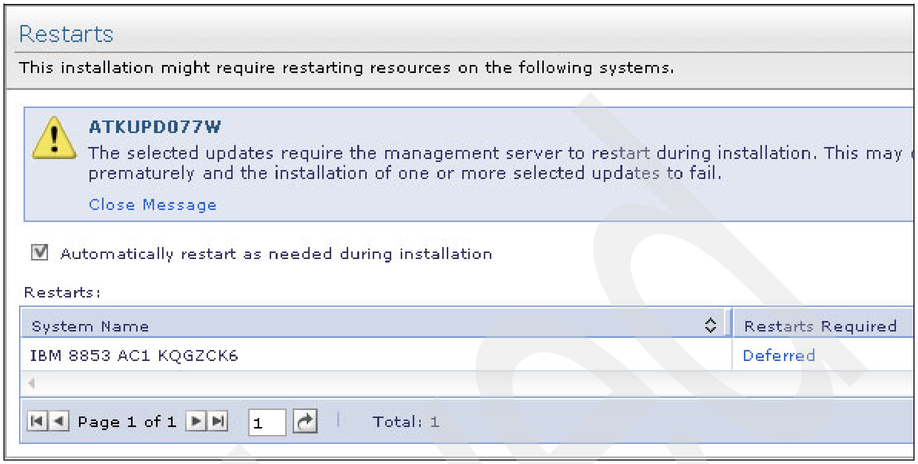


*Figure 3-22 Restarts section of the wizard*

A summary page displays, similar to that shown in Figure 3-23.

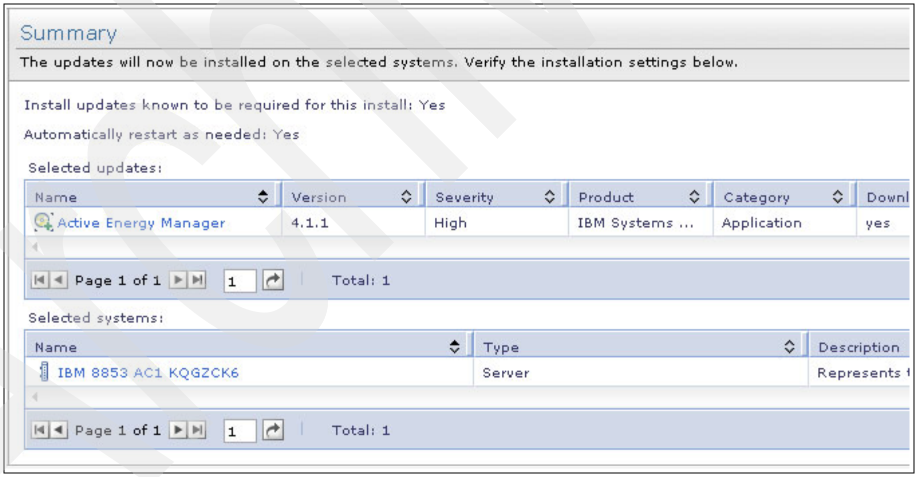


*Figure 3-23 Summary of the updates that will be applied*

17. Restart IBM Systems Director Service, and then verify in the Manage tab of the Welcome page that the version is 4.1.1 for Active Energy Manager, as shown in Figure 3-24.

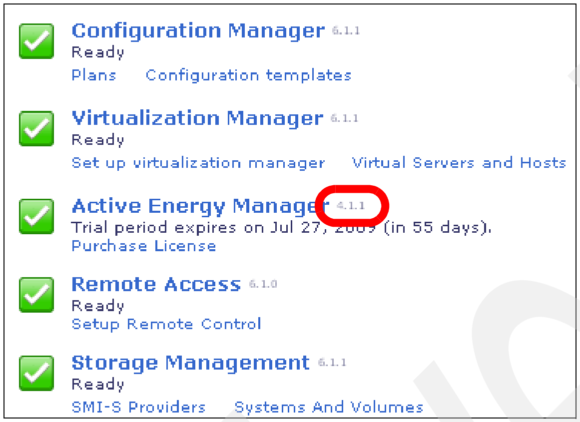


*Figure 3-24 Welcome window of the Web interface*

## 3.7 Troubleshooting the installation

During the Active Energy Manager installation or updating, you might see error messages. In this section, we list some of the common error messages, what the messages mean, and what you need to do to correct the situation:

- Error message: `No X11 DISPLAY variable was set, but this program performed an operation which requires it.`

  Active Energy Manager was not able to open the installation wizard window in Linux or AIX. You need to run the installation wizard from an X Window System desktop environment (KDE, Gnome, CDE, Xfce, and so forth) or enable X11 forwarding from the linux or AIX server.

- Problem: The Welcome page in the IBM Systems Director web interface is not showing the value 6.1.1 after applying the IBM Systems Director updates.

  Sometimes restarting the Systems Director Server service is not enough. Reboot the server to apply correctly the updates.

- Error message: `ATKUPD081I No updates are needed on the selected systems.`

  After Checking for updates, no updates are shown to apply in the View Updates menu. To correct this issue, follow these steps:

  a. Verify that you have collected the inventory.

  b. After importing the updates, the IBM Systems Director Server needs to compare the new updates with the currently installed updates to determine what updates are needed. So, there is a delay after the import finishes before the updates display in the “Show need updates” or “View Updates” views. If the updates still do not display, you can try closing and re-opening the view.

  c. Go to Update Manager menu, and click **Install Updates** to apply the updates.

- Problem: Updates fail to install because pre-requisites are missing.

  An error message such as this one might display in the install updates job activation log:

  ```
  ATKUPD265W Update "com.ibm.director.console.helps.doc_6.1.1"
  requires update "com.ibm.director.main.helps.doc_6.1.1" which does
  not have any metadata acquired yet and will be removed from the list
  of updates to install.
  ```

  This message means that IBM Systems Director could not install the updates because some pre-requisites for that updates were not available. To resolve the issue, try the following fixes:

  – If you imported the updates, verify that the updates were downloaded successfully or retry with the online update.

  – Install the updates as before, and if the missing prerequisite error occurs, retry the installation only on the updates that failed to install.

  – Run the update procedure again using the CLI to get more error messages. For a reference, you can follow 3.6.1, “Using the command-line interface to update from 4.1 to 4.1.1” on page 92.

4

# Navigating the IBM Systems Director Web console

Active Energy Manager runs as a component of IBM Systems Director. Thus, Active Energy Manager inherits powerful features from IBM Systems Director and takes advantage of the integration between the two products.

To use Active Energy Manager effectively, it is important to gain an understanding of the IBM Systems Director Web-based user console. This chapter explains the basics of navigating the console and includes the following topics:

## 4.1 Supported Web browsers

The Web console requires that you use a supported Web browser. The following Web browsers are supported for use with the IBM Systems Director Web console:

- Firefox versions
  - Firefox 2.0
  - Firefox 3.0

  > **Note:** Firefox 3.0 is the minimum required version on SUSE Linux Enterprise Server 10.

- Microsoft Internet Explorer versions
  - Microsoft Internet Explorer 6.0
  - Microsoft Internet Explorer 7.0

## 4.2 Logging in to and out of the Web console

To log in to the IBM Systems Director Web console do the following.

1. From a Web browser, use the following URL:

   `http://`*`system_name`*`:`*`port_number`*`/ibm/console`

   Where:

   - *system_name* is the host name (or IP address) of the system on which IBM Systems Director Server is installed.
   - *port_number* is the first (lower) of two consecutive port numbers that you configured the Web server to use during installation. (Default ports for the Web server are 8421 for http and 8422 for https.)

   > **Note:** If you prefer to use SSL for security, use the second port (by default port 8422) and use https in the URL:
   >
   > `https://`*`system_name`*`:`*`port_number`*`/ibm/console`

2. Enter the user ID and password that corresponds to an authorized IBM Systems Director administrator user ID and password, and click **Log in** (Figure 4-1).

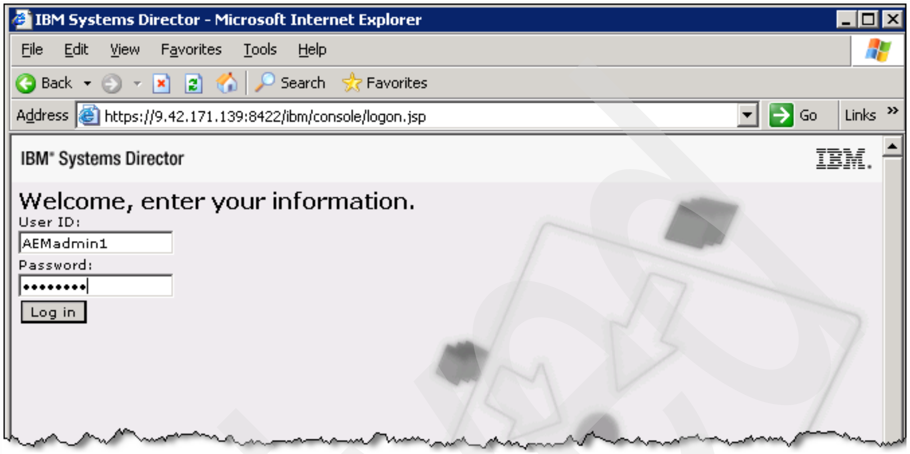


*Figure 4-1 Web console login window*

> **Note:** A security alert window might display before you log in because of an incorrect configuration of the Secure Socket Layer (SSL) certificate.

3. The first time that you log in, you are presented with the welcome page shown in Figure 4-2 on page 112. On the IBM Systems Director Welcome page, you can complete first-time setup steps to make sure that IBM Systems Director and its plug-ins are set up and configured to manage your environment. For a description of the various areas in the Web console, see 4.3, “Layout of the Web console” on page 113.

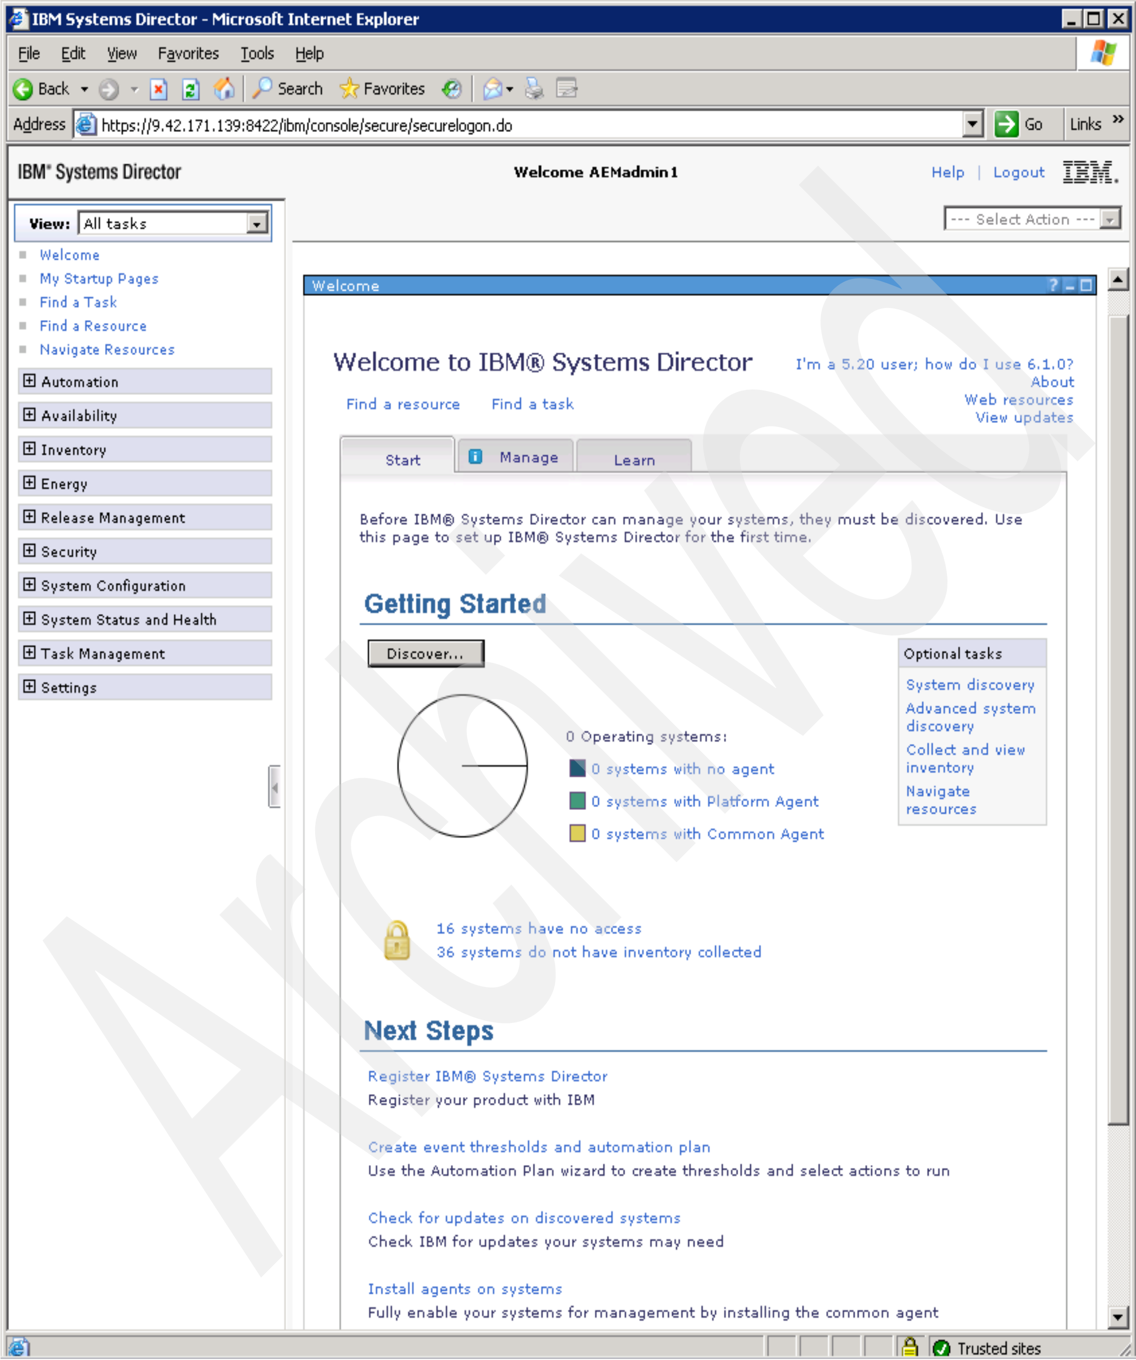


*Figure 4-2 IBM Systems Director welcome view at initial login*

To log out of the IBM Systems Director Web console, locate the link in the top, right corner of the Web console, and click **Logout**.

**Note:** If you do not use the IBM Systems Director Web console actively for 30 minutes, the automatic timeout feature logs you out.

## 4.3 Layout of the Web console

The IBM Systems Director Web console provides tasks and unique views to help you manage your environment. The Web console is divided into six areas, which are highlighted in Figure 4-3.

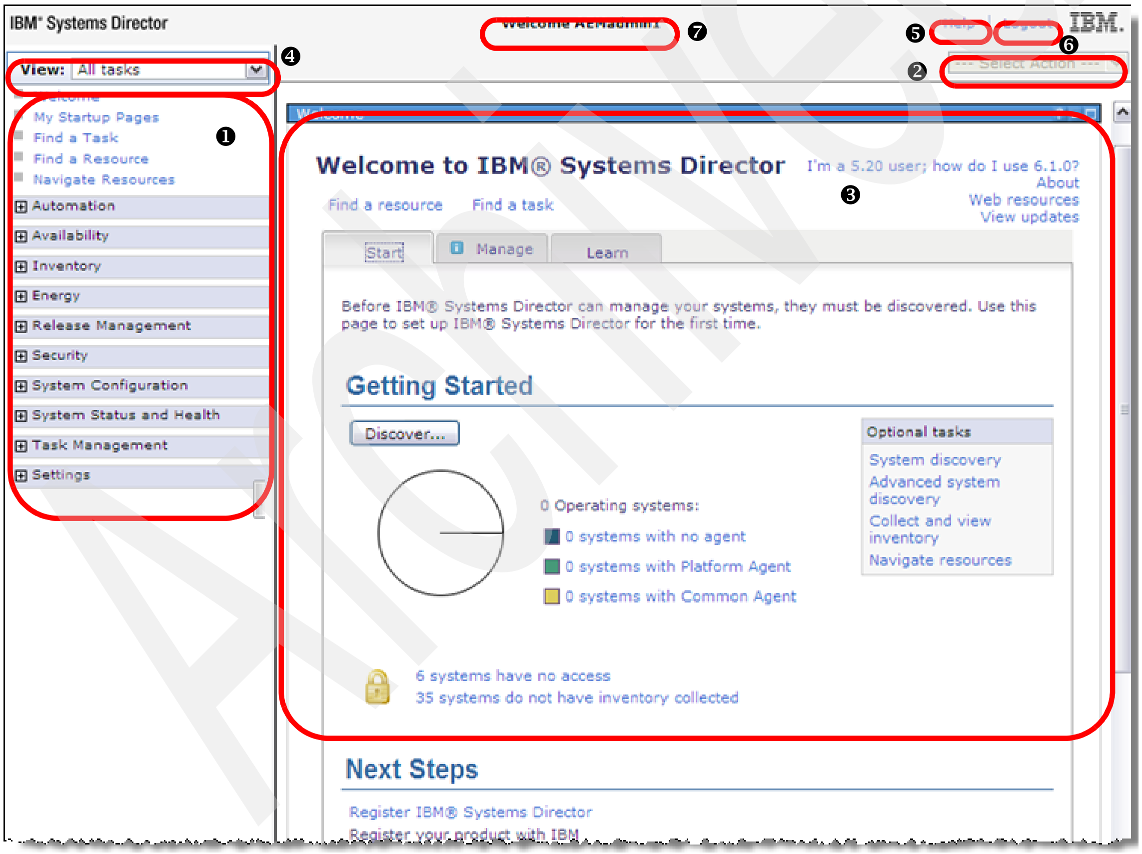


Figure 4-3 Web console layout

The areas indicated in Figure 4-3 on page 113 are:

**1.** Navigation area

The navigation area of IBM Systems Director Web console provides categories of tasks that can vary depending on your IBM Systems Director installation. The navigation area provides links to that you can perform these tasks on your resources. Examples of typical tasks are navigate resources, inventory, health summary, automation, and settings.

**2.** Select Action pull-down menu

The Select Action menu provides the following ways to work with task pages:

- My Startup Pages: Customizes the pages that are started automatically when you log in to IBM Systems Director, the first page that is displayed of these automatically started pages, and the default navigation area view.
- Manage Open Pages: Provides a way to manage and close one or more open pages.
- Close Page: Closes the page that you are viewing.

**3.** Content area

When you open the Web console, by default you see the Welcome page for IBM Systems Director in the content area. The content area changes depending on the item that you select in the navigation area. You can customize aspects of the content area using the Navigation Preferences.

**4.** View drop-down menu

The navigation area of the IBM Systems Director Web console provides categories of tasks that can vary depending on your IBM Systems Director installation. The navigation area provides links to tasks that you can perform on your resources. Examples of typical tasks might include navigate resources, inventory, health summary, automation, and settings. The View menu provides the options shown in Figure 4-4.

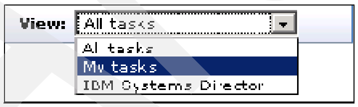


*Figure 4-4 Task options for IBM Systems Director*

**5.** Help

Displays the help information for IBM Systems Director.

**6.** Logout

Logs out of the IBM Systems Director Web console.

**7.** User name identification

You can use multiple IBM Systems Director user IDs so that different types of administrators can manage the system at the same time. At the top of the Web console window, the user ID for the administrator who is logged in displays. For details about setting up multiple user IDs, see the next section, 4.4, “Creating more user IDs” on page 115.

## 4.4 Creating more user IDs

To allow different types of administrators to access IBM Systems Director server at the same time, you can create other user IDs that have different authority. The user IDs that IBM Systems Director uses are operating system user accounts. Therefore, you can create an account in the operating system, and then use IBM Systems Director to grant specific roles to user ID and to specify to which IBM Systems Director group that user ID belongs.

In this example, we use Windows Server 2003. To create multiple user IDs, follow these steps:

1. Open the Windows Computer Management Console shown in Figure 4-5. To create a user ID, select the `Users` folder, and then from the menu bar, click **Action** → **New User**.

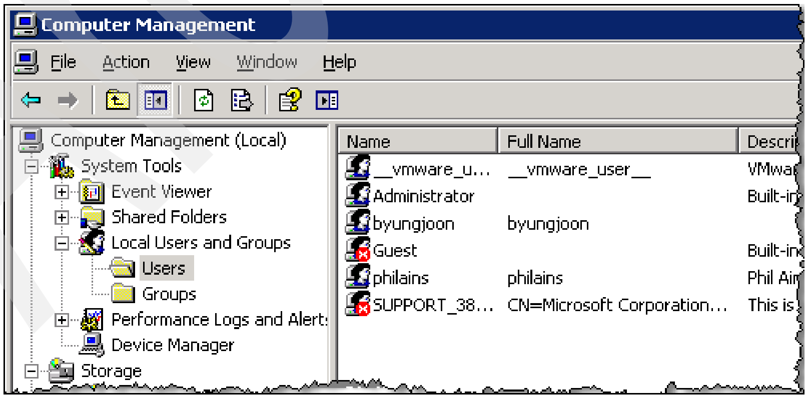


*Figure 4-5 Local User and Group panel on Windows Server 2003*

2. In the New User dialog box, enter a suitable user name and password, and then click **Create**, as shown in Figure 4-6.

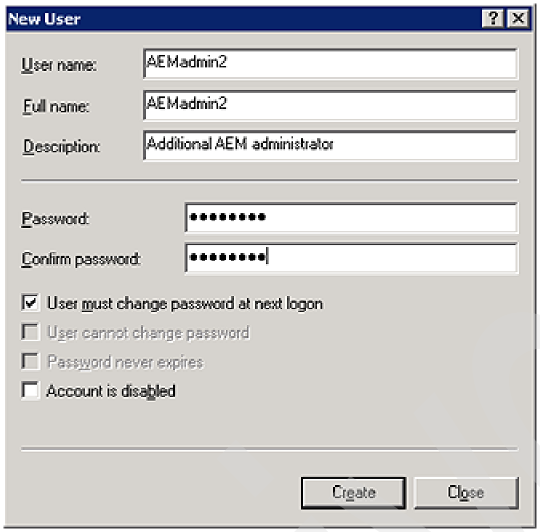


*Figure 4-6 Enter a user name and password*

3. To fully control IBM Systems Director features, the user ID needs to belongs to member of the Administrators group. Click **Action** → **Properties** on the user ID that you just created, and go to the Member Of tab to add the Administrator role, as shown in Figure 4-7.

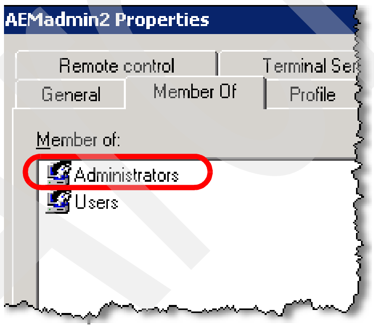


*Figure 4-7 User ID Properties panel*

4. After creating a new account, log in to Windows using this new account at least one time. Otherwise, IBM Systems Director server cannot collect the necessary account information from the operating system.

5. Log in to the Director Console with an administrator ID. Click **Security** → **Users** in the Navigation area as shown in Figure 4-8.

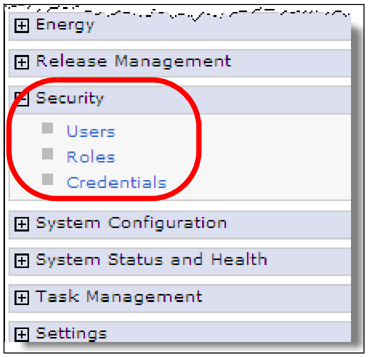


*Figure 4-8 Security tab in the Navigation area*

6. Select the user ID to which you want to give permissions (in this example, Blade43\AEMadmin2 as shown in Figure 4-9), and then click **Assign Role**.

> **Note:** A default role is *GroupRead*, which grants a user the ability to view or open a group. To have a full authority to manage and access system resources, you need to assign a role of *SMAdministrator*.

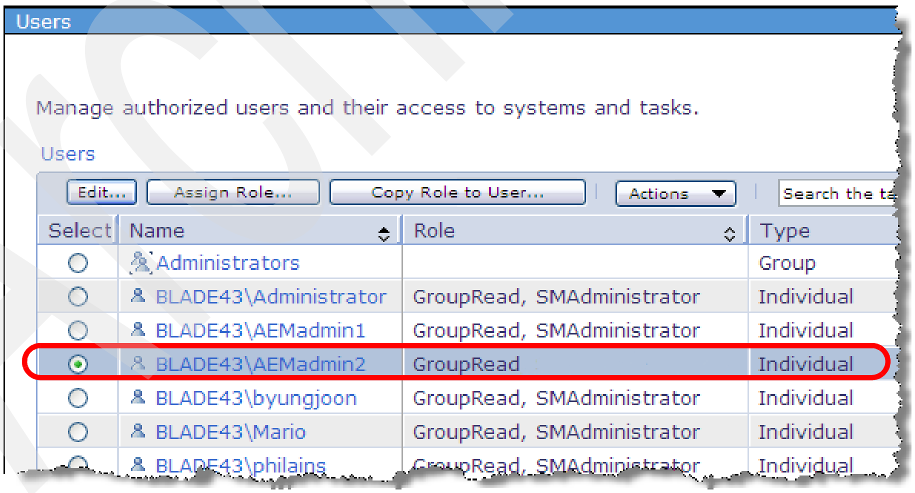


*Figure 4-9 Managing authorized users and their access*

7. The Assign Roles wizard starts. Click **Next** to pass the Welcome panel. In the Roles window, shown in Figure 4-10, you can select the role or roles as appropriate. The *GroupRead* role is a default.

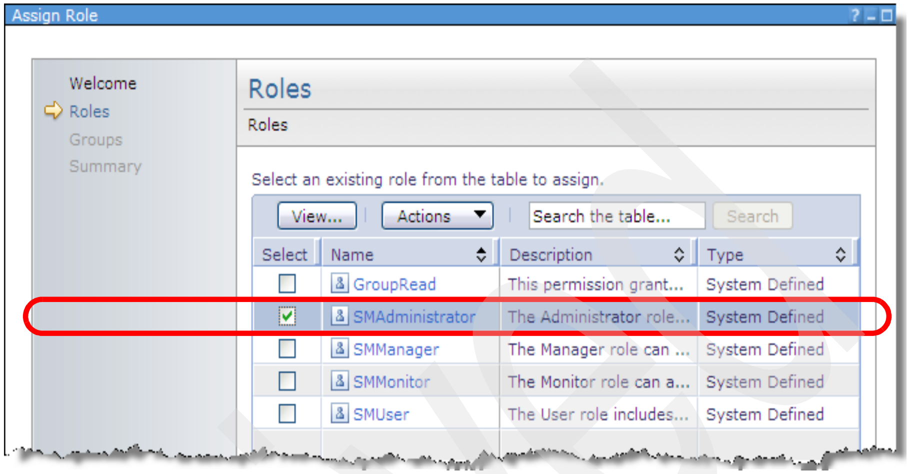


*Figure 4-10 Assigning specific roles*

Available roles include the following categories:

– GroupRead

  Grants a user the ability to view or open a group.

– SMAdministrator

  Has full authority to all resources and tasks, including security administration, product installation, and configuration.

– SMManager

  Can perform a subset of the tasks that an Administrator can perform, which are typically system administration, system health management, and configuration tasks.

– SMMonitor

  Can access those administrative functions that provide read-only access, which are primarily, monitoring, notifications, and status tasks.

– SMUser

  Includes any authenticated user and allows only basic operations, such as viewing resources and properties.

In this example, SMAdministrator is assigned to give full authority to access and manage system resources.

8. As shown in Figure 4-11, an assigned role is only activated under designated groups. To enable access to all resources, click **All Resources**, and then click **Next**. To enable access to only a subset of the available resource groups, click **Selected Groups**, select the group or groups that you want, and then click **Add**. When finished click **Next**.

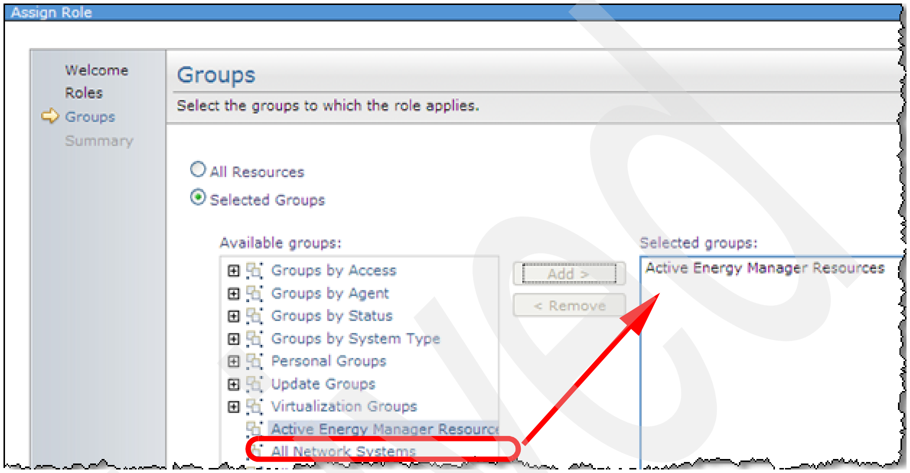


*Figure 4-11 Selecting specific groups*

With these steps complete, although another administrator operates IBM Systems Director server, you can simultaneously log in to IBM Systems Director server with a new account (in this example AEMadmin2).

## 4.5 Features related to Active Energy Manager

Because Active Energy Manager is a plug-in tool that is integrated into IBM Systems Director, it is important to understand the features of Systems Director, such as discovery, events, thresholds, and topology maps, before working with Active Energy Manager.

This section describes the major features of Systems Director that are related to Active Energy Manager.

### 4.5.1 System discovery

Before any system can be managed by IBM Systems Director, it must be discovered. It is also important to gain management access to the system and collect inventory from it before other capabilities in Systems Director can be configured and used to manage it.

*Discovery* is the process by which IBM Systems Director identifies and establishes a connection with objects, such as systems that it can manage. A *system* is one type of resource that Systems Director can manage. A system is an operating-system-based or hardware-based endpoint that has an IP address and host name and can be discovered and managed by Systems Director.

Before Systems Director can manage a system or device, that resource must be discovered by IBM Systems Director Server. IBM Systems Director manages the following types of systems:

- Blade administrative server
- Boot server
- Cluster
- Fabric
- Farm
- Hardware Management Console
- Management controller
- Operating system
- Print server
- SAN
- Server
- Storage system
- Switch
- System chassis

IBM Systems Director provides the following options for discovery:

- System discovery
- Advanced system discovery

> **Note:** Active Energy Manager also uses system discovery and advanced system discovery to discover systems that need power monitoring and management. For more information, refer to 5.2, “Discovering Active Energy Manager endpoints” on page 146.

## Discovery Manager

Discovery Manager is the central repository for all the discovered systems and also serves as the central control point for performing all the discovery and inventory-related actions.

To launch Discovery Manager from the IBM Systems Director Web console, go to **Welcome** → **Discovery Manager** to open the Discovery Manager page, as shown Figure 4-12.

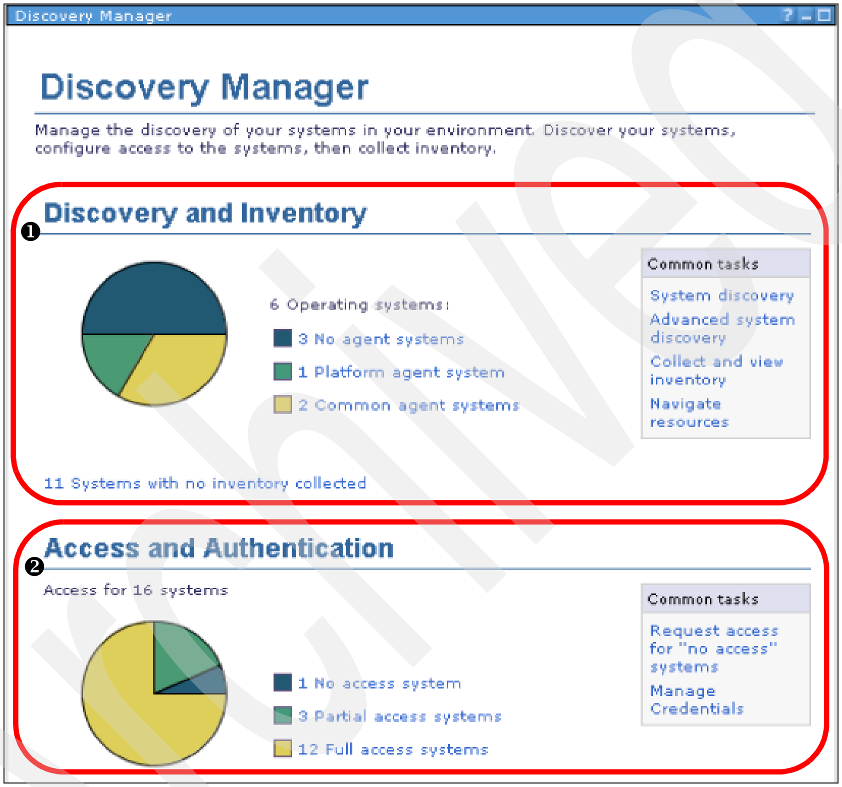


*Figure 4-12 Discovery Manager page*

The Discovery Manager page has two main sections, as shown in Figure 4-12:

**1.** Discovery and Inventory

This section lists all the systems that IBM Systems Director discovers. The discovered systems are classified into three categories based on the agent that is used for discovering:

- No agent systems
- Platform Agent systems
- Common Agent systems

The pie chart in Figure 4-12 shows a visual representation of these classified discovered systems.

You can also launch the following common tasks in this section of the page:

- System discovery
- Advanced system discovery
- Collect and view inventory
- Navigate resources

**2.** Access and Authentication

This section lists all the systems that Systems Director discovers. The discovered systems are classified into three categories:

- No access systems
- Partial access systems
- Full access systems

You can click each classification title to see the list of systems under each respective classification. For example, you can click **No access systems** to see a list of all discovered systems currently having no access, as shown in Figure 4-13.

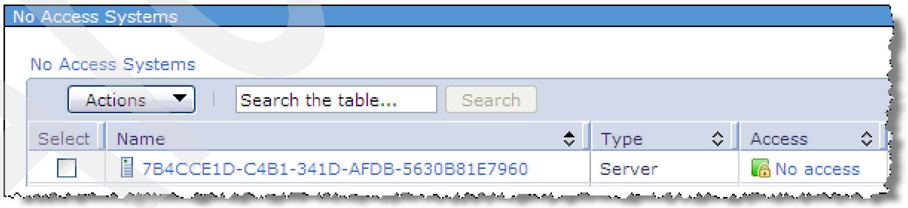


*Figure 4-13 No Access Systems list*

You can also perform the following common tasks for requesting access to a system or configuring the authentication credentials:

- Request access for *no access* systems
- Manage credentials

### 4.5.2 System events

An *event* is an occurrence of significance to a task or system, such as the completion or failure of an operation. IBM Systems Director Server receives events from many sources.

Event sources include, but are not limited to, the following programs and protocols:

- IBM Systems Director native events generated by Common Agent (such as a fan failure or a power supply failure)
- Common Information Model (CIM) indications from the Common Information Model Object Manager (CIMOM) that is installed as part of Common Agent and Platform Agent
- Microsoft Windows event log events
- Windows Management Instrumentation (WMI) events
- SNMP traps through out-of-band communication (communication that is not through Common Agent or Platform Agent)
- Platform Event Traps (PET) through out-of-band communication with systems that support Alert Standard Format (ASF) and Intelligent Platform Management Interface (IPMI) such as the Baseboard Management Controller on IBM System x and BladeCenter Servers
- IBM service processor notifications through out-of-band communication, such as from the Remote Supervisor Adapter II (RSA II) on IBM System x servers

#### Event Automation Plans

One of the most powerful features of IBM Systems Director is the Event Automation Plan tool. This function allows you to react to events in your managed environment in a variety of ways. An *Event Automation Plan* is instructions that which system events need to be acted upon and what actions are to be taken as a result. The Event Action Plan is a combination of two objects: one or more filters used to define exactly what events are to be the trigger, and the one or more actions that are to be performed.

#### Event filters

A *filter* specifies one or more events that you want your Event Automation Plan to process. The Event Automation Plan ignores any event instances that do not meet the specifications of the filter.

In IBM Systems Director Server there are a number of default filters already created, as shown in Figure 4-14.

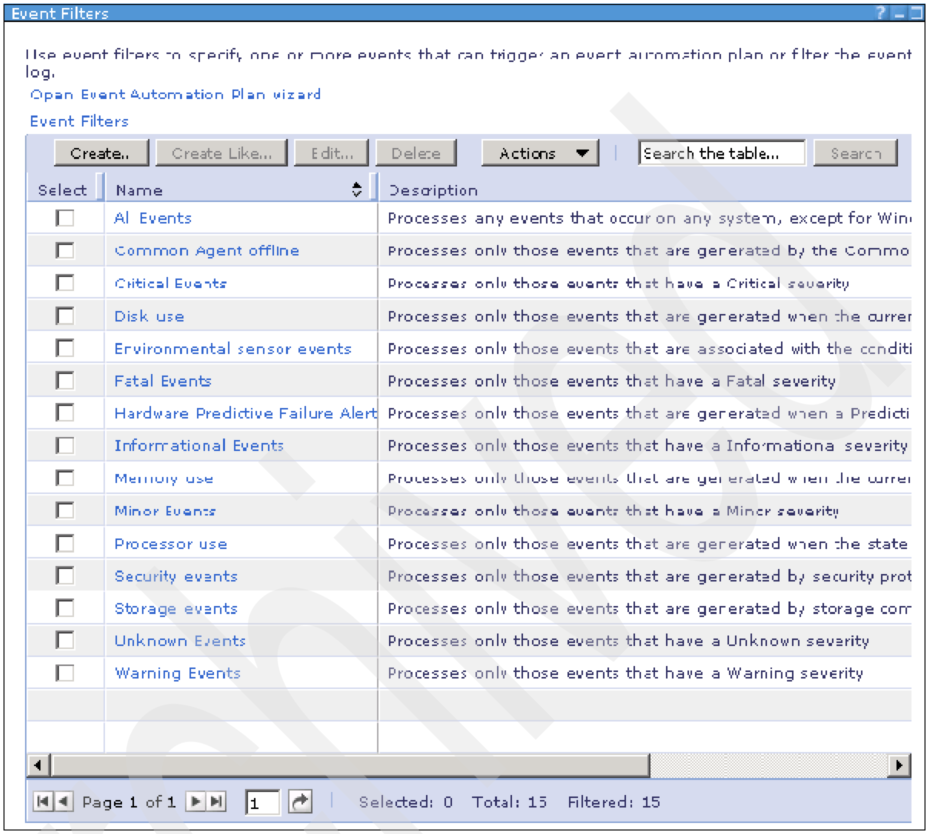


*Figure 4-14 Default Event Filters*

Occasionally, there are situations in which you will want to create a sophisticated event filter. Using these filters, you can specify details for an event such that it covers very specific problems or occurrences. To create filters quickly, default values are provided. However, you can customize the settings.

## Event actions

*Event actions* specify the action that IBM Systems Director takes in response to specific events. For example, using the Event Automation Plan wizard, you can easily create an Event Automation Plan that will send an e-mail or pager notification in response to an event. Additionally, the Event Automation Plan wizard provides other advanced event actions that you can use in response to an event.

IBM Systems Director has several predefined types of event actions, as shown in Figure 4-15. With the exception of *Add to the event log*, you must customize each action that you want to use.

| Name | Type |
|---|---|
| Start a program on a system | Common |
| Start a program on the system that generated the event | Common |
| Send an e-mail to a mobile phone | Common |
| Start a program on the management server | Common |
| Send an e-mail (Internet SMTP) | Common |
| Send an alphanumeric page (using TAP) | Common |
| Static group: add or remove group members | Advanced |
| Post to a newsgroup (NNTP) | Advanced |
| Send an SNMP trap reliably to a NetView host | Advanced |
| Send a Tivoli Enterprise Console event | Advanced |
| Static group: add or remove the event-generating system | Advanced |
| Send an SNMP inform request to an IP host | Advanced |
| Send an SNMP trap to an IP host | Advanced |
| Modify an event and send it | Advanced |
| Timed alarm that starts a program | Advanced |
| Set an event system variable | Advanced |
| Start a task on a specified system | Advanced |
| Timed alarm that generates an event | Advanced |
| Start a task on the system that generated the event | Advanced |
| Log to a log file | Advanced |
| Send a numeric page | Advanced |

*Figure 4-15 Default predefined customizable event actions*

For more detailed information about each of the default actions refer to the IBM Systems Director Information Center:

http://publib.boulder.ibm.com/infocenter/director/v6r1x/index.jsp?topic=/director.automate_6.1/fqm0_c_ea_actions.html

## 4.5.3 Threshold

*Thresholds* are the method that IBM Systems Director uses to generate an alert when a monitor exceeds a limit that you set, either high or low. The Thresholds page, shown in Figure 4-16, shows a list of all thresholds set on resources in your environment.

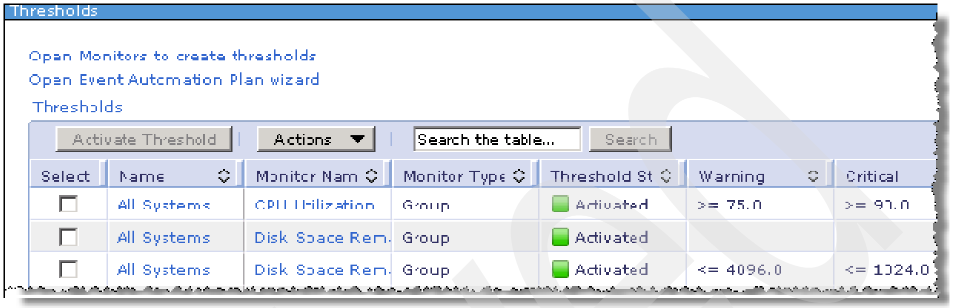


*Figure 4-16 Thresholds page*

Clicking a threshold shows its properties. Also, you can click **Open Monitors to create thresholds** to create additional thresholds. In addition, clicking **Open Event Automation Plan wizard** takes you to the Event Automation Plan wizard, where you can define how events are handled, including those from monitor thresholds.

## 4.5.4 Topology map view

You can use the topology map view to help you visualize the relationships and dependencies between physical and virtual resources. In addition, the topology map includes health status icons to help locate the source of a particular problem.

To access the topology map view, complete the following steps:

1. In the IBM Systems Director Virtualization Manager Web console, navigate to the resource whose relationships you want to see in the topology map view. For example, you can expand **Hardware and Software**, and then select **Virtual Servers and Hosts**.

2. Click **the context menu icon** next to the desired resource, and select **Topology Map**. The topology map view opens, as shown in Figure 4-17.

> **Note:** You can move the blue rectangle (as indicated at the number 1 in Figure 4-17) to navigate the topology map and also to enlarge topology view to look up the systems in detail. For more information, refer to:
>
> http://publib.boulder.ibm.com/infocenter/eserver/v1r2/index.jsp?topic=/eica7/eica7_concept_topology_map_view.htm

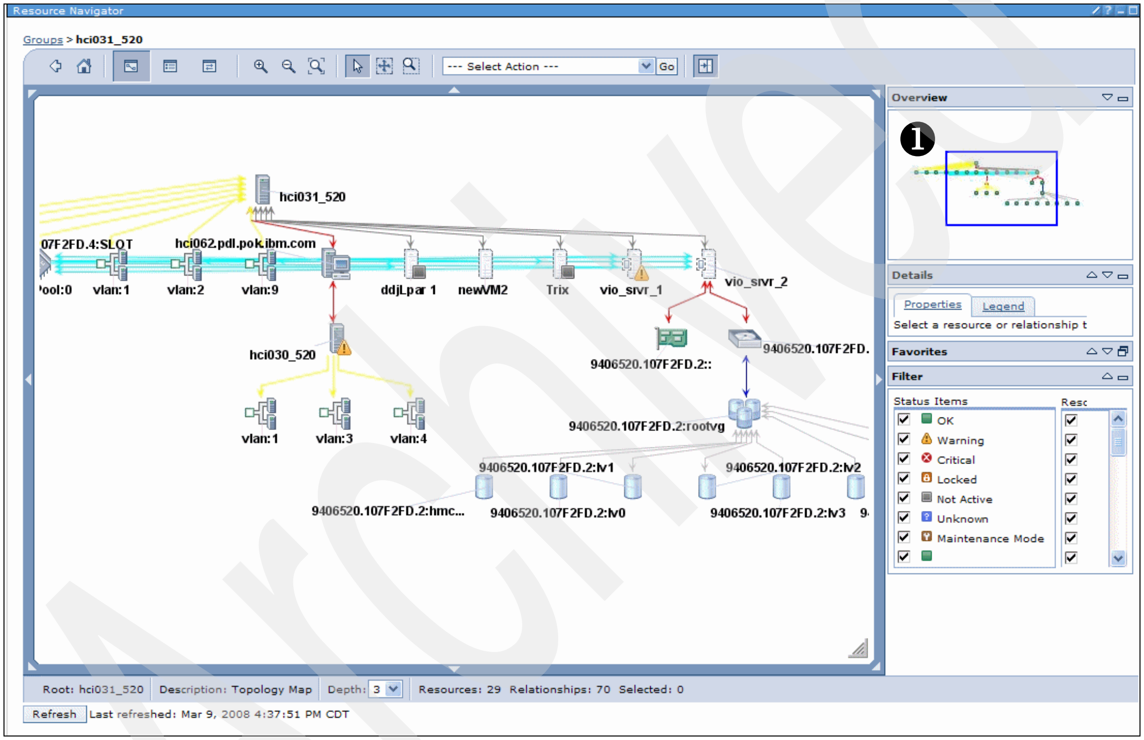


*Figure 4-17 Topology map view*

## 4.6 Using groups

Because Active Energy Manager is integrated with IBM Systems Director as a plug-in tool, it inherits grouping features from Systems Director. Thus, Active Energy Manager provides many different methods to monitor and trend power consumption data using the grouping feature. Understanding how to use groups is important in using Active Energy Manager effectively.

## 4.6.1 The concept of grouping

Groups are a powerful instrument within IBM Systems Director. You can use groups for easier administration, for security purposes, and for event automation. For example, you can use groups to enable a Web administrator to access only a group of Web servers or for the management server to send an e-mail alert to the Wed administrator when an event is received from systems in that group.

We highly recommend that you plan groups properly, especially when using groups for security purposes. If you do not plan attributes thoroughly, it is possible, for example, that a mission-critical database server can be manipulated by someone who does not have the authority.

### Navigating the Groups console

To see current available Groups, click **Navigate Resources** in the Navigation area and Figure 4-18 appears, showing the default groups.

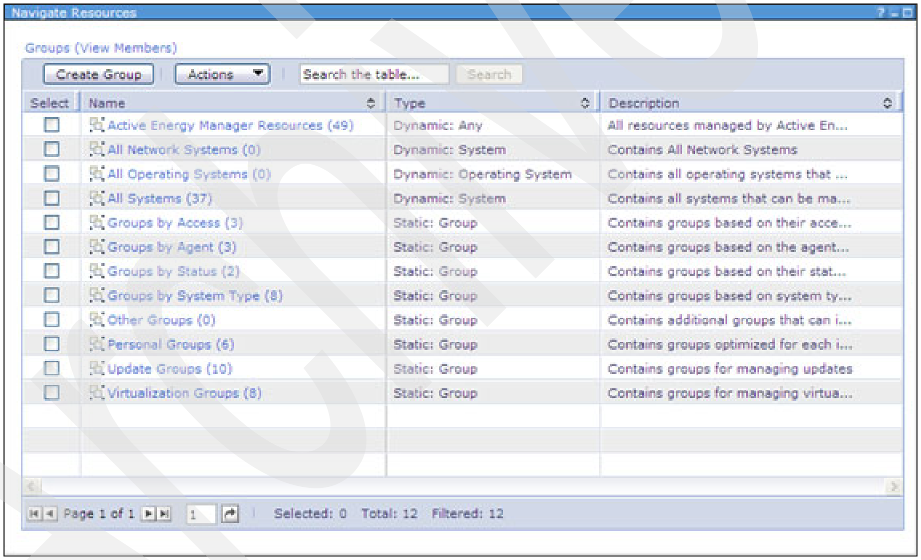


*Figure 4-18 Navigate Resources Groups view*

### Understanding the default groups

When you first start using groups, the discovered resources are categorized and displayed in *default groups*. You can click a group to view subgroups that further categorize the resources for ease-of-use.

> **Note:** You cannot edit or delete default groups.

The default groups are as follows:

- All Operating Systems: All operating systems that can be managed in the IBM Systems Director.
- All Systems: All discovered systems in IBM Systems Director, including servers, chassis, operating systems, switches, blades, and storage systems.
- Groups by System Type: Systems categorized into subgroups by hardware and operating system platform. The available subgroups vary, depending on the plug-ins that are installed in the IBM Systems Director environment.
- Groups by Agent: Systems categorized into subgroups by the degree of management capability available in the IBM Systems Director systems management environment. This default group provides the following subgroups:
  - Systems with No Agent
  - Systems with Platform Agent
  - Systems with Common Agent
- Groups by Status: Dynamic groups that contain systems for which there are unresolved hardware status events. This default group provides the following subgroups:
  - Systems with Problems
  - Systems not in Compliance
- Groups by Access: Dynamic groups that contain systems depending on their current access state. This default group provides the following subgroups:
  - Systems with No Access
  - Systems with Partial Access
  - Systems with Full Access
- Groups with Thresholds: Any groups to which you have applied a threshold.

  **Note:** If you have migrated Event Automation Plans from versions of IBM Director earlier than IBM Systems Director 6.1, Event Automation Plans that use threshold values are migrated to this group.

- Personal Groups: Any groups that you have created or are exclusively associated with your IBM Systems Director user ID. These subgroups include the favorites group.
- Other Groups: Group definitions migrated from versions of IBM Director earlier than IBM Systems Director 6.1.

- Update Groups: All groups provided by update manager. Update groups can be static or dynamic, and both types can be used in compliance policies:
  - Static update groups

    Contain individual updates that were explicitly chosen. When established, the membership changes only when you manually add or delete updates. Static update groups can be used as a baseline for future comparison or update deployment.
  - Dynamic update groups

    Automatically includes updates based on selected update types. The membership of this group changes as update information changes.

> **Note:** The membership of an update group is not resolved at the time that a task using the group is scheduled. The membership of the update group is resolved at the time that the task runs.

- Virtualization Groups: All groups provided by Virtualization Manager.

  IBM Systems Director organizes logical sets of resources into groups. Virtualization Manager provides a set of default or predefined groups for virtual resources.

## 4.6.2 Active Energy Manager resource group

You can use the Navigate Resources task in IBM Systems Director to access the Active Energy Manager Resource group.

As shown in Figure 4-19, when Active Energy Manager is installed, a dynamic group is created containing all discovered resources that Active Energy Manager can monitor and manage. It also includes resources that could be monitored if a firmware update was performed. As new supported resources are discovered, they are added to the group.

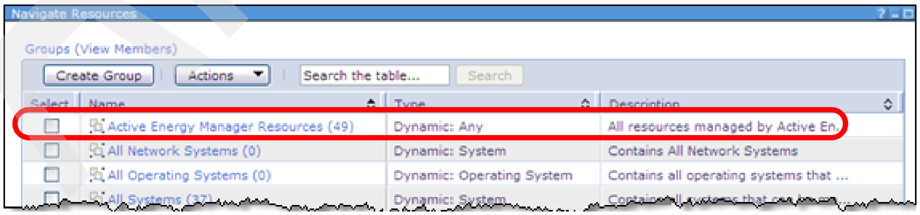


*Figure 4-19 Navigation Resources Group task*

**Tip:** Some discovered resources require a restart of IBM Systems Director before power and thermal data can be displayed in the navigation tables. If you discover a power or thermal metering-capable resource and if you do not see data (or the pending indicator) in power and thermal columns of a navigation table, restart IBM Systems Director Server to see the data.

To access Active Energy Manager resources group using the IBM Systems Director Navigate Resources task:

1. In the left-hand navigation area of IBM Systems Director, click **Navigate Resources.**
2. Select the Active Energy Manager Resources group.

### 4.6.3 Customizing the Active Energy Manager resource group

With IBM Systems Director, you can create a personal group so that you can manage system resources or apply the policies that are followed by each group.

As shown in Figure 4-20 on page 132, you can create a group by joining individual systems but also by integrating groups. This grouping provides the flexibility to monitor and manage power in the data center.

Monitoring by grouping includes the following levels:

- Individual system level grouping
  - Rack 1 includes three IBM System x servers, one IBM Storage system, and one older x86 server
  - Rack 2 includes two External switch, one IBM Power system, one IBM system x server, and one older Storage system
  - Rack 3 includes one x86 server, three IBM System x servers, and two External switch
- Rack level grouping
  - Active Energy Manager Monitoring Group#1 includes Rack 1 and Rack 2
  - Active Energy Manager Monitoring Group#2 includes Rack 2 and Rack 3

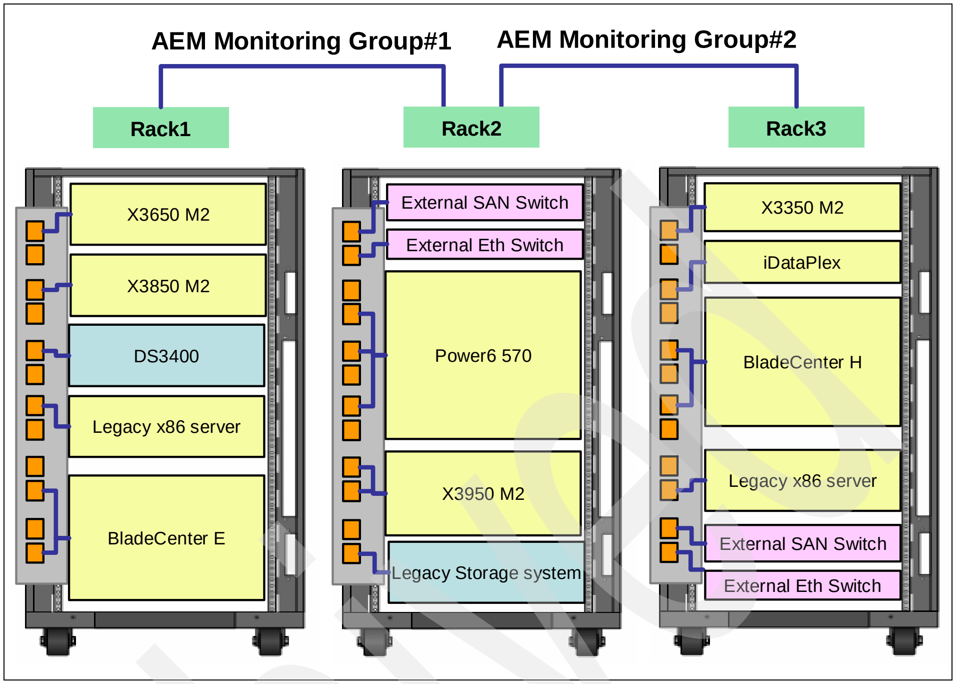


*Figure 4-20 Different level (Individual system and rack) of monitoring by grouping*

To create a group, follow these steps:

1. In the Systems Director navigation area, click **Navigate Resources**, and then click **Create Group** in Navigate Resources.
2. In the Group Editor wizard, the Welcome page displays. Click **Next**. On the Name page, shown in Figure 4-21, enter a unique descriptive name for the group that you are creating. Optionally, you also can enter a description of the group. Click **Next**.

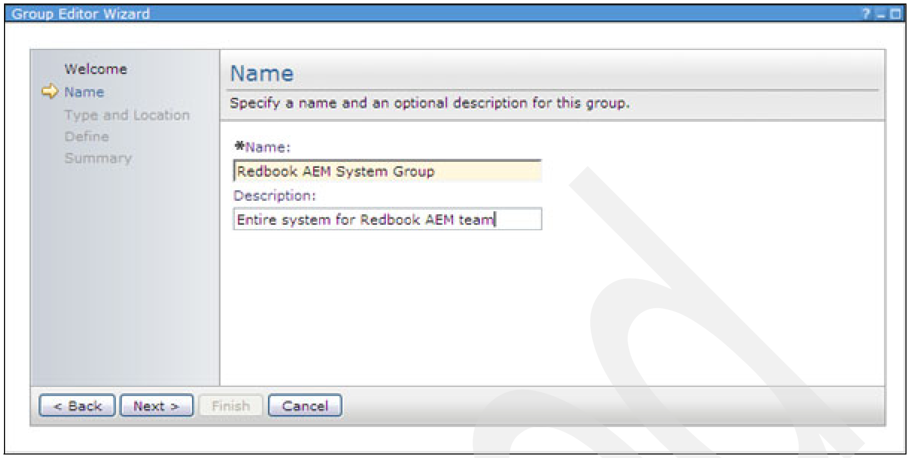


*Figure 4-21 Entering a name an description for a group*

3. On the Type and Location page, select the appropriate Group Type, Member Type, and Group Location from the drop-down menus, as shown in Figure 4-22.

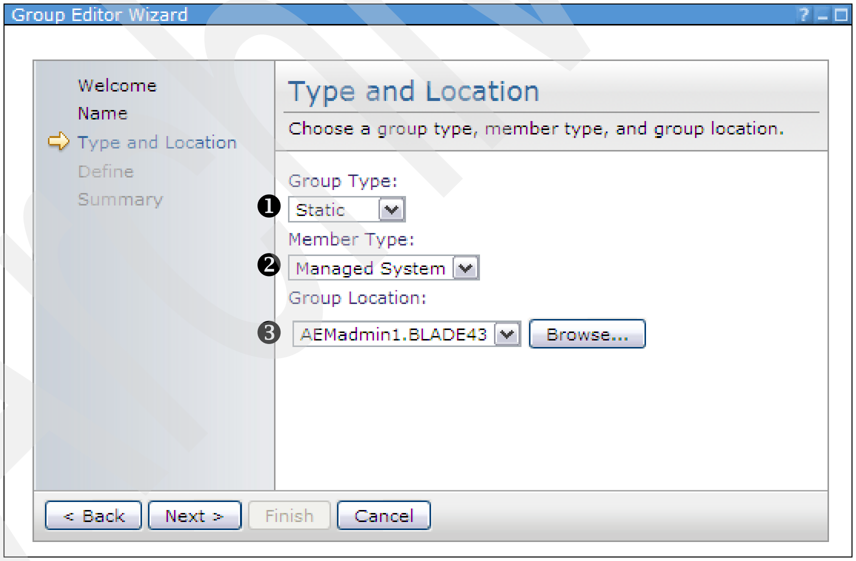


*Figure 4-22 Type and Location*

The fields shown in Figure 4-22 on page 133 are as follows:

**1.** Group Type

You can use groups for organizational purposes and to run tasks on one or more managed resources at a time. A group is either static or dynamic as shown in Figure 4-23.

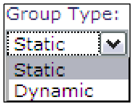


*Figure 4-23 Group Type selection*

You can choose from two types of groups:

- Static group

  *Static groups* are defined by a list of resources. IBM Systems Director Server does not automatically update the contents of a static group. The members of a static group are fixed unless you change

- Dynamic group

  *Dynamic groups* are based on specified system criteria. You can create a dynamic group by specifying criteria that the attributes and properties of the systems must match. IBM Systems Director adds to or removes systems from the group automatically when their attributes and properties change and affect matches to the group criteria.

**2.** Member Type

You can select the type of member that you want to include in the group from the Member Type menu (Figure 4-24). A member type acts as a filter. Only resources of the specified type can be part of the group that you are creating.

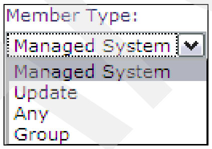


*Figure 4-24 Member Type selection*

Options that you can select in this menu include:

- **Any:** Group membership is unlimited. Any resource can be in the group, including systems, software, and management applications.
- **Managed System:** Group membership is limited to system types such as different type of servers, fabric, farms, hardware control points, controllers, operating systems, chassis, switches, and storage.
- **Update:** Group membership is limited to update types such as for firmware, IBM Systems Director, and operating systems.
- **Group:** Group membership is limited to other existing groups.

**3.** Group Location

From the Group Location list, you can select the parent group that will contain the group that you are creating. In Navigate Resources, a parent group is created and is located under Personal Groups.

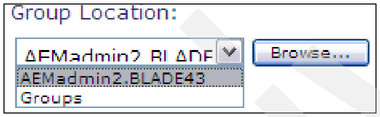


*Figure 4-25 Group Location selection*

Click **Next** to continue.

4. On the Define page, shown in Figure 4-26 on page 136, select one or more groups of resources from the Available list, and click **Add**. You can drill down into a group and select one or more resources. If you want to remove a group or resource, select it from the Selected list, and click **Remove**. Click **Next**.

> **Note:** You cannot add a group's parent to itself. For example, if you define the parent group location for Group1 to be Personal Groups, then you cannot add Personal Groups to Group1.
>
> If you select a resource to add, but the Add button is unavailable, then the selected resource is not a valid selection due to its member type.

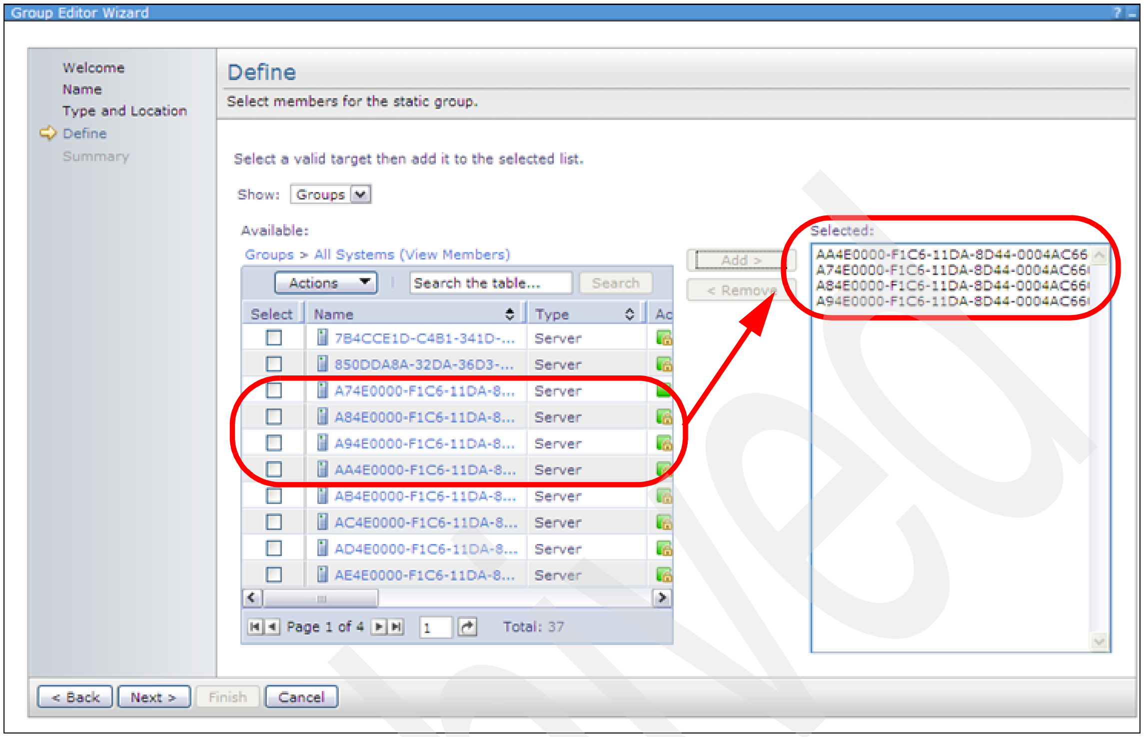


*Figure 4-26 System selection to locate in Group*

5. The static group is created and displays in Navigate Resources as shown in Figure 4-27. A confirmation message about the group creation displays also. Click the static group that you just created (in this example **Redbook AEM System Group**) to manage and apply policies only to the systems under its Group.

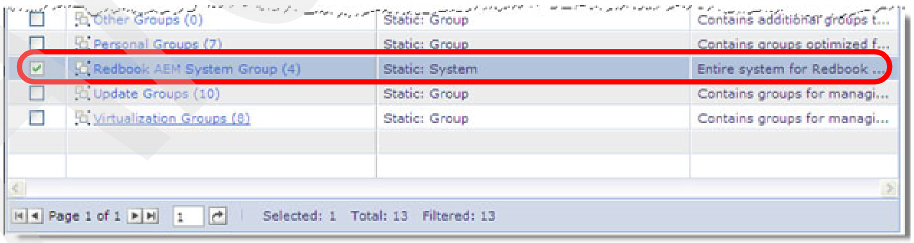


*Figure 4-27 The group is created in the Navigate Resources task*

# 5

# Getting started with Active Energy Manager

This chapter introduces the Active Energy Manager user interface and the discovery of managed resources for monitoring and management. In this chapter, we discuss the following topics:

## 5.1 Active Energy Manager summary page

You can launch the Active Energy Manager summary page from the IBM Systems Director Web console by expanding the **Energy** group and selecting **Active Energy Manager** as shown in Figure 5-1.

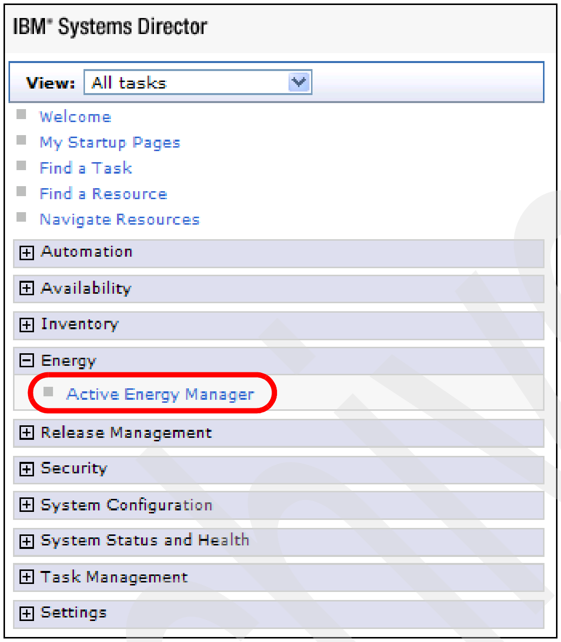


*Figure 5-1 IBM Systems Director Tasks pane*

You can also launch Active Energy Manage from the Manage tab of the Welcome page, as shown in Figure 5-2.


*Figure 5-2 Launching Active Energy Manager*

Figure 5-3 shows the summary page of Active Energy Manager. The Monitor portion of the page is refreshed automatically as changes occur.


*Figure 5-3 Active Energy Manager summary page*

The summary page provides the ability to do the following tasks:

- View the recent status of power and thermal values of the managed resources.
- Navigate the list of resources managed by Active Energy Manager and perform the following tasks against each of the managed objects:
  - View trend data
  - Calculate energy cost
  - View Active Energy monitors
- Set power caps and power savings mode. Configure power and environmental metering devices.
- Create automation plans to take action automatically when a resource reaches a power or environmental threshold or when another Active Energy Manager event is generated.
- View licensing information.

The Active Energy Manager summary page consists of the following sections:

- Status
- Monitor
- Manage
- Automate
- License

In the next sections, we provide a brief introduction to each of these areas of the Active Energy Manager summary page.

## 5.1.1 Status

The Status section of the summary page provides a view of the top five highest average input power values and the top five highest ambient temperatures for Active Energy Manager resources, shown at 1 in Figure 5-4. The names of the resources are links. You can click them to see their properties and to perform actions on the resources.

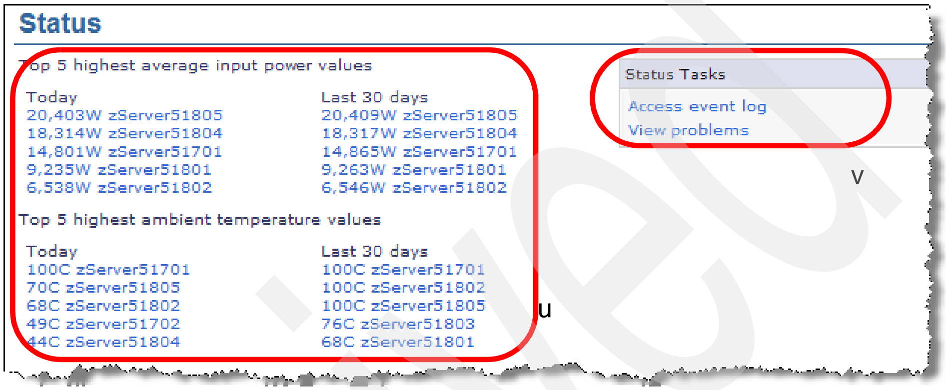


*Figure 5-4 Status section of Active Energy Manager summary page*

In the Status Tasks area, shown at 2 in Figure 5-4, provides the following links:

- Access event log

  Click to access the Event Log page, where you can view events according to filters that you select. By default, all Active Energy Manager events are shown.

- View problems

  Click to access the Problems page, which allows you to view the active problems reported for all discovered resources.

> **Note:** When a resource's power use or temperature has been measured by an external device such as an intelligent PDU or sensor, the name of the resource itself is displayed, not that of the measuring device. However, if power use of a single resource is measured by multiple devices, such as one existing server plugged in to two intelligent PDUs, that device will not appear in the highest input power list. Another feature of the highest input power list is that each BladeCenter's overall power use is considered, rather than that of its modules, such as power domains and blade servers.

## 5.1.2 Monitor

The Monitor section of the summary page provides a list of active power-managed resources, as shown at the number 1 in Figure 5-5.

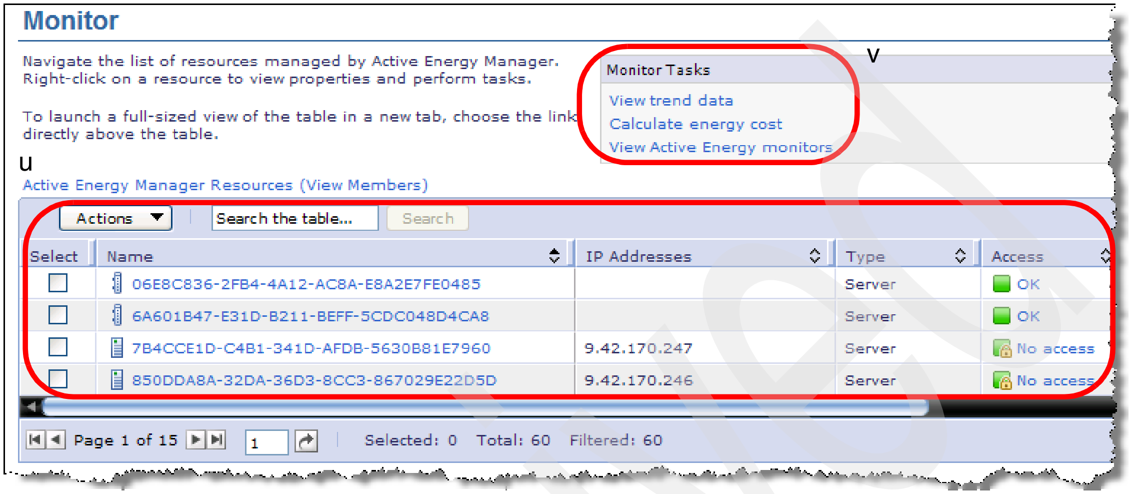


*Figure 5-5 Monitor section in Active Energy Manager summary page*

To launch monitoring tasks, either right-click a resource in the list of Active Energy Manager Resources or use the links in the Monitor Tasks area (at 2 in Figure 5-5). To launch tasks on a resource, you can either use the scroll buttons at the bottom of the table to locate the resource or use the Search option. You can also click the table name link to launch a full-sized view of the table for navigating.

The Monitor Tasks area provides the following links:

- View trend data

  Click to view power and environmental trend data over a specified period of time.

- Calculate energy cost

  Click to calculate energy costs for specified resources.

- View Active Energy monitors

  Click to access the Monitors page, which allows you to monitor power, environmental and CPU values. Thresholds can also be activated to trigger an event whenever a monitored value reaches a specified value.

## 5.1.3 Manage

The Manage section of the summary page allows you to configure the power cap and power savings options, as well as configure power and environmental metering devices and work with power policies. This section also points out the number of power management functions that are used within last 24 hours, as shown at **1** in Figure 5-6.

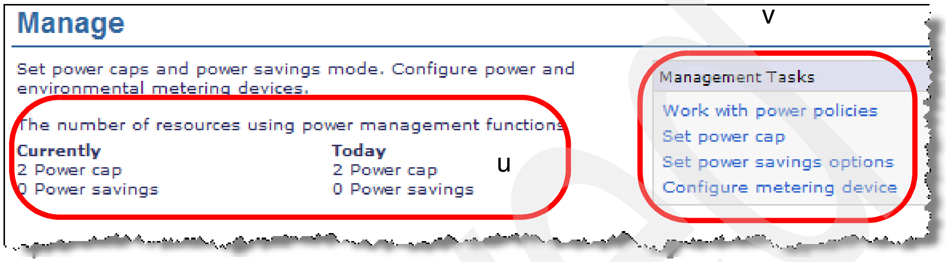


*Figure 5-6 Manage section in Active Energy Manager summary page*

The Management Tasks area, shown at **2** in Figure 5-6, provides the following links:

- Work with power policies

  Click to view, create, edit, or apply power policies.

- Set power cap

  Click to specify the power cap setting on resources that support power capping. You can also disable the power cap capability on a resource.

- Set power savings options

  Click to specify the power savings option on resources that support power savings.

- Configure metering device

  Click to configure metering devices, such as a sensors and PDUs, to associate them with devices. This allows the metering device to monitor and manage data (power, temperature, humidity, and dew point) for resources that are associated with the metering device.

## 5.1.4 Automate

The Automate section of the summary page allows you to specify actions to take in response to power events, as well as view the number of thresholds currently configured. The Automation Tasks area, as shown at 1 in Figure 5-7, provides the following links:

- Create automation plans

  Click to automate energy-related tasks. You can create an IBM Systems Director event automation plan to define energy event criteria (filters) to trigger appropriate event actions.

- Manage thresholds

  Click to create, edit, and activate thresholds.

For more information about how to create automation plans and how to configure thresholds, see Chapter 6, “Monitoring Active Energy Manager resources” on page 179.

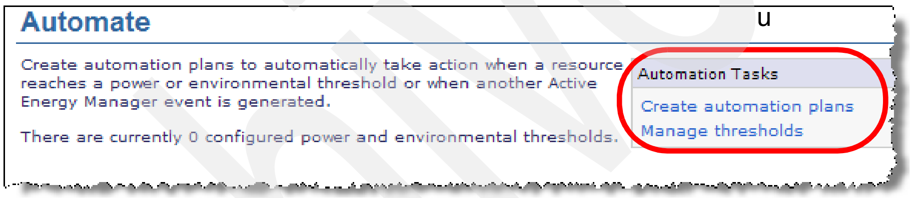


*Figure 5-7 Automate section in Active Energy Manager summary page*

## 5.1.5 License

The License section of the summary page provides status information about the Active Energy Manager license. If you use a trial version of Active Energy Manager, a link displays that points to the Active Energy Manager home page, so that you can purchase a full version license (as shown in 1 in Figure 5-8). The License section also displays the number of times that the power management functions, such as power capping and power savings, are used within the past 24 hours.

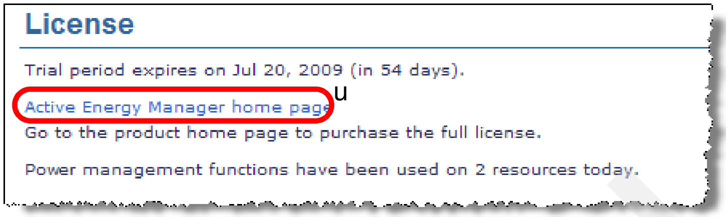


*Figure 5-8 License section in Active Energy Manager summary page*

For more information about Active Energy Manager licenses, see 2.5, “Licensing” on page 64.

## 5.2 Discovering Active Energy Manager endpoints

Active Energy Manager *managed endpoints* (or just *endpoints*) can be thought of as gateways to the Active Energy Manager supported *resources*, whose power consumption can be monitored and managed by Active Energy Manager. Active Energy Manager communicates to each endpoint to obtain the power consumption or thermal state details and to set power management features of the supported resources.

Active Energy Manager supported resources are those systems whose power consumption monitoring and management, as well as thermal monitoring, can be performed by Active Energy Manager. Table 5-1 lists the resources that Active Energy Manager can monitor and manage.

To monitor and manage any Active Energy Manager supported resources, you need to complete these steps:

1. Discover the Active Energy Manager endpoint in the IBM Systems Director.
2. Gain management access to the Active Energy Manager endpoints**.**

You can discover the Active Energy Manager endpoints using either system discovery or advanced system discovery, which we discuss in the following sections:

- 5.2.1, “System discovery” on page 149
- 5.2.2, “Advanced system discovery” on page 155

Before you connect to an Active Energy Manager endpoint, you need to understand the Active Energy Manager supported endpoints as well as the IBM systems that can be managed by each endpoint. Table 5-1 lists the different types of Active Energy Manager endpoint types.

*Table 5-1 Adding Active Energy Manager managed resources*

| Active Energy Manager supported resources | Active Energy Manager endpoint | Discovery resource type | Discovery method (system or advanced system) |
|---|---|---|---|
| System x server | BMC or RSA II[1] | Server | Either |
| System i or System p server | HMC | Server | Either |
| System i or System p server | FSP | Server | Either |
| System z server | HMC | Server | Either[2] |
| BladeCenter chassis | Management Modules | BladeCenter chassis | Either |
| Blade server (x86, Power) | BladeCenter chassis | Not applicable | Not applicable |
| PDU | PDU itself | Power unit | Either[2] |
| SynapSense SNMP Agent and sensor nodes | SynapSense SNMP Agent | Generic system (Generic network device) | Either[2] |
| Various resources (including PDU, power supplies, and sensors) managed by third-party energy management software or devices | Third-party energy management software or device (SiteScan, for example) | Not applicable | Not applicable |
| Emerson-Liebert SiteScan instance[3] | SiteScan software | Operating system | Either |
| Eaton Power Xpert Reporting Database server[3] | Eaton software | Operating system | Either |
| APC InfraStruXure Central server | APC software | Operating system | Either |
| Smart Works gateway | Smart Works software | Operating system | Either |

| Active Energy Manager supported resources | Active Energy Manager endpoint | Discovery resource type | Discovery method (system or advanced system) |
|---|---|---|---|
| Sensatronics, 1-wire, or iButton sensor networks | Sensatronic software | Operating system | Either |

1. Active Energy Manager 4.1.1 now supports the RSA II service processor.
2. You can use basic system discovery only if you are using SNMP v1 or v2c and only if the read community name (and write community name for PDUs) is `public`. Otherwise, you need to use advanced system discovery.
3. Discovery of Emerson-Liebert SiteScan and Eaton Power Xpert Reporting System instances might cause the discovered operating system to have an OS type of *Windows* instead of *Appliance*.

Power Xpert Reporting System considerations:

- The Microsoft SQL Server JDBC driver is required for Active Energy Manager to process Power Xpert Reporting system database files contained on an instance of Microsoft SQL Server. See 10.3.3, “Prerequisites before discovering Power Xpert” on page 466 for details.

  In addition, if you are using Windows integrated authentication, copy the `sqljdbc_auth.dll` file to a directory on the Windows system path on the server where IBM Systems Director is installed.

- A default properties file, `aem_powerxpert.properties`, is created in the `/Director/lwi/conf` directory when Active Energy Manager is installed. This file contains properties that control the discovery process for Power Xpert Reporting System database servers.

In terms of the number of IBM Systems Director server connections that you can have to the different Active Energy Manager endpoints, be aware of the following situations:

- The default number of IBM Systems Director server connections you can have to a BladeCenter management module, RSA, or BMC defaults to one. By using the management interface to the BladeCenter management module, RSA, or BMC, you can increase the limit for the discovery and management functions of IBM Systems Director.
- By default you can have multiple IBM Systems Director server connections to an FSP, HMC, or zHMC.
- By default you can have multiple IBM Systems Director server connections to a PDU+, but we recommend only one. This is because every time an Active Energy Manager server connects to a PDU+ it resets the start time for collecting power information. This can lead to misleading results. For example, if Active Energy Manager server A polls a PDU+ every 4 minutes, and Active Energy Manager server B polls the same PDU+ every minute, then

every time server A reads the power information, it sees only a maximum of the previous minute's worth of data, because server B has reset the PDU+ within the previous minute.

- To add and work with Emerson Network Power and Liebert devices in Active Energy Manager, the user ID used to access SiteScan Web must have `Remote Data Access - SOAP` functional privilege.

### 5.2.1 System discovery

To discover manageable endpoints at a specific network address or range of addresses, use system discovery. Refer to Table 5-1 on page 147 to determine the Active Energy Manager endpoints that can be discovered using the system discovery method.

You can use any of the following methods to discover Active Energy Manager endpoints, which can be discovered using the system discovery method:

- Discovery based on a single IP address
- Discovery based on a range of IP addresses
- Discovery based on a host name

The following sections discuss the steps to discover Active Energy Manager endpoints through the system discovery method.

### Discovering the server resource type Active Energy Manager endpoints

Refer to Table 5-1 on page 147 to determine the server resource type Active Energy Manager endpoints.

Follow these steps to discover server resource type Active Energy Manager endpoints in IBM Systems Director by system discovery:

1. Click **System discovery** under **Optional tasks** on the IBM Systems Director Welcome page, as shown in Figure 5-9.


*Figure 5-9 IBM Systems Director welcome page*

2. In the System Discovery pane as shown in Figure 5-10, select **Single system (IP address)**. Input the IP address and host name of the Active Energy Manager endpoint that you need to discover. Refer to Table 5-1 on page 147 to determine the resource type that you need to select to discover the Active Energy Manager supported resource in IBM Systems Director. Select the appropriate resource type in the drop-down menu (in this case, **Server**). Then, click **Discover**.

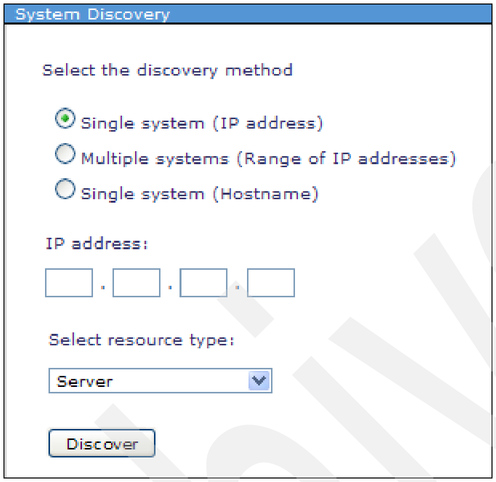


*Figure 5-10 System discovery by IP address and host name*

3. When the system is discovered, it is shown in the Discovered Systems table, as shown in Figure 5-11.

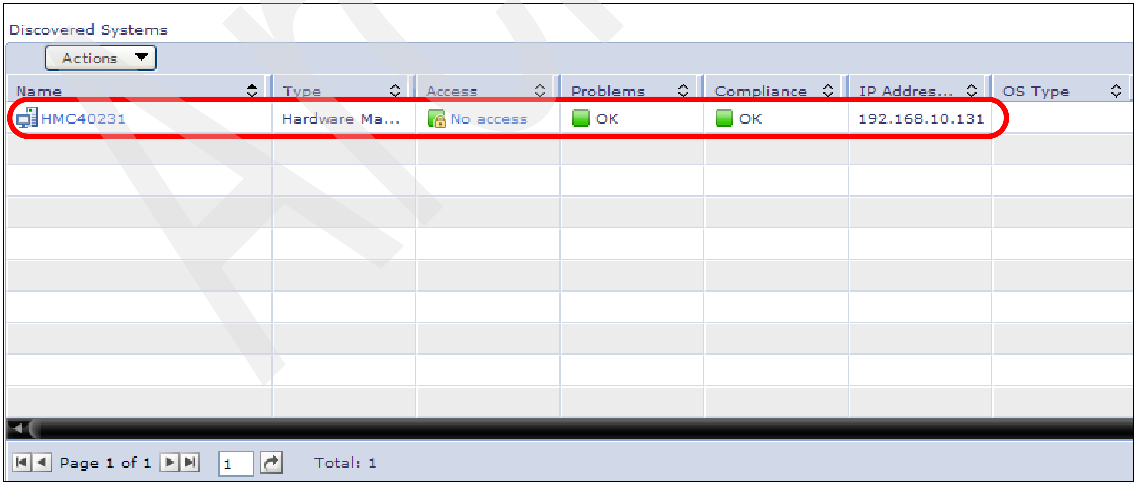


*Figure 5-11 Discovered Systems table*

When you discover any system in IBM Systems Director, by default it has no access. You can check the access status in the Discovered System table under the Access column for each respective systems. For more information about authenticating to Active Energy Manager endpoints, refer to 5.3, “Authenticating to endpoints” on page 172.

### Discovering PDUs using system discovery

Follow these steps to discover the Active Energy Manager supported resources using a range of IP addresses:

1. Click **System discovery** under **Optional tasks** from the IBM Systems Director Welcome page, as shown in Figure 5-9 on page 150.
2. In the System Discovery pane, shown in Figure 5-12, select any of the discovery methods. In this example, **Multiple systems (Range of IP addresses)** is selected, in which case inputting the range of IP addresses of the PDUs is required. Select **Power Unit** as the resource type, and then click **Discover**.

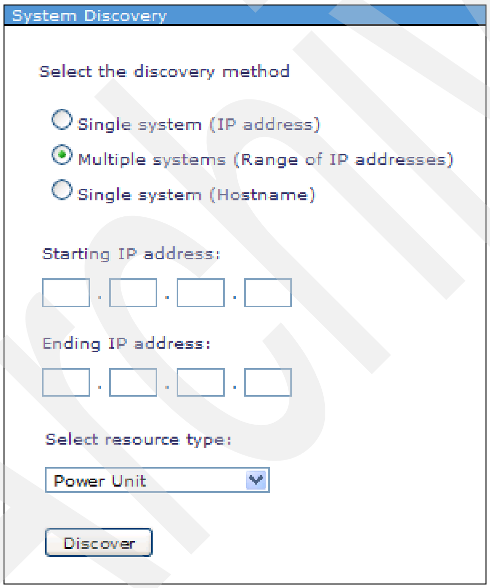


*Figure 5-12 Discovering PDU devices using a range of IP addresses*

3. When the PDU units are discovered, they are shown in the Discovered Systems table, as shown in Figure 5-13.

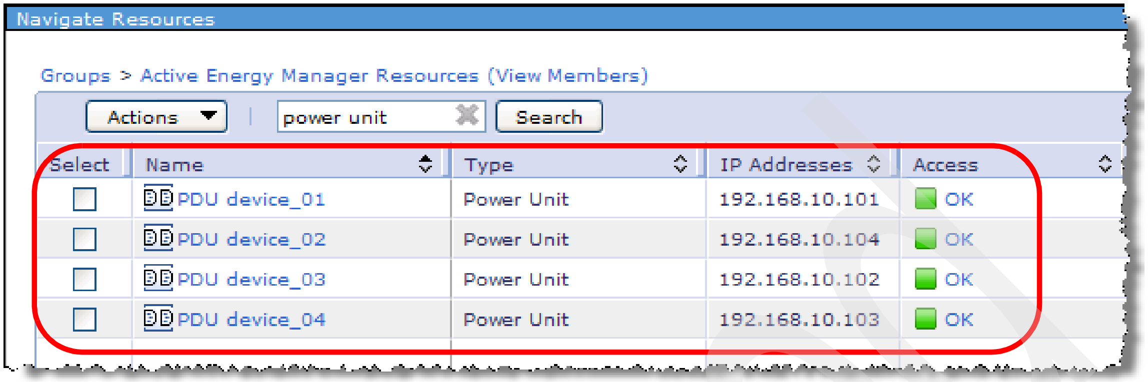


*Figure 5-13 Discovered systems table*

> **Note:** PDU devices that are configured with SNMP v1/2c and a community name of `public` can be discovered using System Discovery. PDU devices that are configured with SNMP v1/v2c with a community name other than public or when SNMP v3 is required for authentication need to be discovered using Advanced System Discovery. For more information about discovering PDUs using Advanced System Discovery, refer to "Discovering PDU devices using advanced system discovery" on page 156.

### Discovering BladeCenter chassis using system discovery

To discover the Active Energy Manager supported resources based on host name, follow these steps:

1. Click **System discovery** under **Optional tasks** in the IBM Systems Director Welcome page as shown in Figure 5-9 on page 150.
2. In the System Discovery pane, shown in Figure 5-14, select **Single system (Hostname)**. Input the host name of Active Energy Manager supported resource that you need to discover. Refer to Table 5-1 on page 147 to determine the resource type to discover the Active Energy Manager supported resource in IBM Systems Director. Select the appropriate resource type in the drop-down menu, and then click **Discover**.

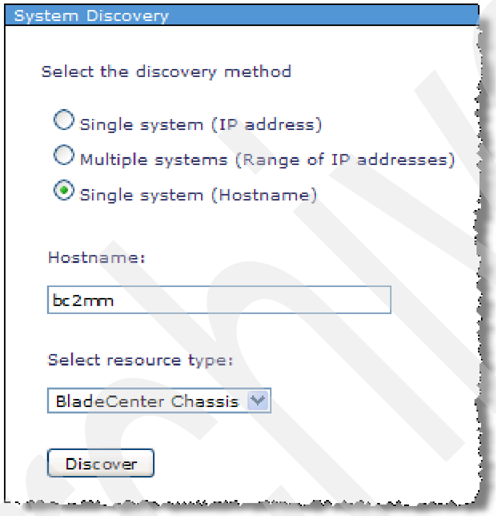


*Figure 5-14 System Discovery: Discover by host name*

3. When the systems are discovered, they are shown in the Discovered Systems table, as shown in Figure 5-15.

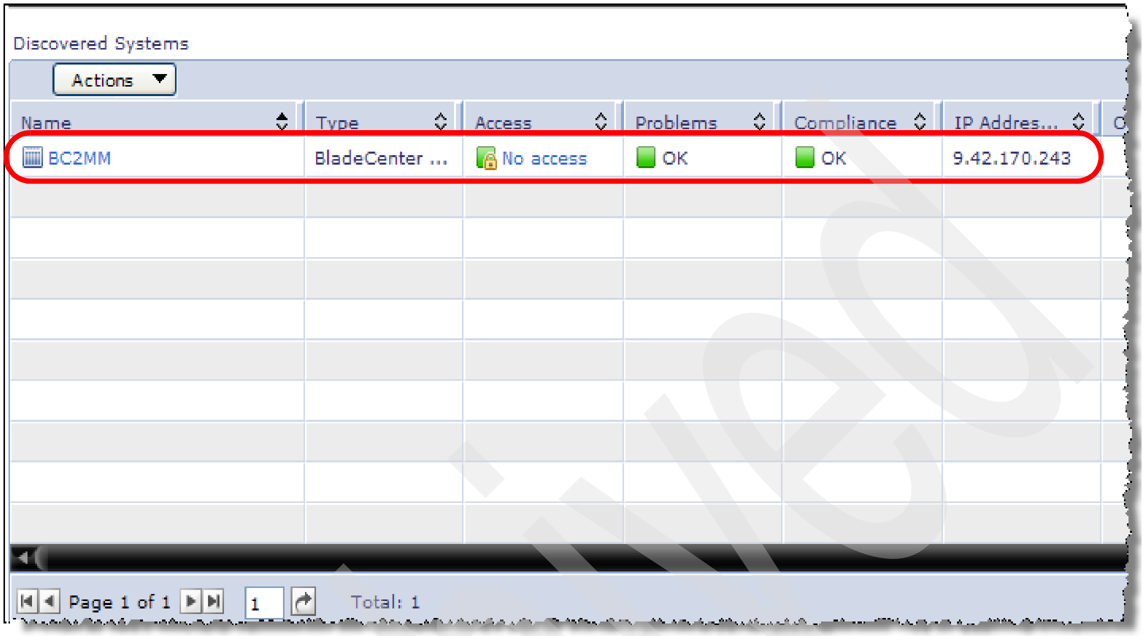


*Figure 5-15 Discovered Systems table*

After discovery, the BladeCenter chassis is in lock mode by default, and authentication is required before it can be managed by Active Energy Manager. For more information about authenticating to Active Energy Manager endpoints, see 5.3, “Authenticating to endpoints” on page 172.

## 5.2.2 Advanced system discovery

You can use advanced system discovery to discover a specific type of resource. Some Active Energy Manager supported resources cannot be discovered using the system discovery method, including:

- Supported PDU devices that do not have default settings for SNMP version and a default community name of `public`
- SynapSense SNMP agents that do not have default settings for SNMP version and a default community name of `public`
- System z HMC

Use the advanced system discovery option to discover these resources in IBM Systems Director. For more information about the type of discovery method to use for each type of Active Energy Manager supported resource, refer to Table 5-1 on page 147.

Advanced system discovery uses discovery profiles to manage the discovery tasks that you want to run. Advanced system discovery provides a list of default discovery profiles. Using the Advanced Discovery Wizard you can:

- Create, copy, edit, or delete discovery profiles.
- Run a discovery profile on specific resources.
- Schedule a profile to run at specific times or in response to specific events.

### Discovering PDU devices using advanced system discovery

> **Tip:** PDU devices that are configured with SNMP v1/2c and a community name of `public` can be discovered using system discovery.

PDUs that are configured with SNMP v2c or v3 or a write community name other than `public` can be discovered only using advanced system discovery, as follows:

1. Click **Advanced system discovery** in the Welcome page of IBM Systems Director, as shown in Figure 5-16.


*Figure 5-16 IBM Systems Director Welcome page*

2. Click **Create** to create a new profile for PDU units discovery in IBM Systems Director, as shown in Figure 5-17.

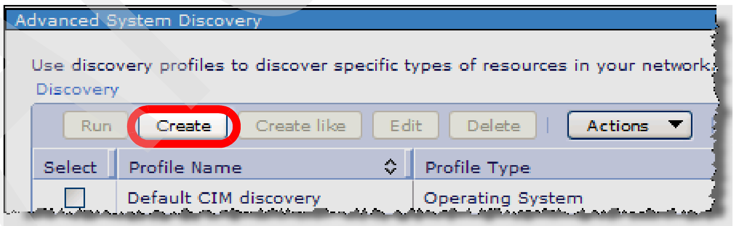


*Figure 5-17 Advanced System Discovery panel*

3. Click **Next** in the Welcome page of Advanced System Discovery Wizard, shown in Figure 5-18, to start the creation process for the discovery profile. Enter the following details in the Profile properties dialog box:

   a. Enter the Profile Name and Profile description for the new profile.

   b. Select the appropriate System type. Refer to the Table 5-1 on page 147 to determine the type of system and subsystem that you need to select to discover the managed resource. For PDU units, the system type is **Power unit**.

   c. Click **Next**.

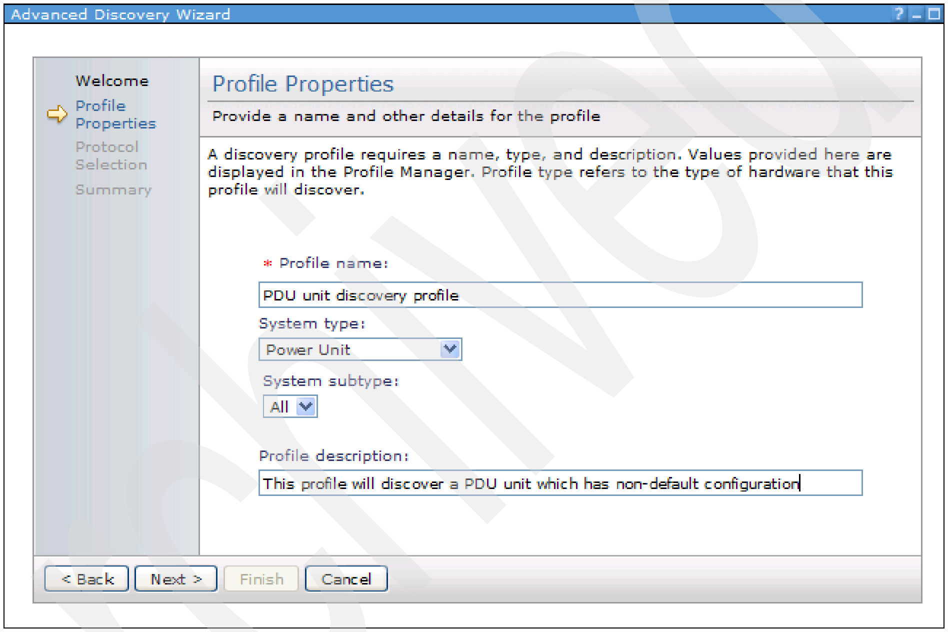


*Figure 5-18 Advanced System Discovery Wizard: Profile properties*

4. Select the communication protocols that need to be used during the discovery process. In this example, SNMP Protocol is used for discovering and managing the PDU unit from Active Energy Manager. Select **SNMP power unit discovery**, and click **Next**.

5. Input the IP address that is used for SNMP discovery. You can choose to input the IP addresses using one of the following methods:
   - Add a single IP address: Use this method when you want to discover a single Active Energy Manager managed resource with this profile.
   - Add a range of IP addresses: Use this method when you want to discover multiple Active Energy Manager managed resources whose IP addresses fall in the same range.
   - Import: Use the import option to list IP addresses from a text file. Use this method when you want to discover multiple Active Energy Manager managed resources that have IP addresses that are not in the same range.

   Click **Next**.

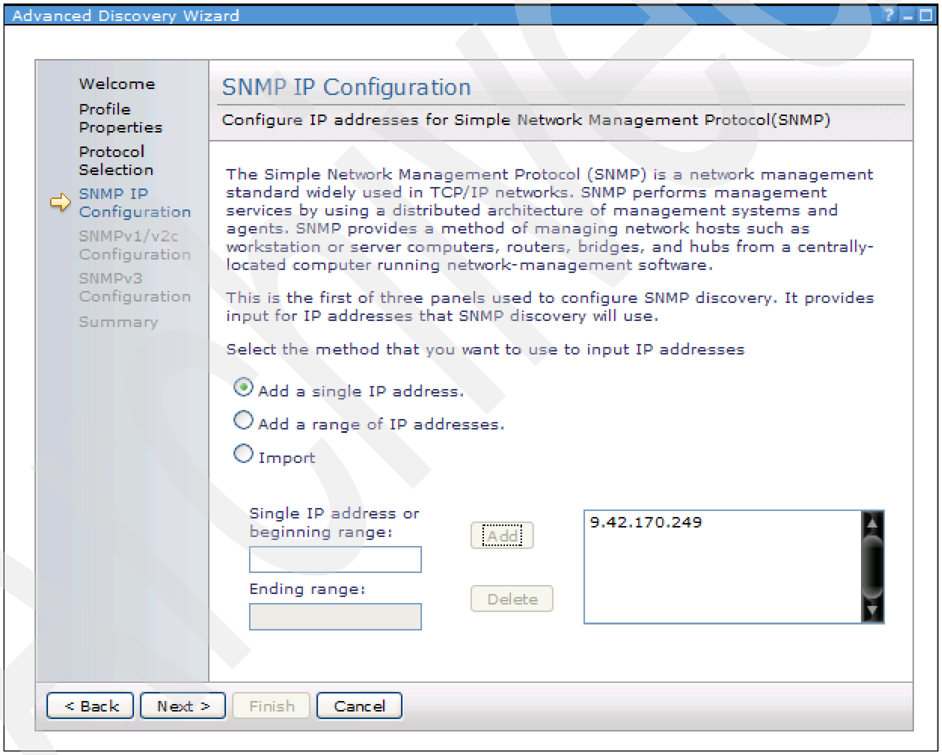


*Figure 5-19 Advanced Discovery Wizard: SNMP IP Configuration*

6. On the SNMP v1/v2c configuration page, shown in Figure 5-20, if the PDU device is configured with SNMPv1/v2c, select the appropriate version of SNMP to discover the PDU unit. Enter the SNMP protocol write community name that is configured in the PDU unit. IBM Systems Director uses this community name to authenticate to PDU unit during the discovery process. Click **Next**.

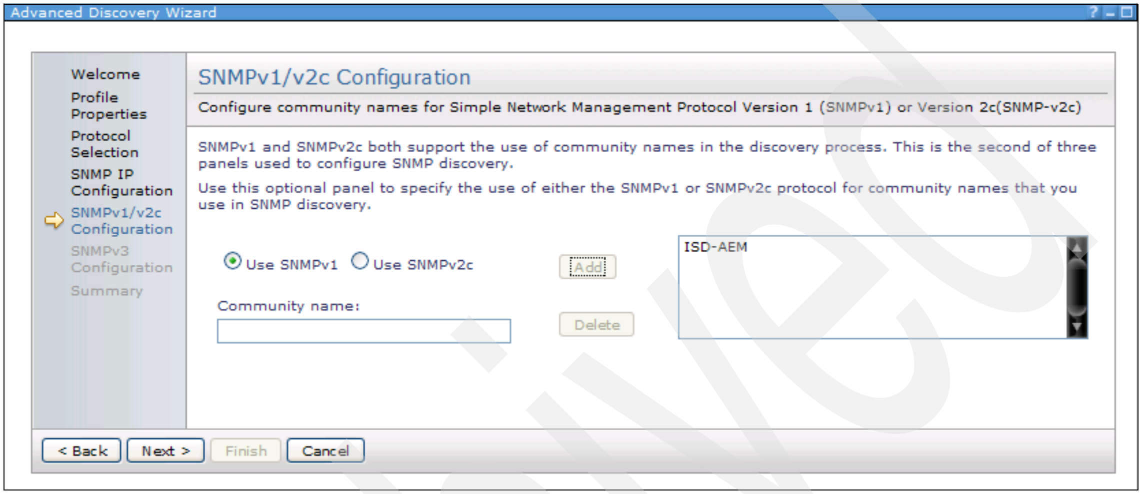


*Figure 5-20 Advanced Discovery Wizard: SNMP 1/v2c configuration*

7. In the SNMP v3 configuration page, shown in Figure 5-21, if the PDU device is configured with SNMP v3, click **Create** to create the SNMP V3 profile used to authenticate to the PDU device during the discovery process. Click **Next**.

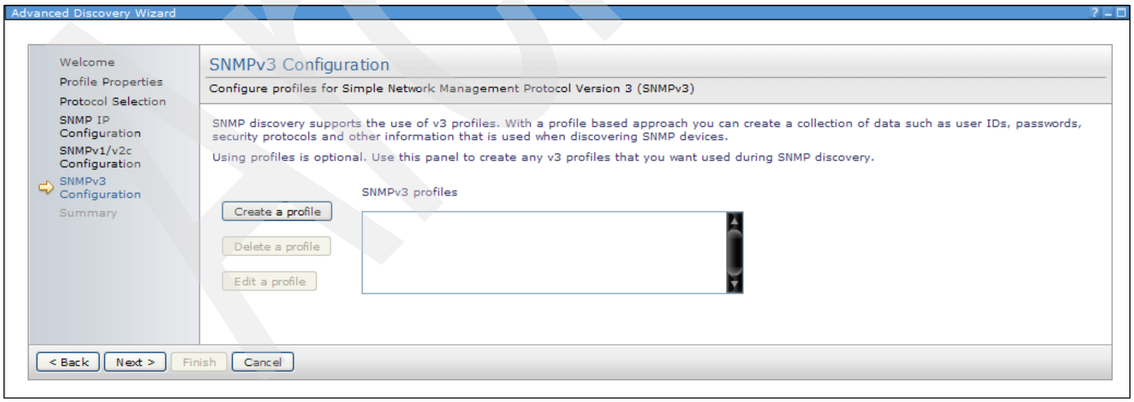


*Figure 5-21 Advanced Discovery Wizard: SNMPV3 Configuration*

8. The Summary page of the Advanced System Discovery Wizard then shows the summary information of the profile selection. Review the summary, and then click **Finish** to create a new profile.
9. Select the new profile created in the Advanced System Discovery page, shown at 1 in Figure 5-22. Click **Run** (at 2) to execute the selected profile.

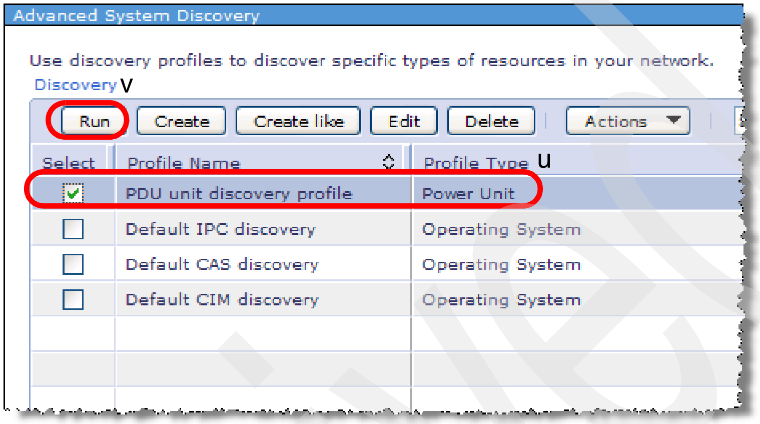


*Figure 5-22 Advanced System Discovery page*

10. As shown in Figure 5-23, you can choose to schedule the Advanced System Discovery task as well as set notification options.

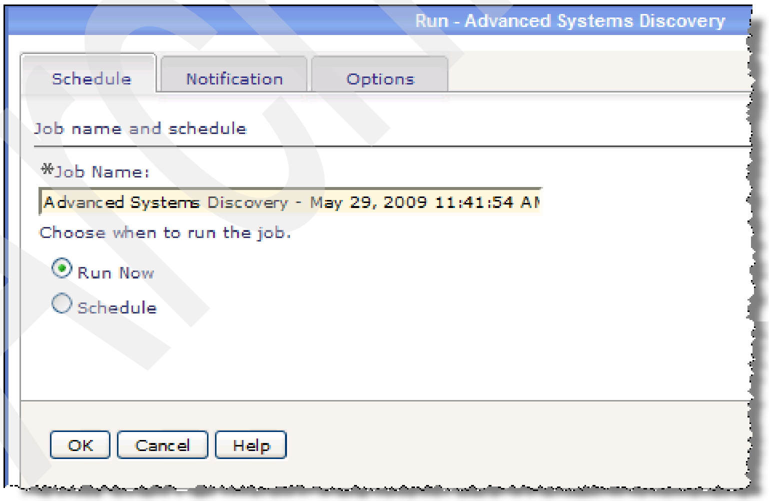


*Figure 5-23 Run Now: Advanced System Discovery*

11. The discovered PDU device displays in the Active Energy Manager resources group, as shown in Figure 5-24.

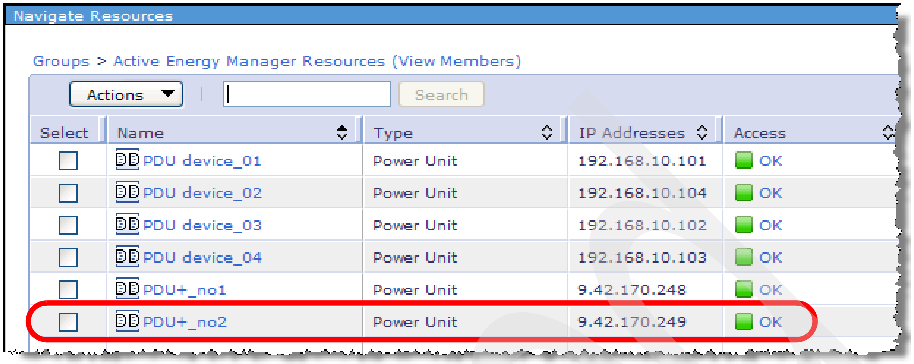


*Figure 5-24 Active Energy Manager Resources (view/Members)*

## Discovering SynapSense sensor networks using advanced system discovery

Follow these steps to discover SynapSense SNMP agents in IBM Systems Director:

1. Click **Advanced system discovery** in the Welcome page of IBM Systems Director as shown in Figure 5-16 on page 157.
2. Click **Create** to create a new profile for SynapSense SNMP agent discovery in IBM Systems Director as shown in Figure 5-17 on page 157.
3. Click **Next** on the Advanced System Discovery Wizard Welcome page to start the creation of the discovery profile.

4. Provide the name and other details for the profile as shown in Figure 5-25. Click **Next**.

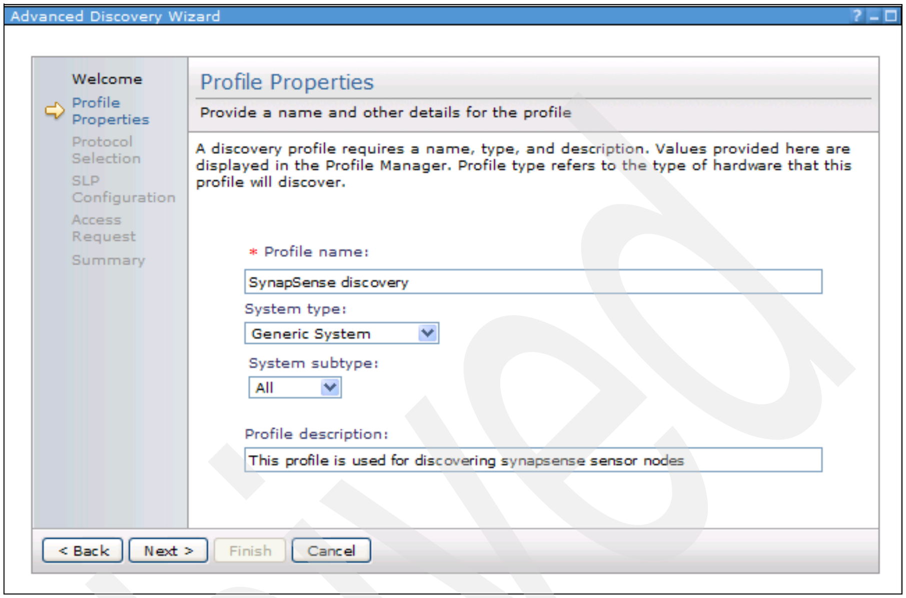


*Figure 5-25 Advanced System Discovery Wizard: Profile properties*

5. Select the protocol that this profile needs to use as follows:

   – SNMP Energy Device Discovery: This protocol is used to discover SNMP devices that should be discovered by Active Energy Manager, like SynapSense SNMP agents.

   – SNMP Device Discovery: This protocol is used to discover SNMP devices that should be discovered by IBM Systems Director.

   We selected **SNMP Energy Device Discovery** as shown in Figure 5-26. Click **Next**.

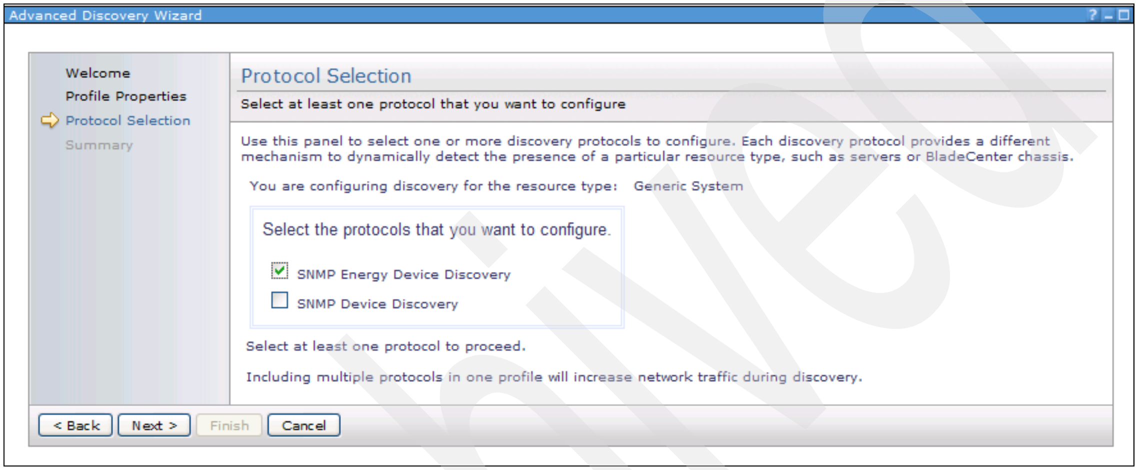


*Figure 5-26 Advanced System Discovery Wizard: Protocol Selection*

6. Input the IP address that is used for SNMP discovery. You can choose to input IP addresses using one of the following methods:

   – Add a single IP address: Use this method when you want to discover single Active Energy Manager managed resource using this profile.

   – Add a range of IP addresses: Use this method when you want to discover multiple Active Energy Manager managed resources whose IP addresses fall in the same range.

   – Import: Used the import option to list IP addresses from a text file. Use this method when you want to discover multiple Active Energy Manager managed resources which have IP addresses which are not in the same range.

   Click **Next**.

7. On the SNMP v1/v2c configuration page, shown in Figure 5-27, if the SynapSense SNMP agent is configured with SNMPv1/v2c, select the appropriate version of SNMP to use to discover the SynapSense SNMP agent. Enter the SNMP protocol community name that is configured in the SynapSense SNMP agent. IBM Systems Director uses this community name to authenticate to SynapSense SNMP agent during the discovery process. Click **Next**.

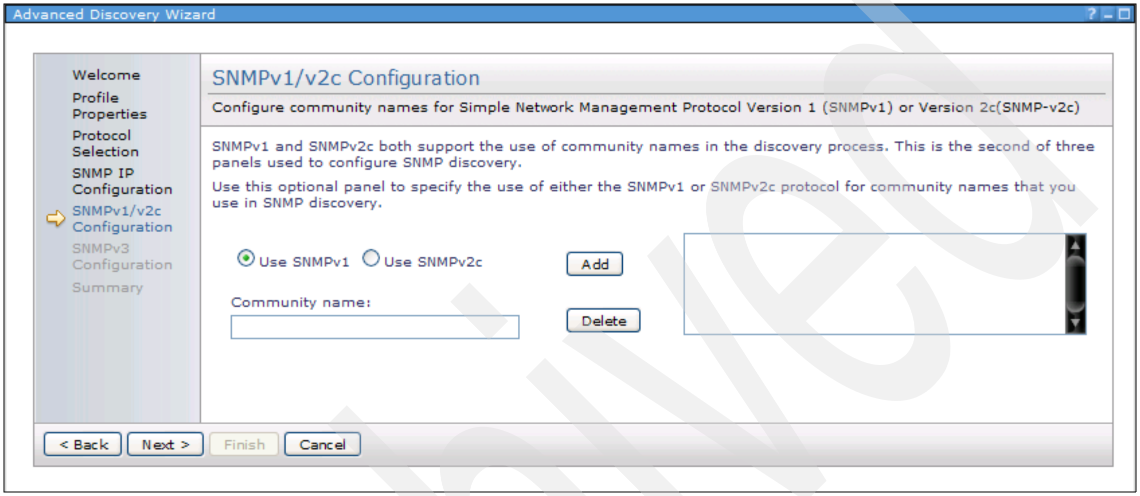


*Figure 5-27 Advanced System Discovery Wizard: SNMP 1/v2c configuration*

8. If the SynapSense SNMP agent is configured with SNMP v3, click **Create** to create the SNMP V3 profile to use to authenticate to the SynapSense SNMP agent during the discovery process. Click **Next**.
9. The Summary page of the Advanced Discovery Wizard shows the summary information of the profile selection. Review the summary, and then click **Finish** to create a new profile.

10. Select the new profile created in the Advanced System Discovery page, as shown in Figure 5-28. Click **Run** to start execution of the selected profile.

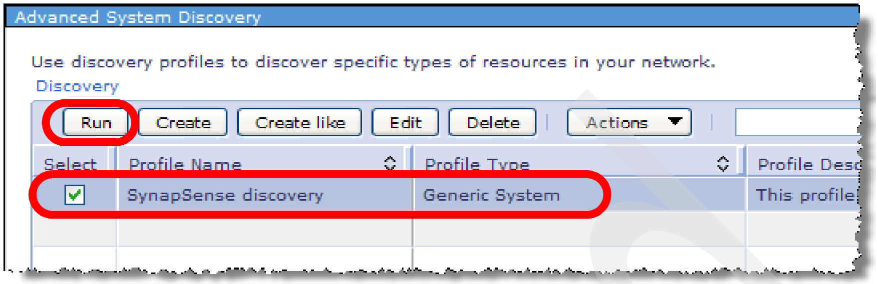


*Figure 5-28 IBM Systems Director Advanced System Discovery page*

11. As shown in Figure 5-29, you can choose to schedule the Advanced System Discovery task as well as set notification options.

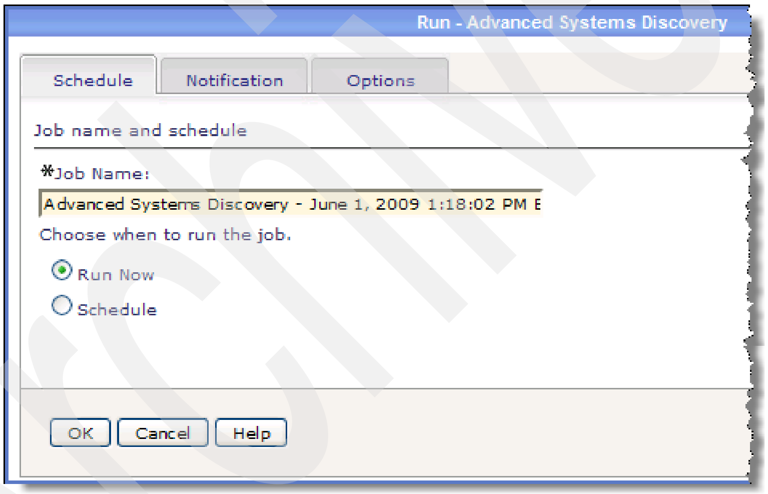


*Figure 5-29 Run - Advanced System Discovery*

12. The discovered SynapSense SNMP agent displays under Generic Systems, as shown in Figure 5-30.

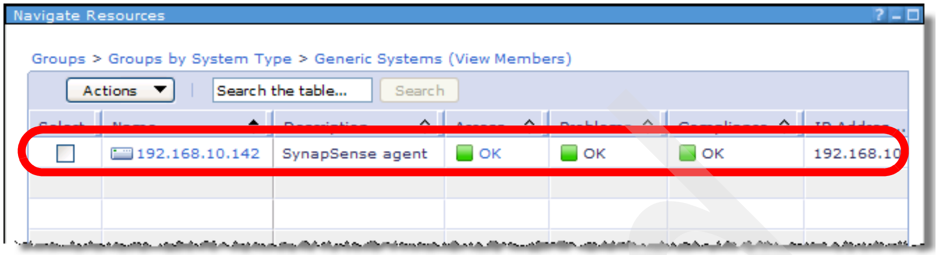


*Figure 5-30 IBM Systems Director: Generic systems group*

After the successful authentication to SynapSense SNMP agent, all the nodes managed by the discovered SynapSense SNMP agent display in the Active Energy Manager resource group. Figure 5-31 shows the topology view of one of the discovered SynapSense sensor nodes.

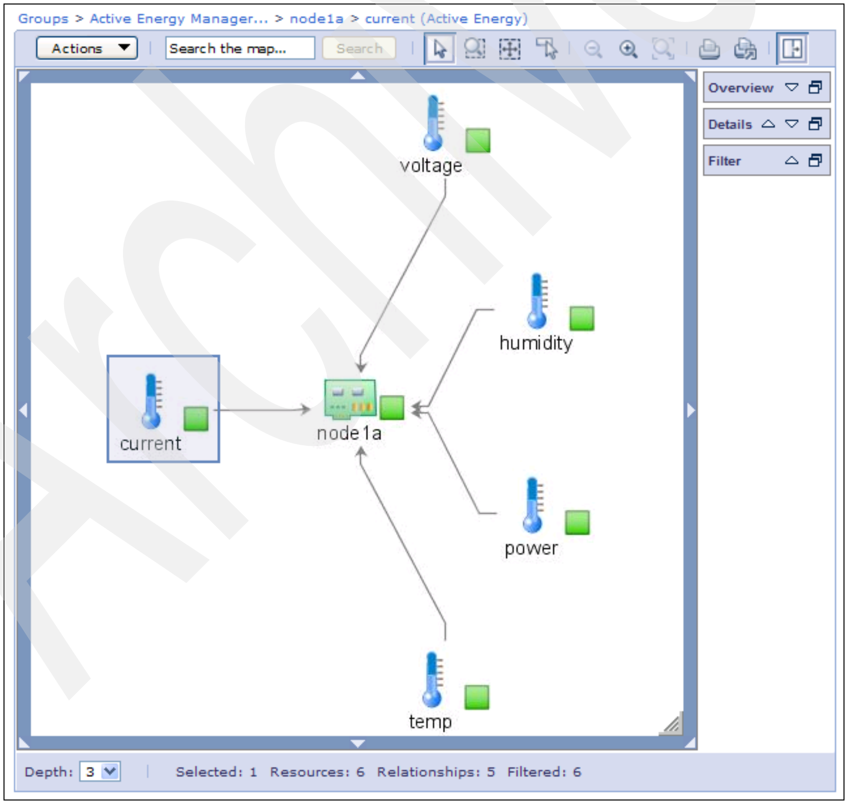


*Figure 5-31 Navigate resources: Topology view of SynapSense node*

### Discovering System z HMC servers using advanced system discovery

Follow these steps to discover the HMC managing a System z server in IBM Systems Director:

1. Click **Advanced system discovery** in the Welcome page of IBM Systems Director as shown in Figure 5-16 on page 157.
2. Click **Create** to create a new profile for discovering System z HMC in IBM Systems Director as shown in Figure 5-17 on page 157.
3. Click **Next** on Advanced Discovery Wizard Welcome page to start the creation of the discovery profile.
4. Provide the name and other details for the profile as shown in Figure 5-32:
   - Profile name: Name of the Advanced System Discovery profile.
   - System Type: Category of the system which will be discovered using this profile. For System z HMC, select system type as **Server**.
   - System subtype: Type of the system that will be discovered using this profile. For System z HMC, select system subtype as **HMC Managing System z**.
   - Profile description: Brief description of the discovery profile being created.

   Click **Next**.

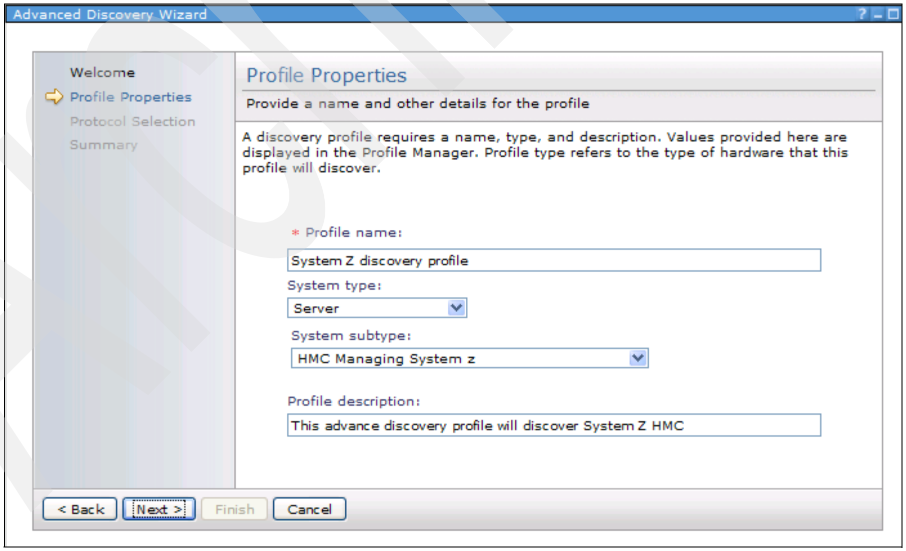


*Figure 5-32 Advanced System Discovery Wizard: Profile properties*

5. Select the protocol that this profile needs to use. For System z HMC, select **SNMP Discovery**, as shown in Figure 5-33. Click **Next**.

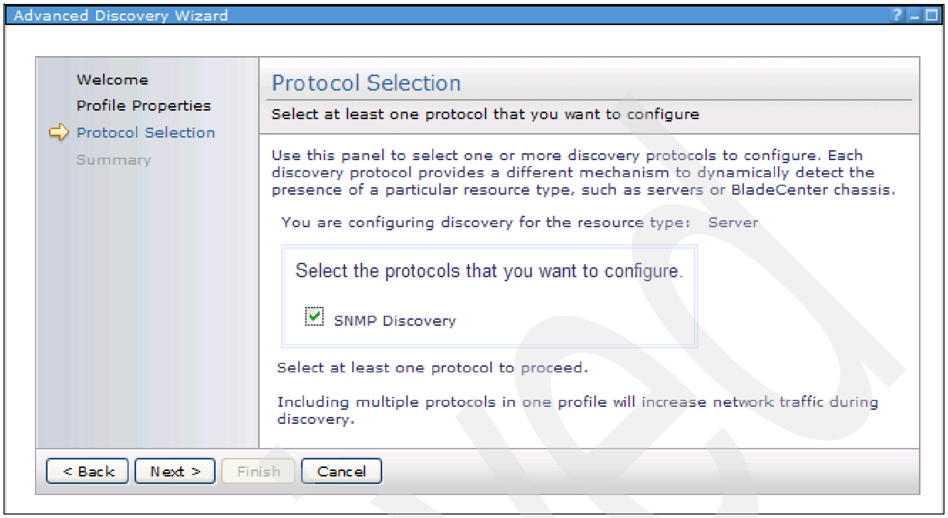


*Figure 5-33 Advanced System Discovery Wizard: Protocol Selection*

6. Input the IP address that will be used for SNMP discovery. You can choose to input IP addresses using one of the following methods:

   – Add a single IP address: Use this method when you want to discover single Active Energy Manager managed resource using this profile

   – Add a range of IP addresses: Use this method when you want to discover multiple Active Energy Manager managed resources whose IP addresses fall in same range.

   – Import: Use the import option to list IP addresses from a text file. Use this method when you want to discover multiple Active Energy Manager managed resources which have IP addresses that are not in the same range.

   Click **Next**.

7. On the SNMP v1/v2c configuration page, shown in Figure 5-34, If the System z HMC is configured with SNMPv1/v2c, select the appropriate version of SNMP to use to discover the System z HMC. Enter the SNMP protocol community name that is configured in the System z HMC. IBM Systems Director uses this community name to authenticate to System z HMC during the discovery process. Click **Next**.

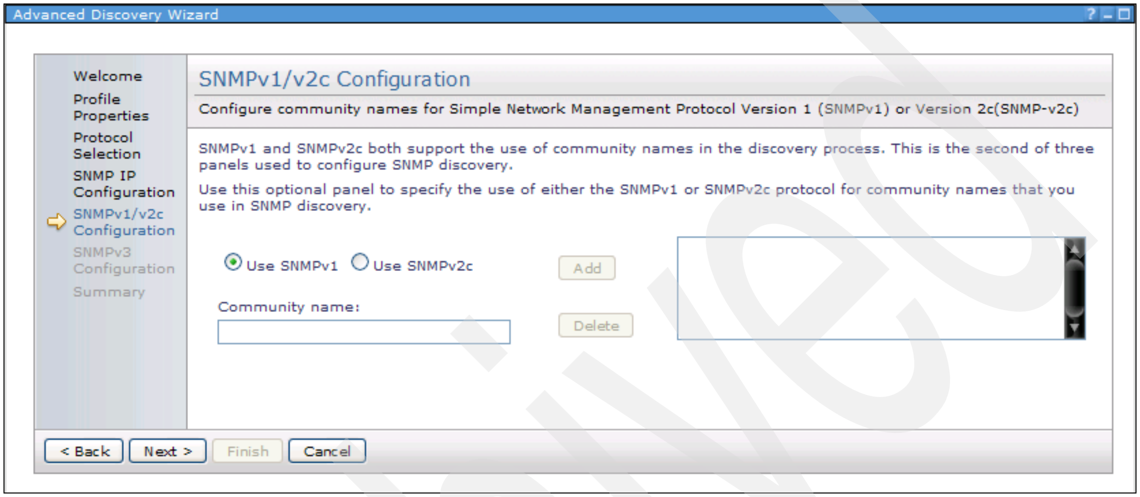


*Figure 5-34 Advanced System Discovery Wizard: SNMP v1/v2c configuration*

8. If the System z HMC is configured with SNMP v3, click **Create** to create SNMP V3 profile that should be used to authenticate to the System z HMC during the discovery process. Click **Next**.
9. The Summary page of the Advanced Discovery Wizard shows the summary information of the profile selection. Review the summary, and click **Finish** to create a new profile.
10. Select the new profile that you just created in the Advanced System Discovery page as shown in Figure 5-35. Click **Run** to execute the profile.

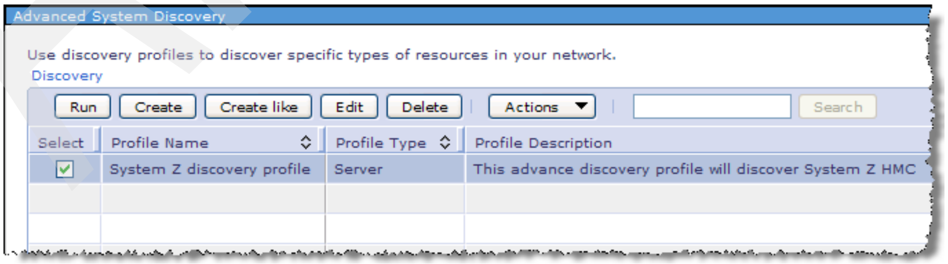


*Figure 5-35 IBM Systems Director Advanced System Discovery page*

11.As shown in Figure 5-36, you can choose to schedule the Advanced System Discovery task as well as set notification options.

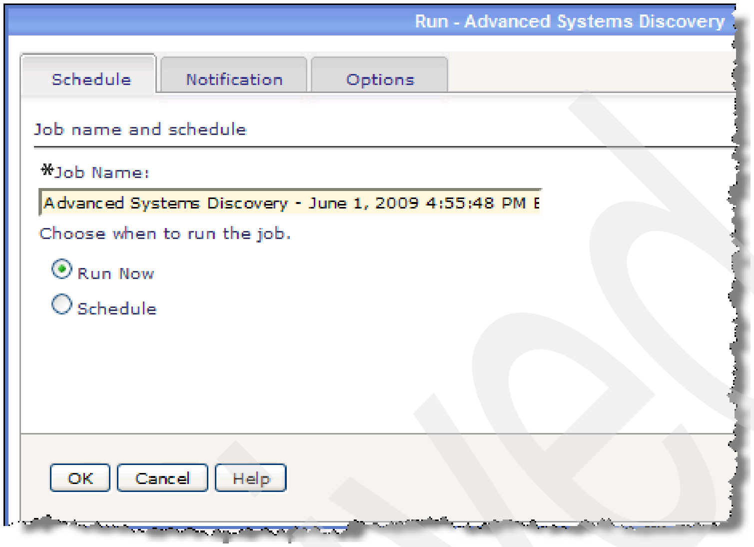


*Figure 5-36 Run Advanced System Discovery*

12.The discovered System z HMC and managed System z servers appears under the “HMC and managed System z servers” group, as shown in Figure 5-37.

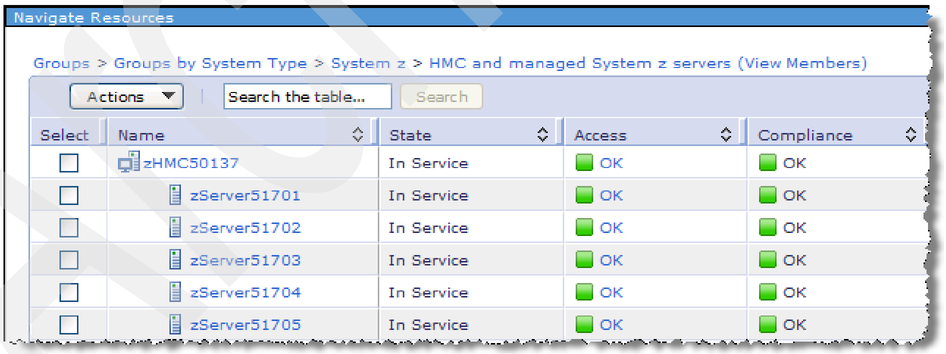


*Figure 5-37 IBM Systems Director: HMC and managed System z servers group*

# 5.3 Authenticating to endpoints

After discovering the Active Energy Manager endpoint, you must authenticate to it before you can gain access to the supported resources that the endpoint controls.

## 5.3.1 Credential requirements

Table 5-2 shows the credentials that you need to authenticate to each Active Energy Manager endpoint. Active Energy Manager endpoints have the capability to monitor and manage the power usage of the power managed objects that they control using a level of functionality that depends on the type of object and the type of endpoint.

*Table 5-2 Active Energy Manager endpoint credential requirements*

| Active Energy Manager endpoints | Credential requirements |
|---|---|
| FSP | User ID: admin[1]<br>Password: |
| HMC | Supervisor-level user ID and password |
| BladeCenter Advanced Management Module | Supervisor-level user ID and password |
| RSA II | Supervisor-level user ID and password |
| BMC | Supervisor-level user ID and password |
| System z HMC | SNMP credentials (Community name or user ID and password) |
| PDU+ | SNMP credentials (Write community name or user ID and password) |
| SynapSense SNMP Agent and sensor nodes | SNMP credentials (Community name or user ID and password) |
| Emerson-Liebert SiteScan instance | SiteScan web service credentials |
| Eaton Power Xpert Reporting Database server | SQL Server authentication credentials |
| APC InfraStruXure Central server | ISXC web service credentials |
| Smart Works gateway | No credentials required. |
| Sensatronics sensor network | Web service credentials |
| 1-wire or iButton sensor networks | Web service credentials |

1. `admin` is the only FSP user ID supported for use with Active Energy Manager.

You can tell whether you have authenticated to a power managed object's Active Energy Manager endpoint by the presence of the lock symbol next to the endpoint's icon in the Managed Objects pane of the IBM Systems Director console. The lock symbol indicates that you have not yet authenticated to the Active Energy Manager endpoint and, therefore, cannot access any power managed objects that it controls. Figure 5-38 shows a managed objects with the lock symbol.

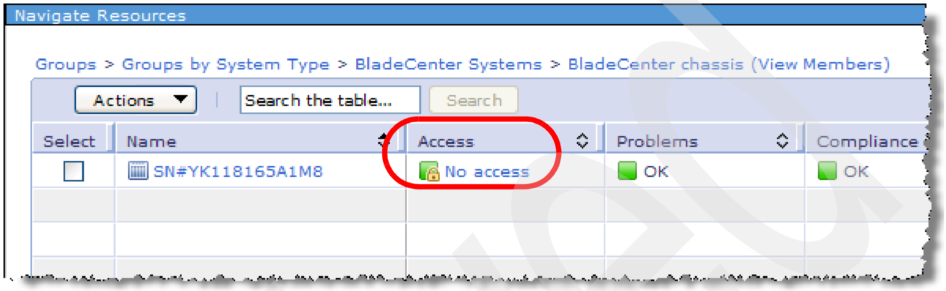


*Figure 5-38 Active Energy Manager endpoints requiring management access*

## 5.3.2 Active Energy Manager endpoints requiring authentication after discovery

The following Active Energy Manager endpoints require you to authenticate *only after discovery* using the system discovery method:

- FSP
- HMC
- RSA II
- BMC
- BladeCenter Chassis
- Emerson-Liebert SiteScan instance
- Eaton Power Xpert Reporting Database server
- APC InfraStruXure Central server
- Sensatronics sensor networks
- 1-wire or iButton sensor networks
- PDU devices
- SynapSense SNMP agent

The following Active Energy Manager endpoints require you to authenticate *only after discovery* using the advanced system discovery method:

- System p FSP managed
- System p HMC managed
- System x RSA II/BMC managed

The following endpoints allow you to provide authentication credentials optionally when attempting to discover them using advanced system discovery:

- BladeCenter Chassis MM/AMM managed
- Emerson-Liebert SiteScan instance
- Eaton Power Xpert Reporting Database server
- APC InfraStruXure Central server
- Sensatronics sensor networks
- 1-wire or iButton sensor networks

For a resource that is controlled by one of these endpoints to appear in the resource view, you must first authenticate to the endpoint through IBM Systems Director Web console as described in 5.3.4, “Authenticating to Active Energy Manager endpoints” on page 174.

> **Note:** PDU devices and SynapSense SNMP agents can be discovered using system discovery only if the PDU devices use SNMP v1/v2c with default community names.

### 5.3.3 Active Energy Manager endpoints requiring authentication during discovery

The following Active Energy Manager endpoints require you to provide authentication credentials or SNMP credentials when attempting to discover them in IBM Systems Director using the advanced discovery method and SNMP v3:

- System z HMC managed servers
- PDU devices
- SynapSense SNMP agents

### 5.3.4 Authenticating to Active Energy Manager endpoints

Before authenticating to Active Energy Manager endpoints, you need to review the following points:

- In the case of authenticating to an HMC:
  - You need a user ID with supervisor level authority.
  - After you have authenticated successfully to the HMC, the resources that the HMC is managing become visible in the Resources group. However, the HMC itself is not shown in the Active Energy Manager Resources page.

– It is the service processor of the resource that the HMC is managing which provides the Active Energy Manager functionality, not the HMC itself.

– When an FSP has been managed by an HMC, you cannot discover the FSP directly as with an Active Energy Manager endpoint. You must connect to and authenticate with the FSPs that manage the HMC first.

**Note:** When logging on to an FSP, `admin` is the only supported user ID that you can use for authentication through IBM Systems Director Web console.

To authenticate to an Active Energy Manager endpoint after discovering, perform the following steps:

1. Locate the discovered objects under their respective group in the Navigate Resources window as shown in Figure 5-39. This example shows a BladeCenter chassis as the discovered object.

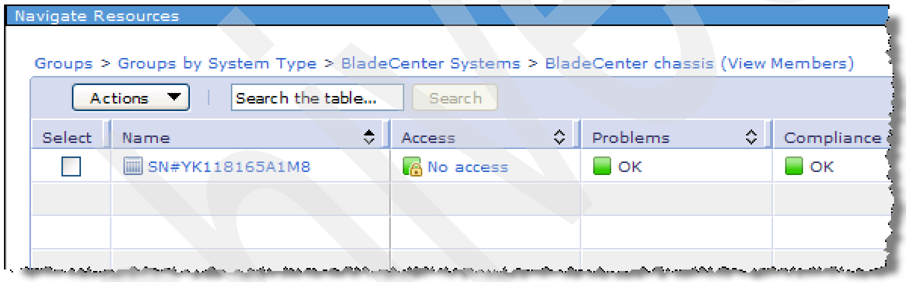


*Figure 5-39 IBM Systems Director Navigate Resource page*

2. Right-click the object in the list and select **Request Access**. Figure 5-40 shows an example of requesting access to the BladeCenter chassis.

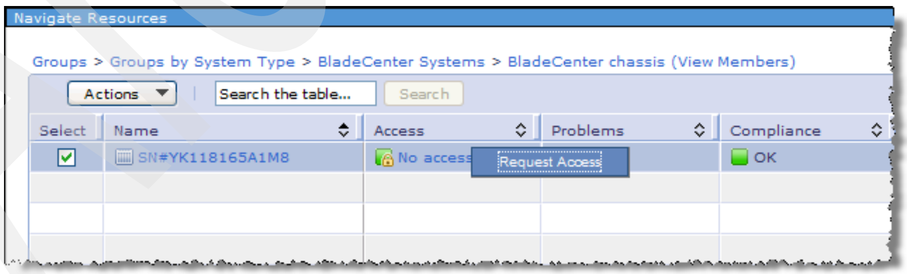


*Figure 5-40 Requesting access to Active Energy Manager endpoint*

3. In the Request Access window (Figure 5-41), enter a valid user ID and password. Refer Table 5-2 on page 172 to determine the credentials that are required to request access to different types of Active Energy Manager endpoints. Enter the valid credentials, and click **Request Access**.

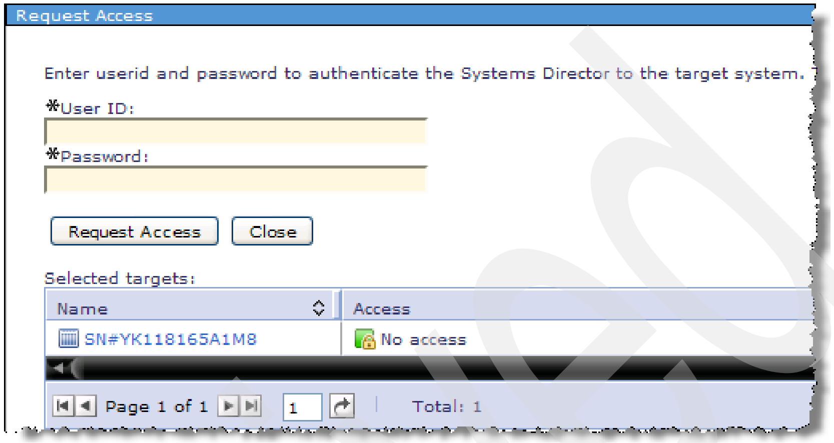


*Figure 5-41 Request access to Active Energy Manager endpoint*

4. The lock icon next to the Active Energy Manager endpoint no longer displays, as shown in Figure 5-42. Click **Close**.

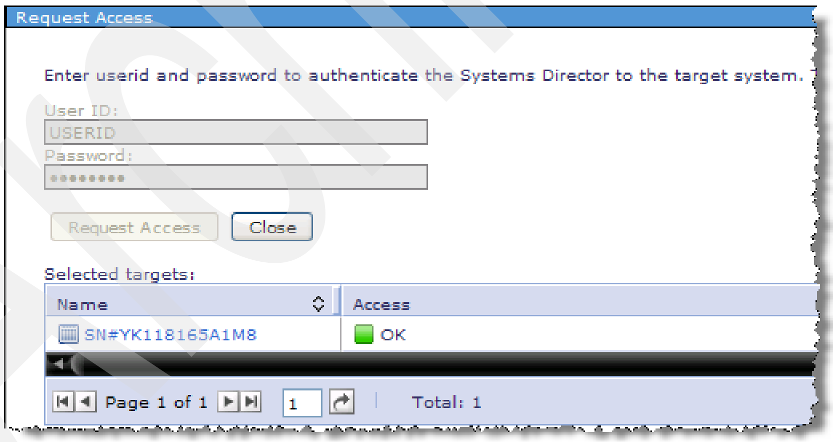


*Figure 5-42 Successful authentication to Active Energy Manager endpoint*

5. After successful authentication to the endpoint, all the Active Energy Manager supported resources that are managed by the endpoint appear in the resources navigation pane, as shown in Figure 5-43.

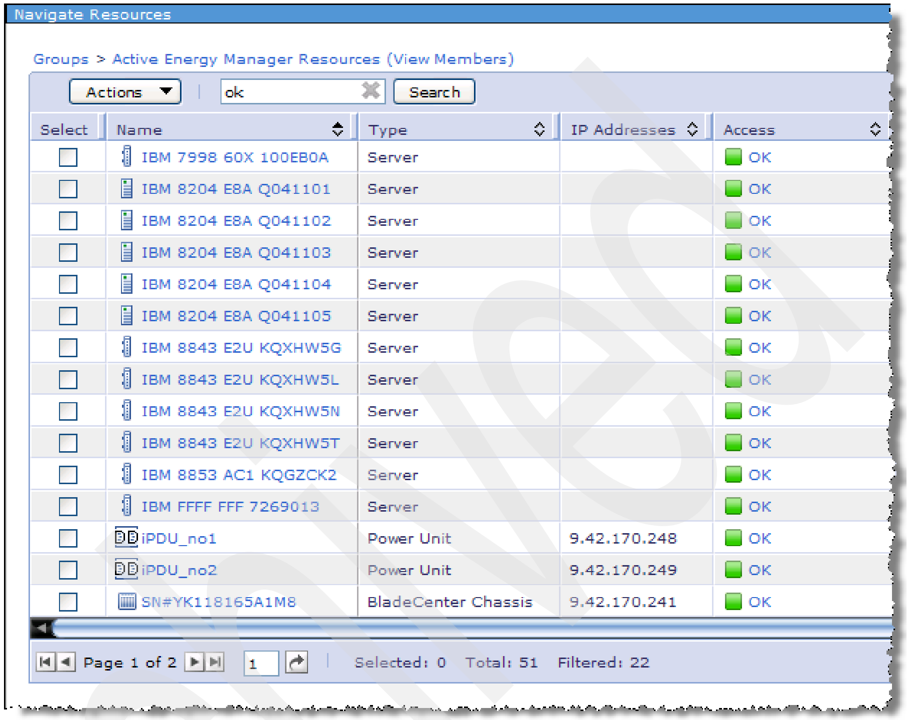


*Figure 5-43 Active Energy Manager resources panel*

## 5.4 Working with Active Energy Manager resources properties

Use the Properties view to display the current active energy information for energy-managed resources in the properties view.

To view active energy properties for an energy-managed resource, follow these steps:

1. In the Director navigation area, click to expand **Energy**, then click **Active Energy Manager**. The Active Energy Manager summary page displays.
2. In the Monitor section, select a resource. Then, click **Actions** → **Properties**.

> **Tip:** You can also access the properties view by right-clicking the resource and selecting **Properties**.

3. Go to the Active Energy tab to view active energy properties for the resource.

You can edit some of the active energy properties. When such properties for a resource are present, an Edit button displays.

To edit active energy properties for an energy-managed resource, follow these steps:

1. Click **Edit** at the bottom of the Active Energy property page. The Edit window opens and displays the properties. The properties vary depending on the resource. Only some of the properties are editable. The editable properties display in fields and lists.
2. Edit one or more of the editable properties.
3. Click **OK** to save the changes.

For descriptions of the active energy properties of all Active Energy Manager supported resources, see the “Working with active energy properties” page in the Information Center:

`http://publib.boulder.ibm.com/infocenter/director/v6r1x/topic/aem_410/frb0_t_viewing_aem_properties.html`

# 6

# Monitoring Active Energy Manager resources

In this chapter, we discuss monitoring Active Energy Manager resources. To provide a foundation for the chapter, we begin with the following basic topics:


- 6.1, "Introducing the Active Energy Manager monitoring tasks" on page 180
- 6.2, "Terminology used in this chapter" on page 183


Then, we cover the following monitoring-related tasks:


- 6.3, "Navigating Active Energy Manager tasks" on page 185
- 6.4, "Creating a group of Active Energy Manager resources" on page 195
- 6.5, "Configuring Active Energy Manager settings" on page 204
- 6.6, "Collecting Active Energy Manager inventory" on page 213


Finally, we describe how to use the following monitoring functions of Active Energy Manager:


- 6.7, "Viewing trend data" on page 221
- 6.8, "Exporting trend data" on page 238
- 6.9, "Using the Energy Cost Calculator" on page 244
- 6.10, "Working with monitors" on page 248
- 6.11, "Working with thresholds" on page 254
- 6.12, "Working with events" on page 257
- 6.13, "Working with automation plans" on page 265


**Note:** All Active Energy Manager monitoring functions are available at no charge.

 

## 6.1 Introducing the Active Energy Manager monitoring tasks

In this section, we provide a high-level introduction to the Active Energy Manager monitoring and monitoring-related tasks.

> **Note:** For a list of power managed systems that support Active Energy Manager monitoring functions, refer to 2.3.2, "Hardware requirements for managed systems" on page 56. For the most up-to-date list, go to:
>
> http://publib.boulder.ibm.com/infocenter/director/v6r1x/topic/aem_410/frb0_r_HW_reqs_managed_systems.html

The tasks that are related to monitoring in Active Energy Manager include the following tasks:

- Navigating Active Energy Manager tasks

  With IBM Systems Director, you can navigate the tasks that you want to perform in several different ways. We describe a single, straightforward way to access the tasks in 6.3, "Navigating Active Energy Manager tasks" on page 185.

  We recommend that you also refer to Chapter 4, "Navigating the IBM Systems Director Web console" on page 109 for a more general discussion of navigating IBM Systems Director.

- Creating a group of Active Energy Manager resources

  You can run many Active Energy Manager functions against a single resource or a group of resources. We describe how to create a group in Active Energy Manager to take advantage of the group functionality in 6.4, "Creating a group of Active Energy Manager resources" on page 195.

- Configuring Active Energy Manager settings

  We explain how to specify Active Energy Manager global default variables in 6.5, "Configuring Active Energy Manager settings" on page 204. These settings include how often the IBM Systems Director server polls the power managed objects for data, how long that data is kept, the default data refresh interval, and information that is related to the calculation of energy costs using the Energy Cost Calculator.

  We also show how to configure Active Energy Manager settings for individual resources by changing the properties of the resource.

> **Note:** In our examples in this chapter, we use watts and degrees Celsius. You can set the power units to BTUs per hour and the temperature to Fahrenheit using the Active Energy Manager settings.

- Collecting Active Energy Manager inventory

  Inventory collection is the process by which IBM Systems Director server establishes connections to network-level resources, such as computers, switches, and printers that are already discovered, and collects information about the hardware and software that is currently installed on those resources. We explain inventory collection in 6.6, “Collecting Active Energy Manager inventory” on page 213.

- Viewing trend data

  You can view the following trend data for a single resource or group over time as either a chart (graph) or a table:

  – Power consumption
  – Power cap
  – Temperature
  – Humidity and dew point
  – CPU speed

  You can set the time interval for viewing trend data from 1 hour up to 1 year, or you can define a custom interval of your choice where start and end times might both be in the past.

  You can also view the current data for a single resource or group of resources. The current data views show you additional detailed information that is not available on the trend data displays.

  We explain how to view trend data in 6.7, “Viewing trend data” on page 221.

- Exporting trend data

  You can export the following types of trend data from the Active Energy Manager database to a file that can be used in an external application:

  – Chart trend data

    Chart trend data can be exported in `.BMP` format (Internet Explorer) or `.PNG` format (Firefox).

  – Table trend data

    Table trend data can be exported in `.CSV` format only.

  We explain how to export trend data in 6.8, “Exporting trend data” on page 238.

- Using the Energy Cost Calculator

  The Energy Cost Calculator can be very useful for estimating the total energy cost for a power managed object. You can specify a cost per kilowatt-hour as well as a factor that takes into consideration the cooling costs that are associated with the energy consumed by the power managed object.

  We describe the functions of the Energy Cost Calculator in 6.9, “Using the Energy Cost Calculator” on page 244.

- Working with monitors

  We describe the monitors that are provided by Active Energy Manager for power managed objects and how to use them in 6.10, “Working with monitors” on page 248. The monitors include:

  – Average input power
  – Average output power
  – Ambient temperature
  – Exhaust temperature
  – Effective CPU speed

- Working with thresholds

  You can set a threshold for a monitor to issue a warning event or a critical event when the monitor has reached a too high level or a too low level. When the threshold is reached, it triggers an event that is logged in the Active Energy Manager event log.

  We describe how to work with thresholds in to 6.11, “Working with thresholds” on page 254.

- Working with events

  Active Energy Manager logs system events and also events that are triggered in response to a threshold being reached. You can use events to generate actions automatically in response to situations that require attention.

  The following types of events are logged in the event log:

  – Critical events
  – Warning events
  – Informational events

  We describe how to use event in 6.12, “Working with events” on page 257.

- Working with automation plans

  As previously mentioned, you can set up a trigger on a monitor to generate an event when a situation arises that requires attention. You can create an automation plan to implement an action in response to a critical or warning event. For example, you could create an automation plan to invoke static power savings on a power managed system in response to a high exhaust temperature reading.

We describe how to use automation plans in to 6.13, “Working with automation plans” on page 265.

## 6.2 Terminology used in this chapter

The following definition can help you to understand the terminology that we use in this chapter and in Chapter 7, “Managing Active Energy Manager resources” on page 283:

- *Device*

  A device is a generic term that we use to represent a computer or other piece of equipment that appears in the Navigate Resources display or the Active Energy Manager (View Members) display on the IBM Systems Director server.

- *IBM Systems Director resource*

  IBM Systems Director resources are devices that can be discovered either directly or indirectly by IBM Systems Director server and that show in the Navigate Resources display of the IBM Systems Director server.

- *Active Energy Manager resource*

  Active Energy Manager resources are devices that can be discovered by IBM Systems Director and appear in the Active Energy Manager (View Members) display on IBM Systems Director server.

- *Active Energy Manager managed endpoint*

  An Active Energy Manager managed endpoint is an IBM Systems Director resource that Active Energy Manager can communicate with to retrieve data from on behalf of devices that the endpoint is connected to. In this book, we shorten the term *Active Energy Manager managed endpoint* to either *Active Energy Manager endpoint* or simply *endpoint*.

  Active Energy Manager endpoints can be thought of as guardians of, or gateways to, systems and metering devices from which Active Energy Manager can retrieve power and thermal data. You must authenticate to the Active Energy Manager endpoint before you can gain access to the power managed objects that the endpoint is connected to.

Active Energy Manager supports the following endpoints:

– System z Hardware Management Console (zHMC)
– Hardware Management Console (HMC)
– BladeCenter management module (BCMM)
– Remote Support Adapter (RSA)
– Baseboard Management Controller (BMC)
– Intelligent Management Module (IMM)
– Flexible Service Processor (FSP)
– IBM Intelligent Power Distribution Unit (PDU+)
– Various non-IBM PDUs, sensors, power units, uninterruptible power supplies, computer room air conditioning units (CRAC), and miscellaneous metering devices as listed at:

  http://publib.boulder.ibm.com/infocenter/director/v6r1x/topic/aem_410/frb0_r_HW_reqs_metering_devices.html

Note that the following endpoints are IBM Systems Director resources but *not* Active Energy Manager resources because they do not show in the Active Energy Manager (View Members) display:

– System z Hardware Management Console (zHMC)
– Hardware Management Console (HMC)

These endpoints act as “passthru” devices that Active Energy Manager uses to retrieve power and thermal data. They are distinguished from other Active Energy Manager endpoints because they do not act as Active Energy Manager “clients.” In other words, they do not contain firmware that can communicate with Active Energy Manager running on the IBM Systems Director server to provide power and thermal monitoring and management functions.

- *Active Energy Manager power managed system*

  Active Energy Manager power managed systems are IBM systems that directly support the power monitoring or management functions of Active Energy Manager. They are connected to an Active Energy Manager endpoint and include the following IBM products:

  – IBM System z10
  – IBM POWER6 processor-based Power 5xx systems
  – IBM POWER6 processor-based JSxx blade servers
  – IBM x86 architecture rack, tower and blade servers (selected models)

  In this book, we shorten the term *Active Energy Manager power managed system* to *IBM power managed system* or simply *power managed system*. Power managed systems are a subset of power managed objects.

- *Active Energy Manager power managed object*

  Active Energy Manager power managed object is a generic term for the IBM systems and IBM and non-IBM metering devices from which Active Energy Manager can retrieve power and thermal data directly. Active Energy Manager power managed objects are a subset of Active Energy Manager resources and comprise devices that are being power monitored or managed by Active Energy Manager endpoints, as well as the following endpoints themselves:

  – BladeCenter management module
  – RSA
  – BMC
  – IMM
  – FSP
  – PDU+

  They do not include devices that have no native metering capabilities or the following Active Energy Manager endpoints:

  – System z Hardware Management Console (zHMC)
  – Hardware Management Console (HMC)

## 6.3 Navigating Active Energy Manager tasks

We cover the following topics in this section:

- Accessing the Active Energy Manager Resources (View Members) display
- Searching the resource list
- Accessing the Groups (View Members) display

> **Note:** In this section, we discuss only the Active Energy Manager Web-based graphical user interface (GUI). If you need to use the command-line interface (CLI), refer to 7.6, “Using the command-line interface” on page 362 for a summary or refer to the information center for more detailed information:
>
> `http://publib.boulder.ibm.com/infocenter/director/v6r1x/topic/aem_410/frb0_r_ref_commands.html`

The IBM Systems Director user interface provides a number of ways to perform a task. As you use Systems Director more and more, you might find one method to perform a given task suits you better than others.

In this chapter, we take a consistent approach by accessing virtually all tasks from the same interface. If you follow this approach, you know that you can always get to the task that you want to complete.

You can access almost all the Active Energy Manager tasks that we cover in this chapter from either of the following displays:

- Active Energy Manager (View Resources), for a single resource
- Groups (View members), for a group

We explain how to navigate to both these displays.

The following Active Energy Manager tasks that we discuss are not available from the single or group resource displays:

- 6.5, “Configuring Active Energy Manager settings” on page 204
- 6.13, “Working with automation plans” on page 265

For a more general discussion of navigating IBM Systems Director and Active Energy Manager, refer to Chapter 4, “Navigating the IBM Systems Director Web console” on page 109.

### 6.3.1 Accessing the Active Energy Manager Resources (View Members) display

The Active Energy Manager Resources (View Members) display lists all of the resources known to Active Energy Manager. It is helpful to become familiar with this display, because from this display you can run almost all of the tasks (for a single resource) that we cover in this chapter and in Chapter 7, “Managing Active Energy Manager resources” on page 283. We use this display extensively.

To access the Active Energy Manager Resources (View Members) display to run a task on a *single* Active Energy Manager resource, follow these steps:

1. From the IBM Systems Director Welcome page, click **Active Energy Manager** (Figure 6-1 on page 187), or expand **Energy** in the left-hand navigation pane and click **Active Energy Manager** (Figure 6-2 on page 187).


*Figure 6-1 Accessing Active Energy Manager tasks - 1a*

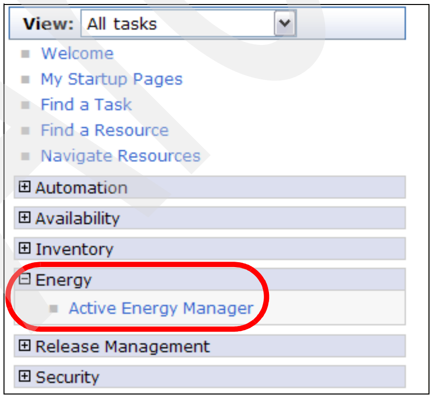


*Figure 6-2 Accessing Active Energy Manager tasks - 1b*

2. In the Active Energy Manager home page (Figure 6-3), click **Active Energy Manager Resources (View Members)**.

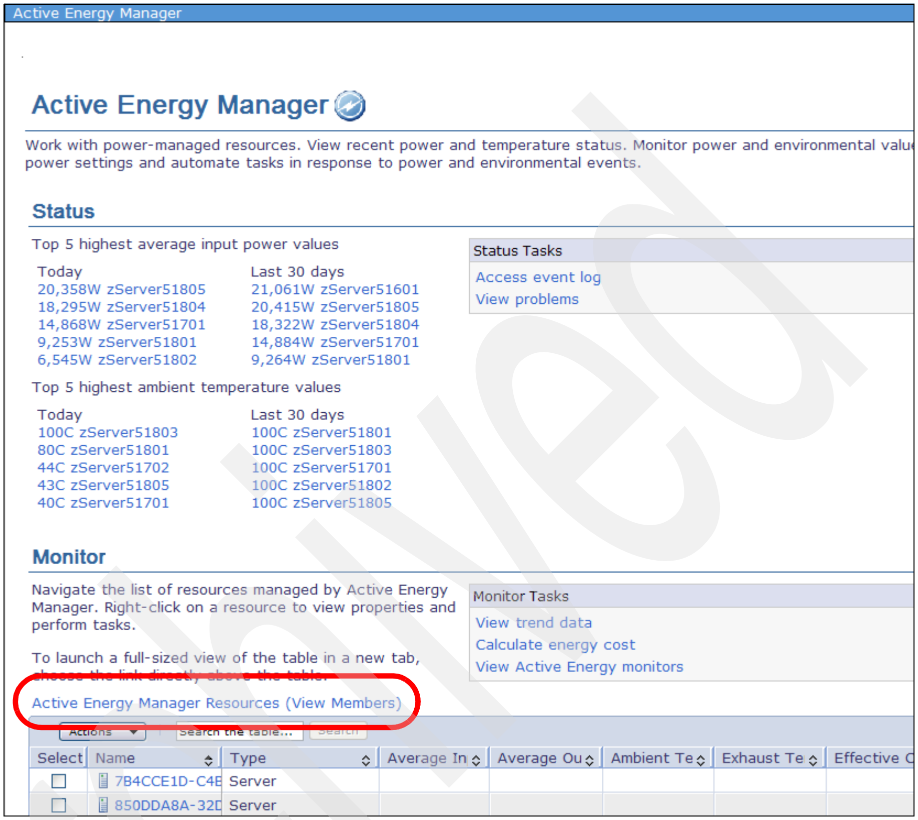

*Figure 6-3 Accessing Active Energy Manager tasks - 2*

3. You see a list of Active Energy Manager resources such as the one shown in Figure 6-4.

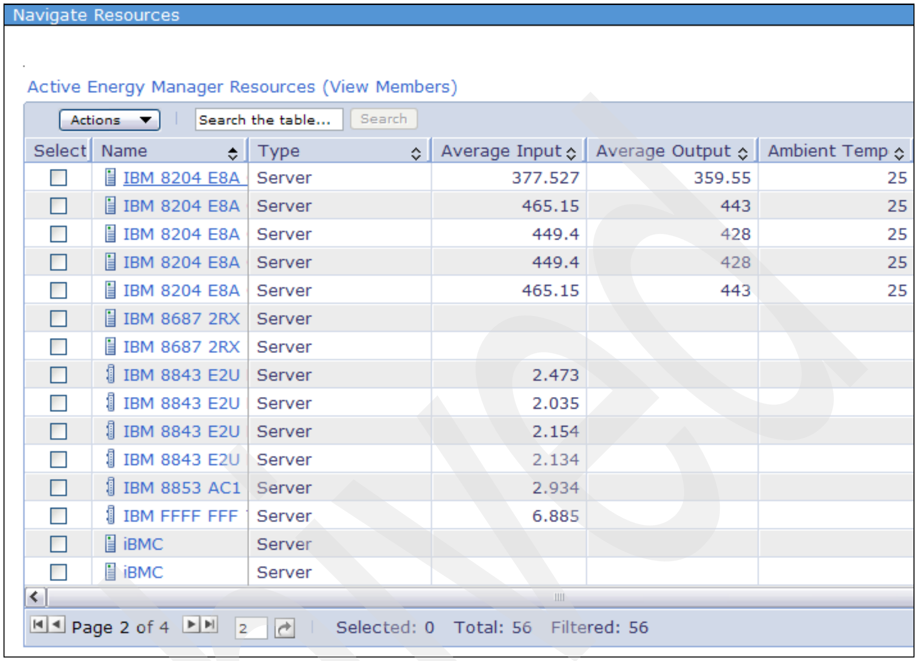


*Figure 6-4 Accessing Active Energy Manager tasks - 3*

4. To change the displayed columns, or to add new ones, click **Actions** and select **Columns** (Figure 6-5 on page 190). Then, click **OK**.

   You can adjust the width of the new column using the Width tab in the dialog box. You can also adjust the width of any column by dragging the column bar in the header of the table as with other applications.

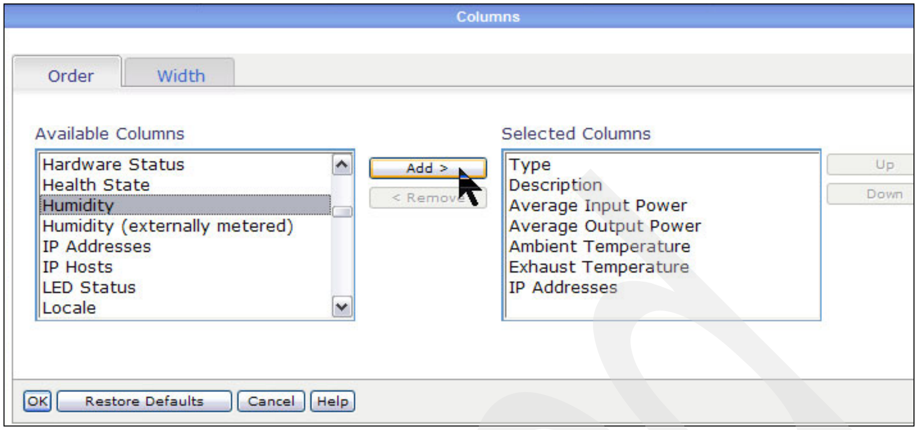


*Figure 6-5 Adding and removing columns*

5. To run most of the Active Energy Manager functions that we discuss in this chapter, simply right-click a resource and select the option that you want from the context menu. Figure 6-6, Figure 6-7 on page 191, and Figure 6-8 on page 191 show some examples of accessing Active Energy Manager functions.

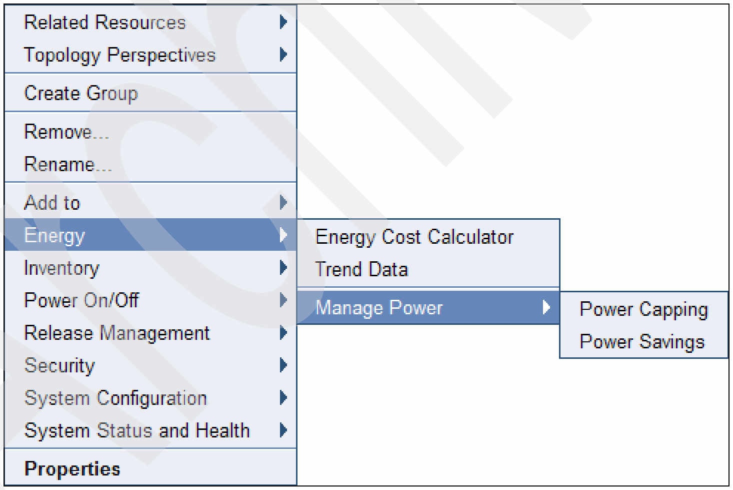


*Figure 6-6 Active Energy Manager functions from the resource context menu - 1*

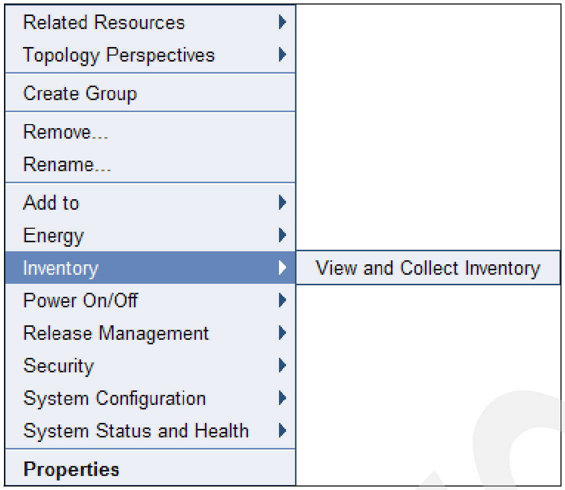


*Figure 6-7 Active Energy Manager functions from the resource context menu - 2*

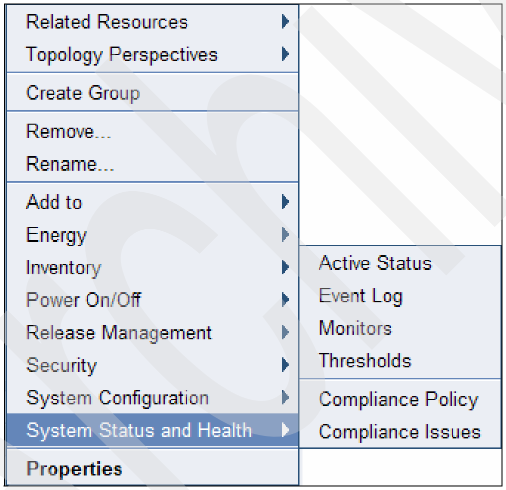


*Figure 6-8 Active Energy Manager functions from the resource context menu - 3*

## 6.3.2 Searching the resource list

If you know the name of the resource that you want, simply type a string into the search field of the Active Energy Manager (View Members) display, and click **Search** (Figure 6-9).

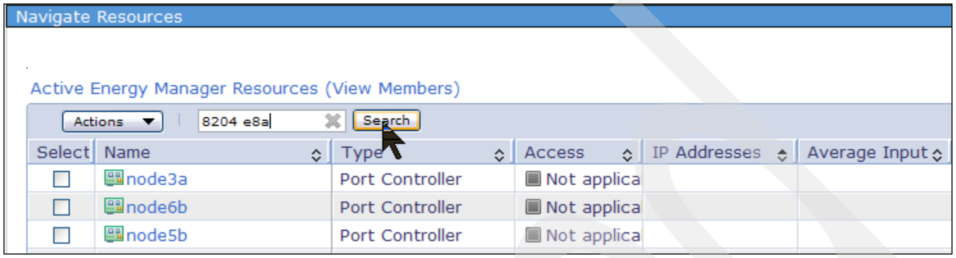


*Figure 6-9 Searching the resource list - 1*

You see a list of the resources that contain the search string (Figure 6-10).

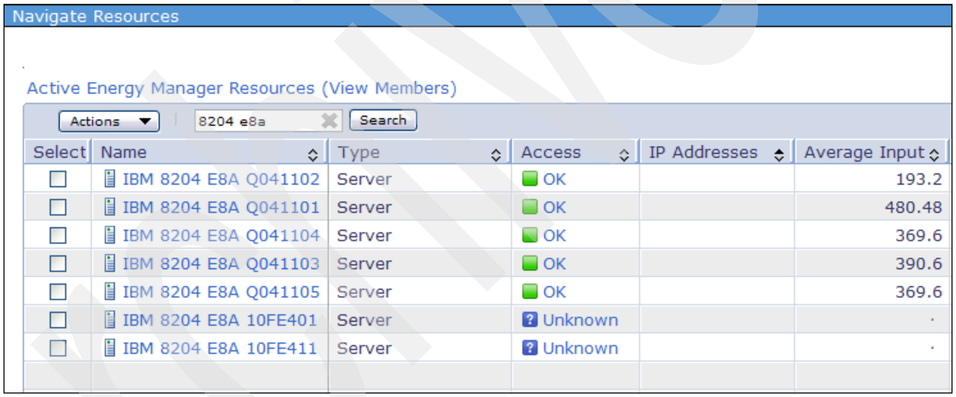


*Figure 6-10 Searching the resource list - 2*

## 6.3.3 Accessing the Groups (View Members) display

To access the Groups (View Members) display to run a task on a group of Active Energy Manager resources, follow these steps:

1. First close any open Navigate Resources tabs. Then, click **Navigate Resources** in the left-hand navigation pane (Figure 6-11).


*Figure 6-11 Accessing the Group (View Members) display - 1*

2. The Groups (View Members) display opens, as shown in Figure 6-12. Note that groups can be nested, so you might need to drill down by clicking one of the groups shown in the display.

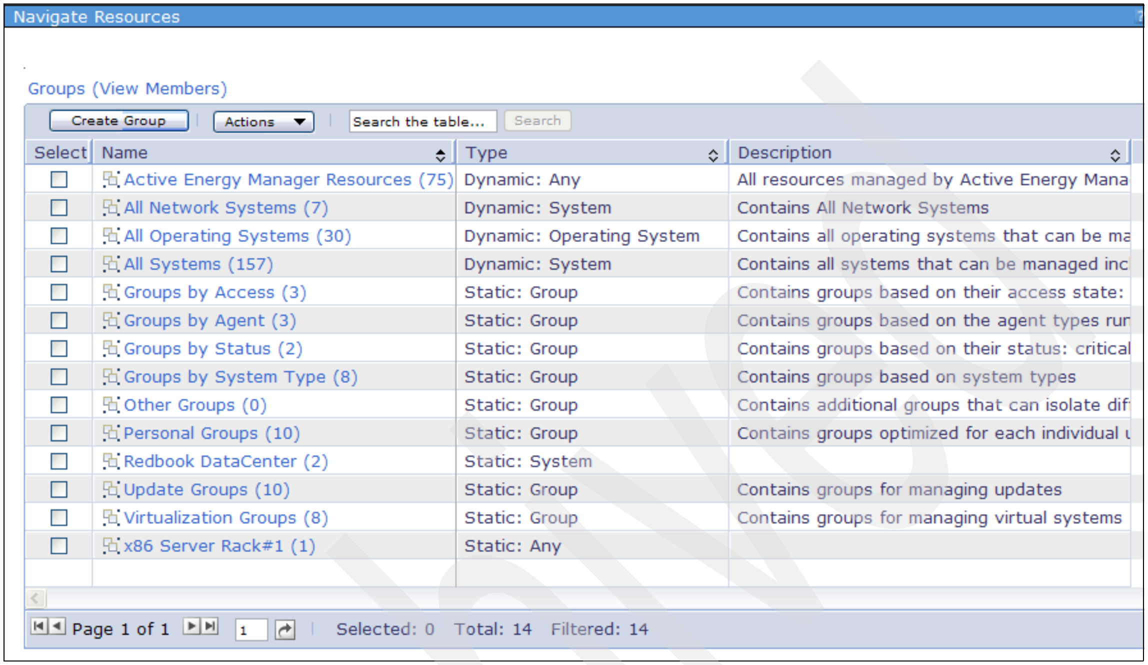


*Figure 6-12 Accessing the Group (View Members) display - 2*

3. To run Active Energy Manager tasks on a group, simply right-click the group to view the context menu. Displays open that are similar to the displays for an individual resource in Figure 6-6 on page 190, Figure 6-7 on page 191, and Figure 6-8 on page 191. However, when you try to run a task (for example power capping) on a group of Active Energy Manager resources, you might see a warning that the function is not supported on all members of the group (as shown in Figure 6-13). All members of a group need to support the function that you are trying to run.

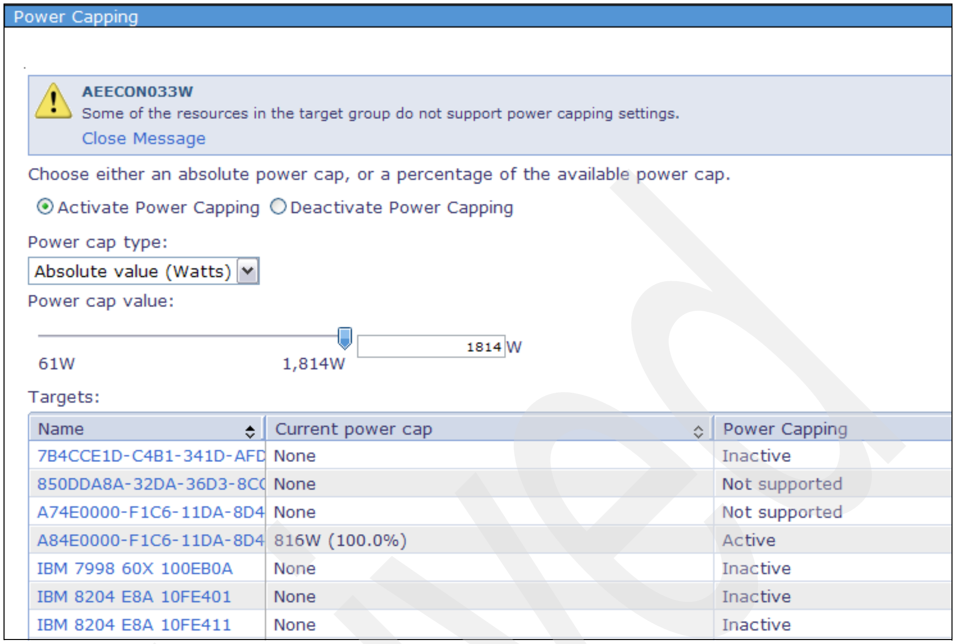


*Figure 6-13 Running an Active Energy Manager task on a group*

## 6.4 Creating a group of Active Energy Manager resources

You normally run an Active Energy Manager task on a single resource, but running a task on a group can be very useful. You can create a group of Active Energy Manager resources and then perform operations on the group versus individual resources. For example, viewing power data for a group shows the aggregated power consumption for all resources in the group.

Most Active Energy Manager functions support groups; however, depending on the makeup of the group, the data might not be meaningful. Also, you need to make sure that all members of the group that you create support the functions that you want to run on that group. Note that a single resource can be a member of multiple groups. Therefore, you can create additional groups (with overlap between the groups) to allow for any differences in Active Energy Manager functionality.

Note that there is a special type of policy called *group power capping* that allows you to set a power cap for a whole group but have Active Energy Manager calculate an individual cap for each member of the group. The individual cap for each managed system is based on the hardware configuration of the particular system, rather than simply averaging the cap over all members of the group. As the configuration of a group member changes, the power cap is recalculated to match the new configuration. Note, however, that the aggregate power cap of all the group members at no time exceeds the specified group power cap.

For a more detailed explanation of group power capping, refer to “Creating a group power capping policy” on page 336.

> **Note:** In this book we discuss how to create only static groups. If you want to create dynamic groups, refer to *Implementing IBM Systems Director 6.1*, SG24-7694.

To create a group of resources, follow these steps:

1. Navigate to the Active Energy Manager Resources (View Members) display. Select the resources for your group (two in this example). Then, right-click one of the resources, and select **Create Group**, as shown in Figure 6-14.

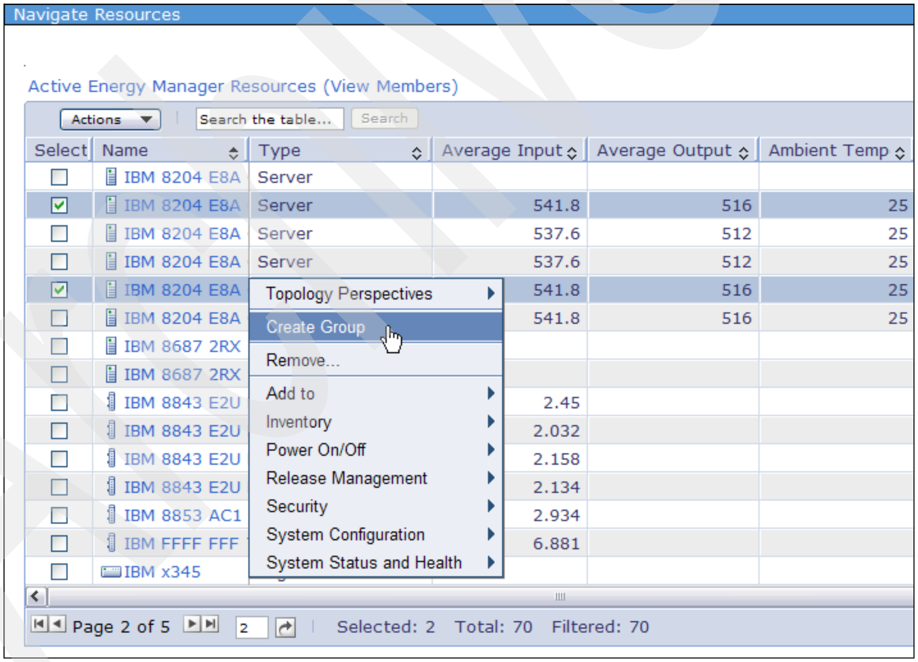


*Figure 6-14 Creating a resource group - 1*

2. The group editor wizard starts. Click **Next** on the Welcome display, and then give your group a name and description as required (Figure 6-15). Click **Next**.

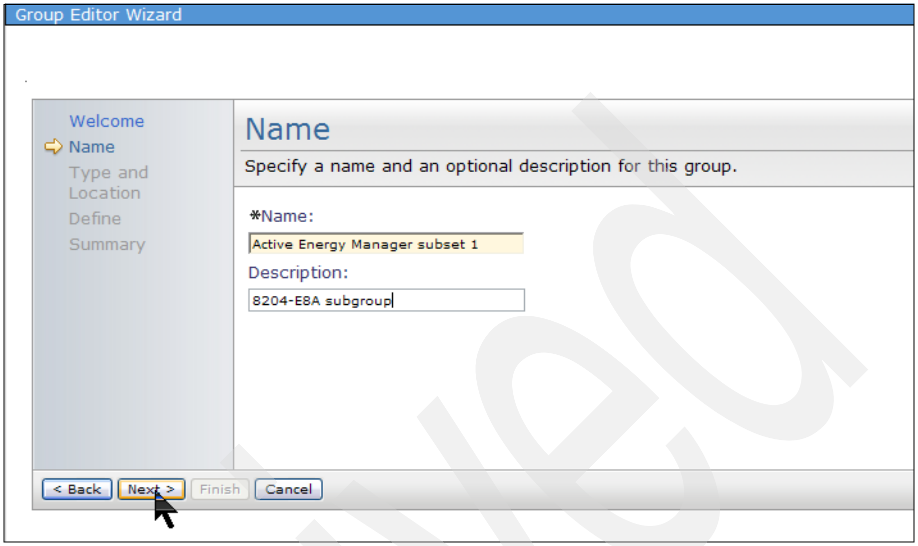


*Figure 6-15 Creating a resource group - 2*

3. In the Type and Location panel, shown in Figure 6-16, select the **Member Type** and a **Group Location**. Locations can be nested within other locations. Think of a group location as a Windows Explorer folder.

   You can click **Browse** to select a group location that is not listed. Notice that a default location, based on your user ID and the name of the IBM Systems Director Server, was created when the user ID was created. Click **Next**.

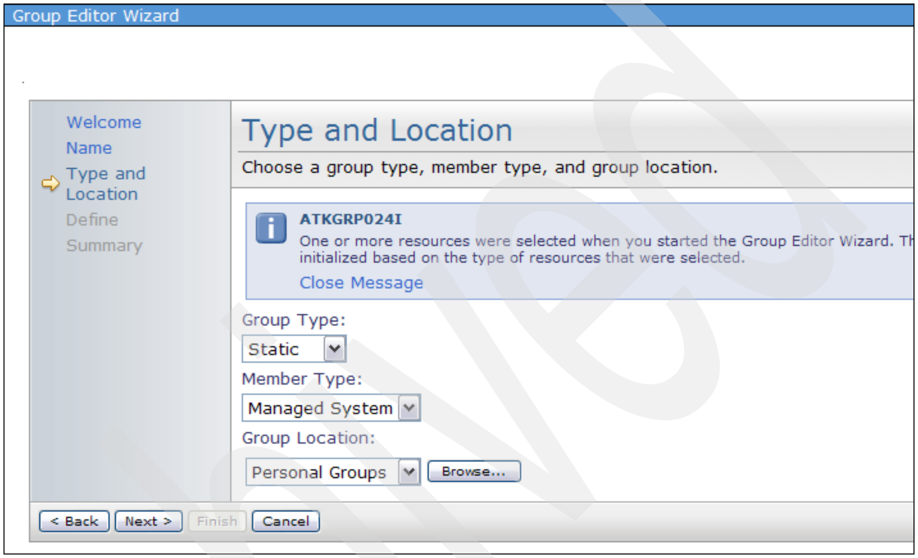


*Figure 6-16 Creating a resource group - 3*

4. The wizard shows the resources that you preselected when you began, as highlighted in Figure 6-17. Click **Next**.

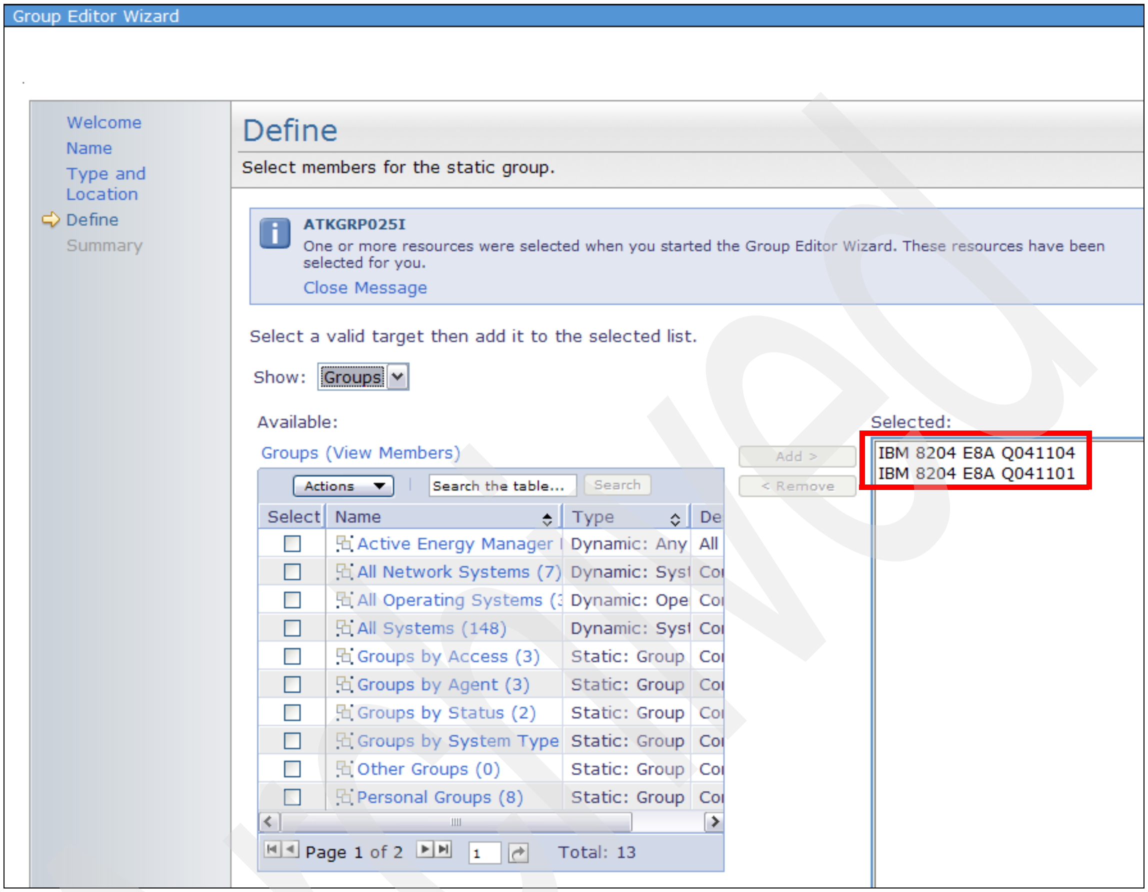


*Figure 6-17 Creating a resource group - 4*

5. The wizard displays the summary window. Click **Finish**.

6. You are returned to the Navigate Resources display (Figure 6-18). Click **View Group** to view summary status information for the members of your group.

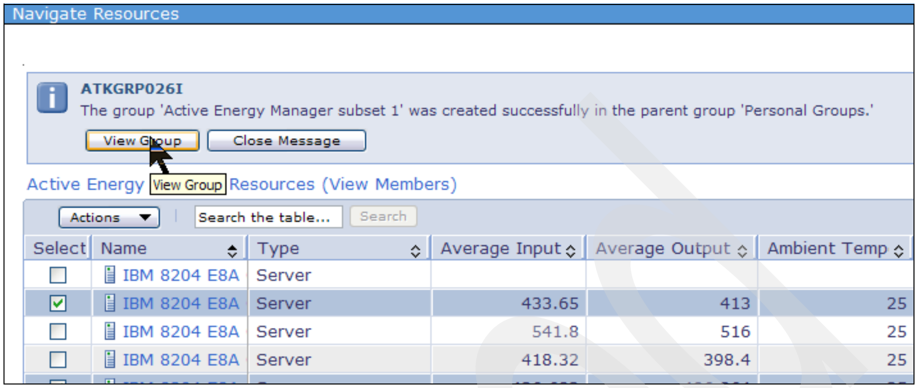


*Figure 6-18 Creating a resource group - 5*

7. A display of the selected resources opens, similar to that shown in Figure 6-19.

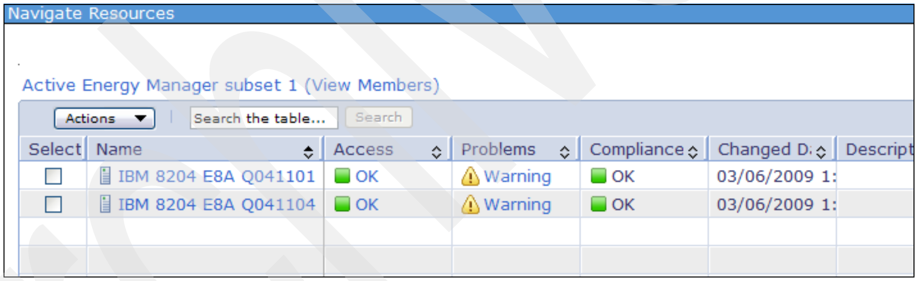


*Figure 6-19 Creating a resource group - 6*

8. To view your new group (at a group level), first close any open Navigate Resources tabs, and then click **Navigate Resources** in the left-hand navigation pane, as shown in Figure 6-20.


*Figure 6-20 Creating a resource group - 7*

9. The display shown in Figure 6-21 opens. Click the location where you created your new group (**Personal Groups** in our example).

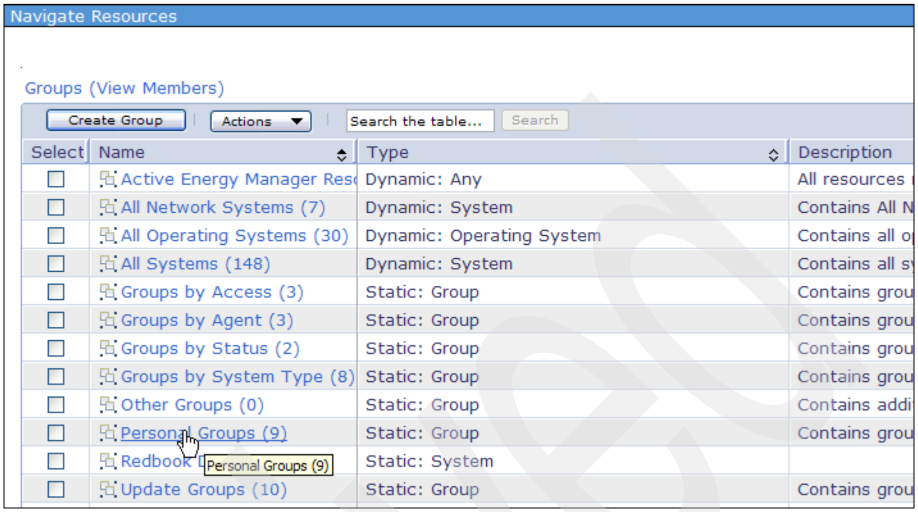


*Figure 6-21 Creating a resource group - 8*

10. The display shown in Figure 6-22 opens. Click the new group that you created to see the group members (**Active Energy Manager subset 1** in our example).

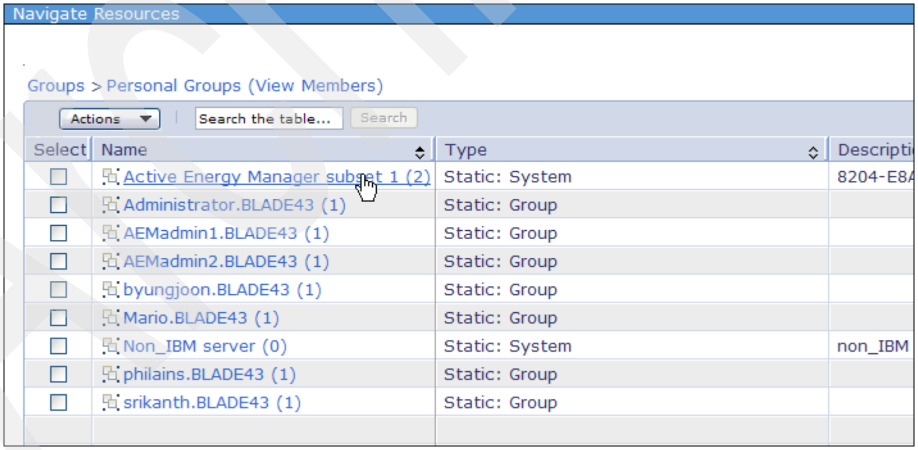


*Figure 6-22 Creating a resource group - 9*

11. A display of the group members opens, similar to that shown in Figure 6-23. (This display is the same display that you saw in Figure 6-19 on page 200.)

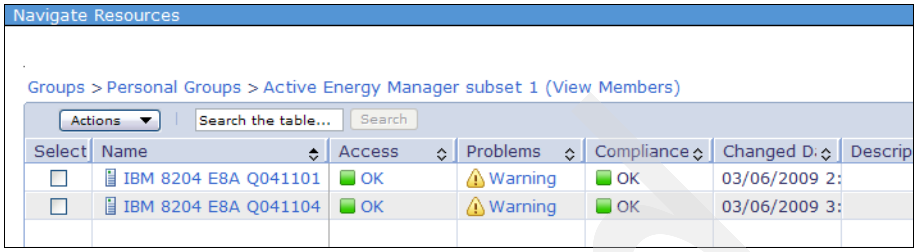


*Figure 6-23 Creating a resource group - 10*

12. You can now use Active Energy Manager groups with tasks such as viewing trend data, power capping, and power savings in a similar way to individual resources. Using the context menu shown in Figure 6-24, you can also manage your groups (editing, exporting, removing, and renaming).

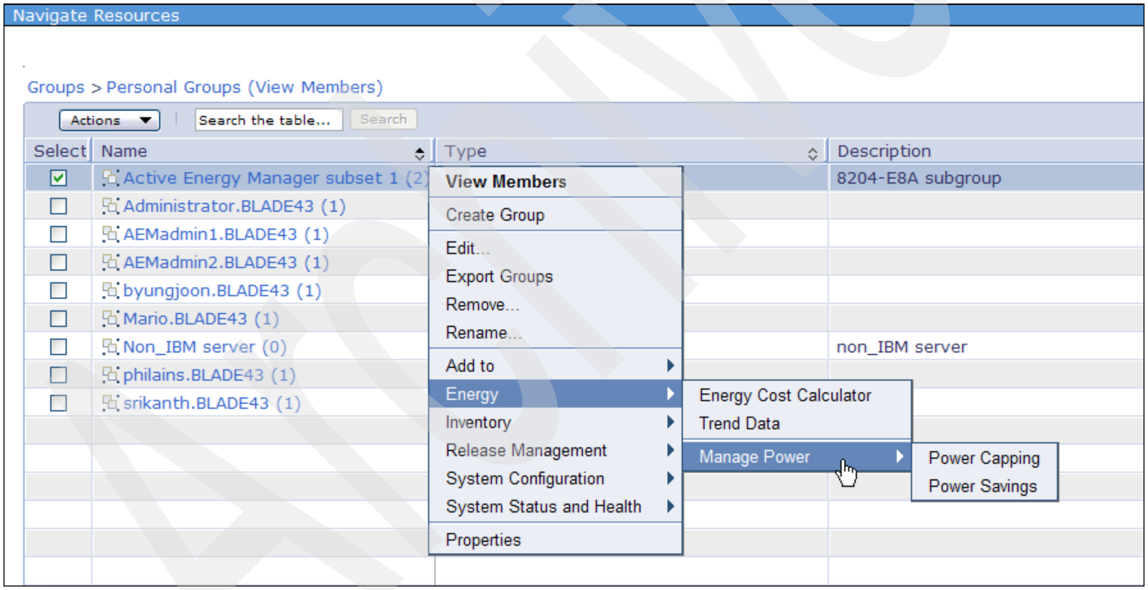


*Figure 6-24 Working with Active Energy Manager groups*

## 6.5 Configuring Active Energy Manager settings

We cover the following topics in this section:

- Configuring settings globally
- Configuring settings for an individual resource

You can configure Active Energy Manager settings at two levels:

- *Global* settings apply to all Active Energy Manager resources and tasks.
- *Resource properties* apply only to the selected Active Energy Manager resource.

Notice that settings that are set in the properties of a resource override global settings.

**Note:** You need SMManager authority or higher to change Active Energy Manager settings.

### 6.5.1 Configuring settings globally

You can change the global default settings by first clicking **Settings** on the Active Energy Manager home page, as shown in Figure 6-25.

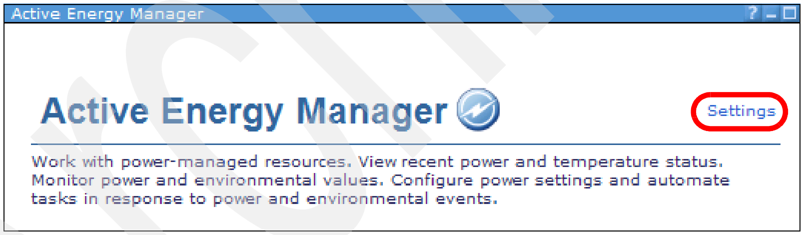


*Figure 6-25 Active Energy Manager home page*

The Active Energy Manager Setting display opens, as shown in Figure 6-26.

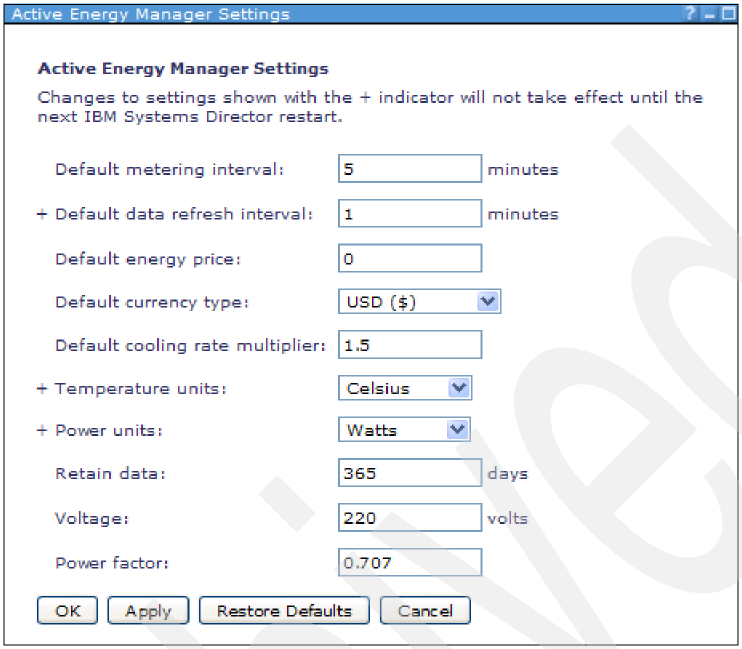


*Figure 6-26 Active Energy Manager settings*

> **Tip:** The metrics preceded by a plus sign (+) require that you restart IBM Systems Director before any changes become active.

You can change the following settings:

- Default metering interval
  - Valid values: 1-60 minutes
  - Default: 5 minutes

  Use this setting to specify how often you want to collect trend data for a resource. The metering interval is specified in minutes and is how often the resource is polled for trend data.

  > **Note:** For a BladeCenter management module, Active Energy Manager uses a minimum polling interval of 10 minutes, regardless of the value specified.

Take the following factors into consideration when setting the polling interval:

– How often does data need to be gathered to meet your requirements? Are there power managed objects for which you need to closely monitor, or others that you can monitor less often if you know that the data changes little over time? If only basic power consumption data is required, polling every 5 or 10 minutes is probably sufficient.

– Depending on the environment in which Active Energy Manager is running (server CPU speed, disk speed, network bandwidth, number of systems being polled, and polling interval), Active Energy Manager might not be able to maintain the polling rate at the set value. In this case, Active Energy Manager polls the devices as fast as it is able. Increasing the polling interval can decrease the CPU, disk, and network load on the IBM Systems Director server.

– The more frequently systems are polled, the more data that is collected and the more disk space that is needed to hold this data.

– If the metering interval for a specific resource is configured, it, rather than the default metering interval, is used for that resource.

- Default data refresh interval

  – Valid values: 1-60 minutes
  – Default: 1 minute

  Use this setting to specify the interval, in minutes, that information is refreshed on trending charts and tables. Changing the refresh interval can be useful when you are viewing large amounts of data.

  On BladeCenter systems, trending information is refreshed at a default rate of once every 10 minutes. On SynapSense sensor networks, trending information is refreshed at a default rate of once every 5 minutes. You can specify shorter metering intervals, but Active Energy Manager meters these resources at these minimum default intervals, regardless.

  **Note:** Changes to this setting do not take effect until you restart the IBM Systems Director server.

- Default energy price

  Use this setting to specify the price of energy per unit. If the value of the Power units setting in Figure 6-26 on page 205 is watts, the energy unit is a kilowatt-hour. If the value of the Power units setting is BTUs/hour, the energy unit is a British Thermal Unit (BTU).

- Default currency type

  Use this setting to specify the currency type for documentation purposes. No currency conversion is done. The default is based on the IBM Systems Director environment locale settings.

- Default cooling rate multiplier

  – Default: 1.5

  Use this setting to specify the multiplier for calculating the cost of energy to cool Active Energy Manager resources in relation to the amount of heat that they produce. For example, if you use the default cooling rate factor of 1.5, for every unit of energy consumed, you allow an additional 1.5 units of energy to be used in removing the heat that is produced. To estimate the cost of the energy consumed by an Active Energy Manager resource alone, set this factor to 0.

- Temperature units

  – Valid values: Celsius or Fahrenheit
  – Default: Celsius

  Use this setting to specify the unit that is used to display temperature values in Active Energy Manager.

  > **Note:** Changes to this setting do not take effect until you restart the IBM Systems Director server.

- Power units

  – Valid values: watts or BTUs/hour
  – Default: watts

  Use this setting to specify the unit that is used to display power values in Active Energy Manager.

  > **Note:** Changes to this setting do not take effect until you restart the IBM Systems Director server.

- Retain data

  – Valid values: 7 - 3,650 days
  – Default: 365 days

  Use this setting to specify the number of days to keep historical data. The default setting in Active Energy Manager is to keep trend data for 365 days and then delete it. Active Energy Manager can keep the data for a maximum of 3650 days.

Note that after 7 days Active Energy Manager compresses automatically trend data for each metric into a single value that is averaged over 1 hour.

When managing a large number of systems, a large amount of data is gathered and saved, often in the range of gigabytes. Decreasing the number of days that data is saved, therefore, decreases the amount of disk space that is needed to store that data.

If the Active Energy Manager database cannot expand due to disk storage constraints, the oldest data is deleted automatically to make room for new data.

- Voltage
  - Default: 220 Volts

  Use this setting to specify the voltage value to use when converting current values to power values. Current values are retrieved from current sensors and are converted to power values in order to view them as trend data.

  It is important that you set this value correctly, because incorrect settings cause large errors in the reported power consumption values.

  Note that you can set the voltage for a specific resource, which overrides the default voltage setting for that resource.

- Power factor
  - Valid values: 0 - 1
  - Default: 0.707

  Use this setting to specify the power factor that is related to the AC current that is consumed by power using devices, or in this case Active Energy Manager resources. Use this setting when converting current values (amps) to power values. Current values (amps) can be retrieved from current sensors and converted to power values in order to view them as trend data.

## 6.5.2 Configuring settings for an individual resource

You can change the settings for individual Active Energy Manager resources. Notice that settings for an individual resource take precedence over global settings. For a complete list of properties that are available for each type of resource, refer to:

http://publib.boulder.ibm.com/infocenter/director/v6r1x/topic/aem_410/frb0_t_viewing_aem_properties.html

To change settings for an individual resource, follow these steps:

1. Click **Active Energy Manager Resources (View Members)** on the Active Energy Manager home page. Then, right-click a resource, and select **Properties** as shown in Figure 6-27.

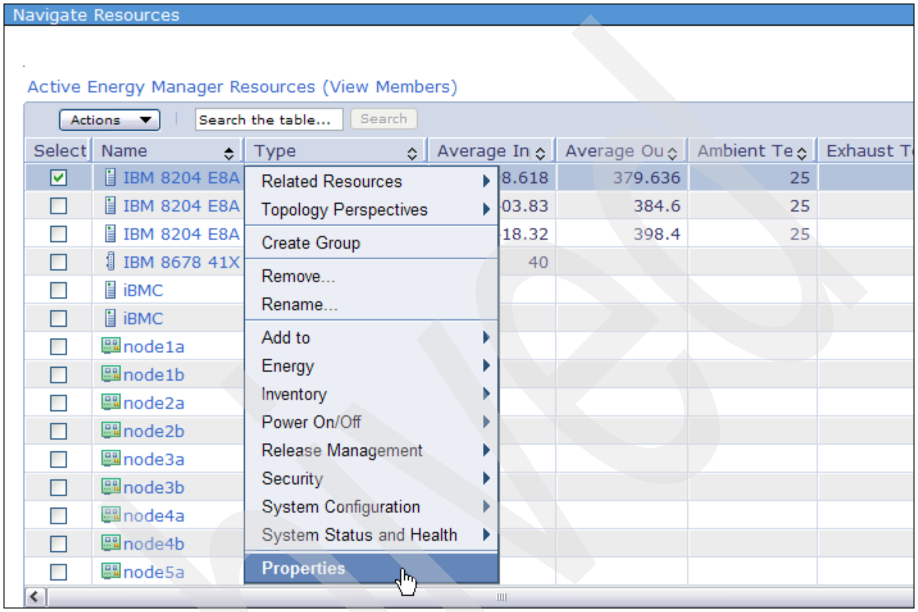


*Figure 6-27 Active Energy Manager resources display*

2. Click the Active Energy tab, and then click **Edit**. The display shown in Figure 6-28 opens.

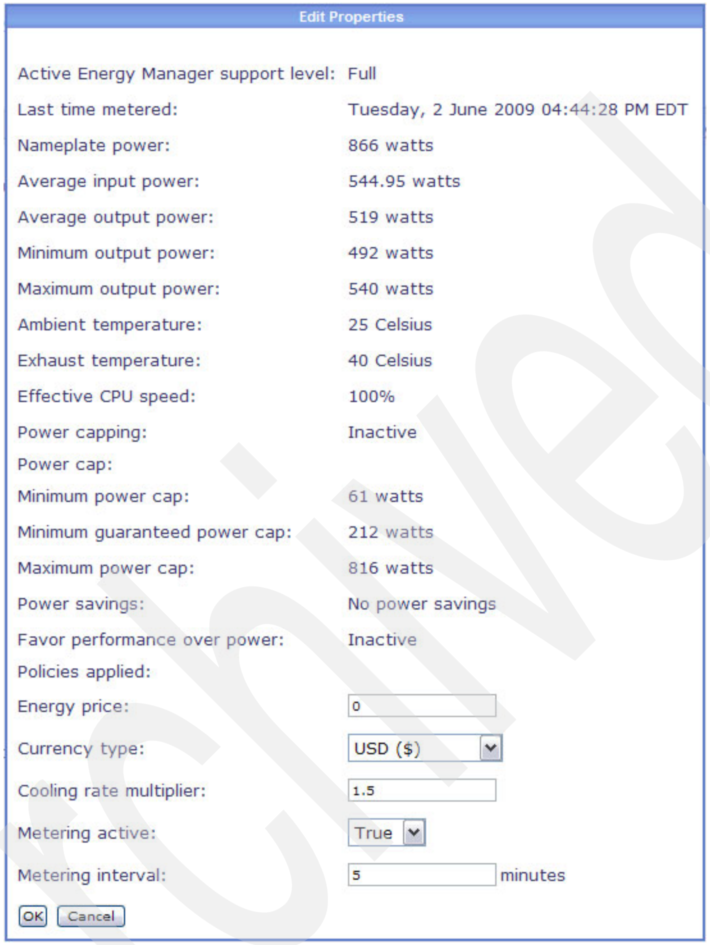


*Figure 6-28 Edit the properties of a resource*

3. Make any changes to the editable fields, and then click **OK**.

   Notice that you are limited in the number of properties that you can change and that the properties you can change differ from resource to resource.

## Disabling and enabling metering for a resource

You can disable metering (data collection) for an Active Energy Manager resource by editing the properties for the resource as shown in Figure 6-28 on page 210 and setting **Metering Active** to **False**. Click **OK**.

Active Energy Manager resources are enabled for metering by default.

> **Note:** Data that is collected during metering includes power and environmental data as well as state data (for example whether power management functions such as power capping are currently enabled). Disabling metering prevents Active Energy Manager from detecting changes in these states.
>
> If you disable metering, do not rely on Active Energy Manager to display such state information accurately. If you want this information to remain current for a specific resource, do not disable its metering.

## Setting the altitude for a resource

Setting the system altitude might improve energy-optimized fan control. POWER6 systems are designed to operate in worst-case conditions, that is at high altitude.

By default, fan control firmware in each system assumes this high altitude until static power savings or dynamic power savings are enabled. In this case, the system firmware attempts to reduce power by monitoring power consumption, ambient temperature, thermal signature, and altitude to reduce fan speeds to save energy.

If no altitude is set in Active Energy Manager, the system firmware assumes an elevation of 350 meters above sea level. In this case, the fans operate at a higher speed than necessary, because the air is thinner at the higher the altitude. Therefore, by setting the altitude to a more realistic value (0 to 100 meters in most cases), you can save some energy.

You can set the altitude only at the resource level. There is no global option.

To set the altitude for a resource, follow these steps:

1. Right-click the resource in the Active Energy Manager Resources (View Members) display as shown in Figure 6-27 on page 209. Select **Properties**.
2. On the properties display for the resource, click **Location** (Figure 6-29).

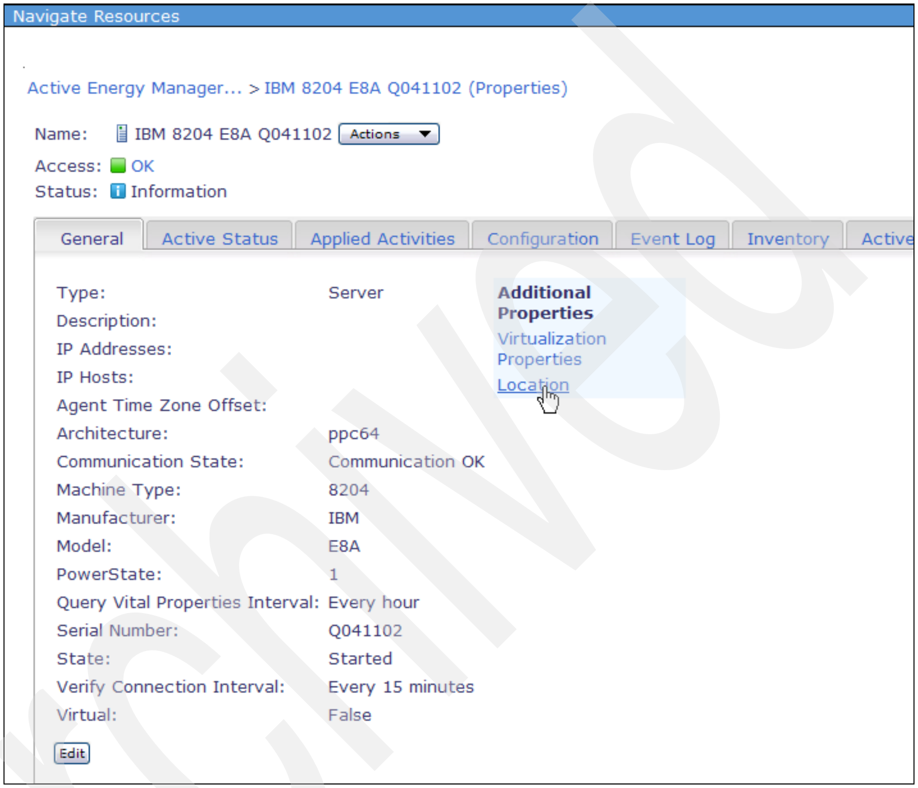


*Figure 6-29 Setting the altitude for a resource - 1*

3. When the location properties display, click **Edit** to edit the location properties, and change the altitude, as shown in Figure 6-30. Click **OK**.

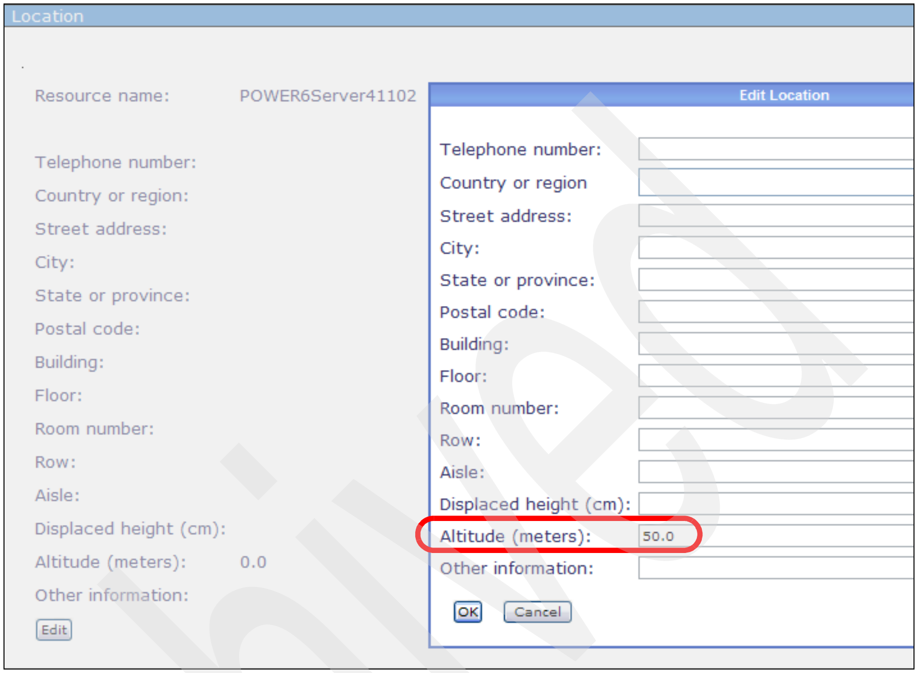


*Figure 6-30 Setting the altitude for a resource - 2*

## 6.6 Collecting Active Energy Manager inventory

We cover the following topics in this section:

- Collecting inventory at the individual resource level
- Collecting inventory at the group level
- Monitoring energy usage at the component level

After a resource is discovered and access is granted to it, you can collect inventory data relating to the resource. You can collect inventory for Active Energy Manager resources at the group level or by selecting one or more individual resources. We use the View and Collect Inventory task to view and manage an extended set of resources and relationships for network-level resources that are already discovered.

You can use inventory collection to monitor energy usage at a more granular level than is available through the Active Energy Manager Resources (View Members) display.

## 6.6.1 Collecting inventory at the individual resource level

To collect inventory at an individual resource level, follow these steps:

1. Click **Active Energy Manager Resources (View Members)** on the Active Energy Manager home page.
2. Right-click a resource and select **Inventory** → **View and Collect Inventory** as shown in Figure 6-31. In this example, we selected a BladeCenter chassis.

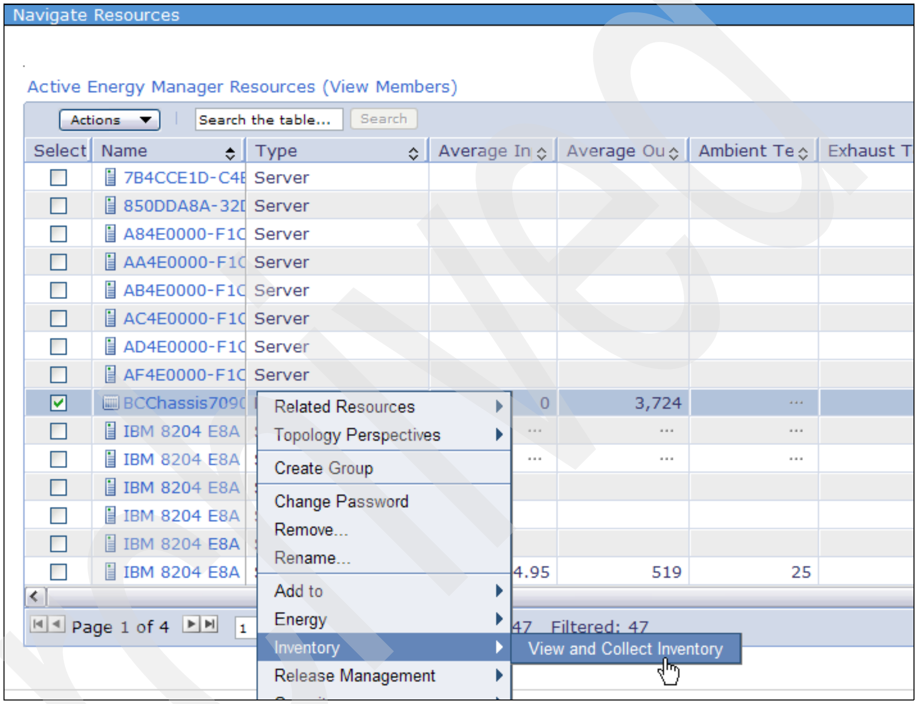


*Figure 6-31 Viewing collected inventory for an Active Energy Manager resource - 1*

3. The View and Collect Inventory view opens similar to that shown in Figure 6-32.

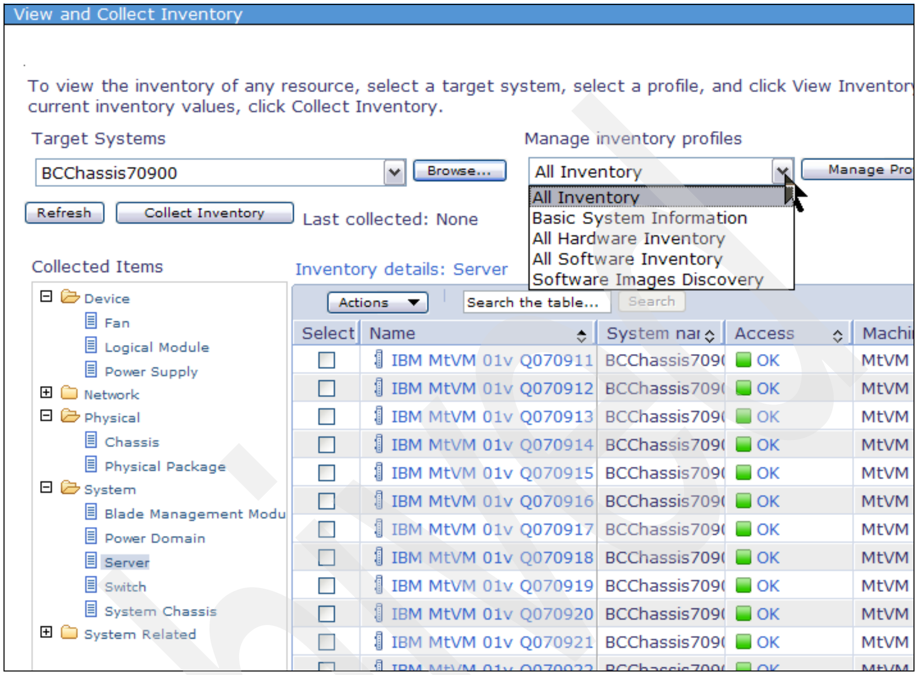


*Figure 6-32 Viewing collected inventory for an Active Energy Manager resource - 2*

In Figure 6-32 notice the following features:

– In this example, we clicked **Server** in the navigation pane to display the blade servers in the BladeCenter chassis that we selected in the previous step.

– The default for inventory collection is **All Inventory**, which shows all possible inventory classes for the target system in the Collected Items navigation pane.

– You can subset the inventory data that displays by selecting one of the profiles in the Manage inventory profiles drop-down menu.

– You can define your own inventory profile by clicking **Manage Profiles**, and then clicking **Create** on the Manage Inventory Profiles display to start the Create Inventory Profile wizard.

– You can use the Refresh button to update the view.

4. You can drill down further by clicking one of the selectable resources shown in Figure 6-32 on page 215. In this case, you see the properties for one of the blade servers as shown in Figure 6-33.

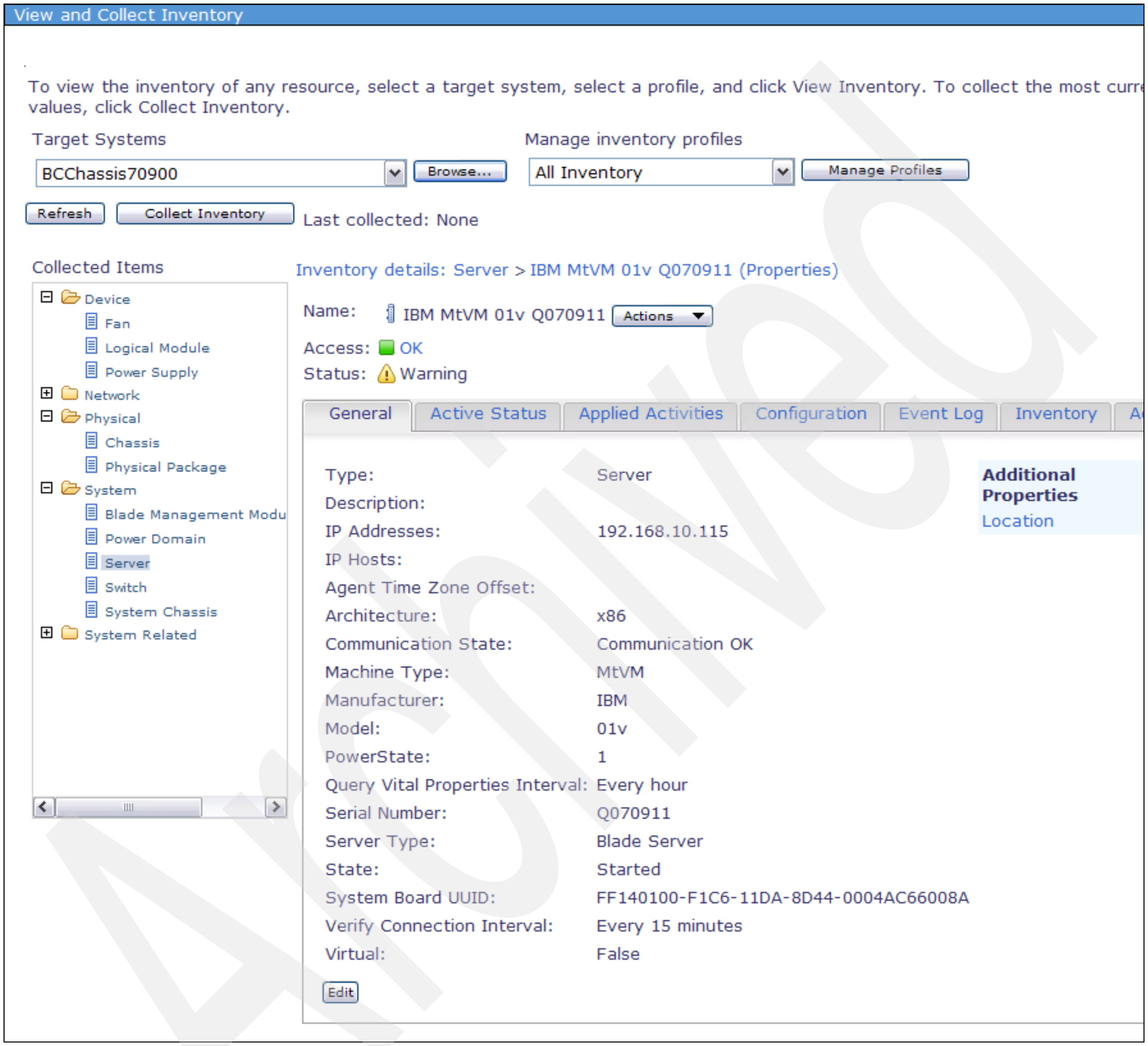


*Figure 6-33 Viewing collected inventory for an Active Energy Manager resource - 3*

5. Instead of drilling down, you can also right-click the name of an inventory item, and select an option from the context menu, as shown in Figure 6-34. Notice the options to power on or power off the blade server.

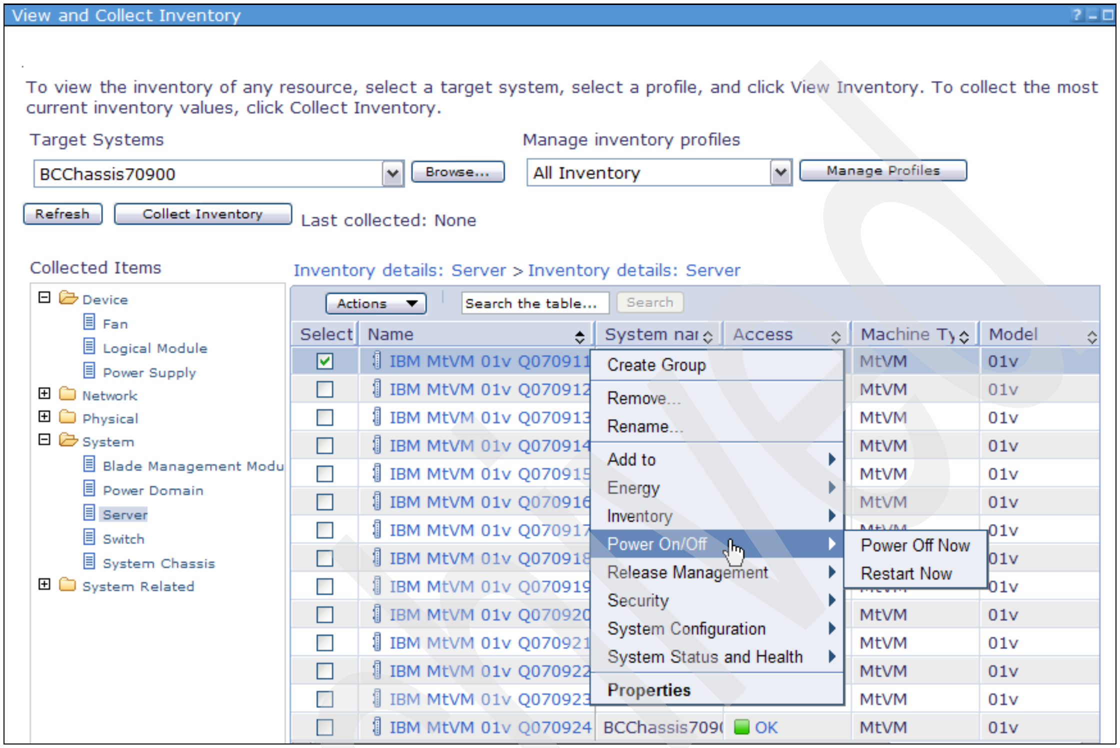


*Figure 6-34 Using the context menu for an inventory item*

6. If you click **Collect Inventory** in Figure 6-34, you can select to schedule the collection of inventory as shown in Figure 6-35.

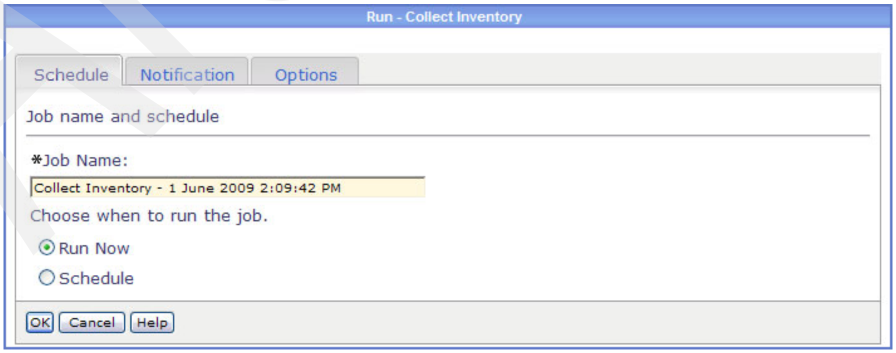


*Figure 6-35 Scheduling inventory collection*

## 6.6.2 Collecting inventory at the group level

To collect inventory at a group level, follow these steps:

1. Close any open Navigate Resources tabs, and then click **Navigate Resources** in the left-hand navigation pane of the IBM Director Server display. On the Navigate Resources display, right-click an Active Energy Manager group, and select **Inventory** → **View and Collect Inventory** (Figure 6-36).

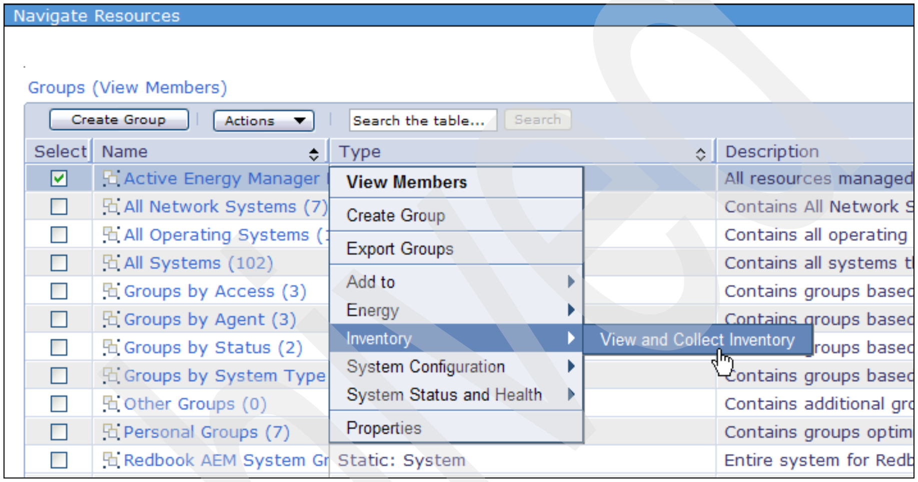


*Figure 6-36 Collecting inventory for the Active Energy Manager group - 1*

2. A display opens that is similar to that shown in Figure 6-37 on page 219.

   Notice that from this display that you have access to all Active Energy Manager resources that IBM Systems Director has discovered. However, due to the number of resources displayed it might be a better option to display the inventory for one or two resources only, as described in 6.6.1, “Collecting inventory at the individual resource level” on page 214.

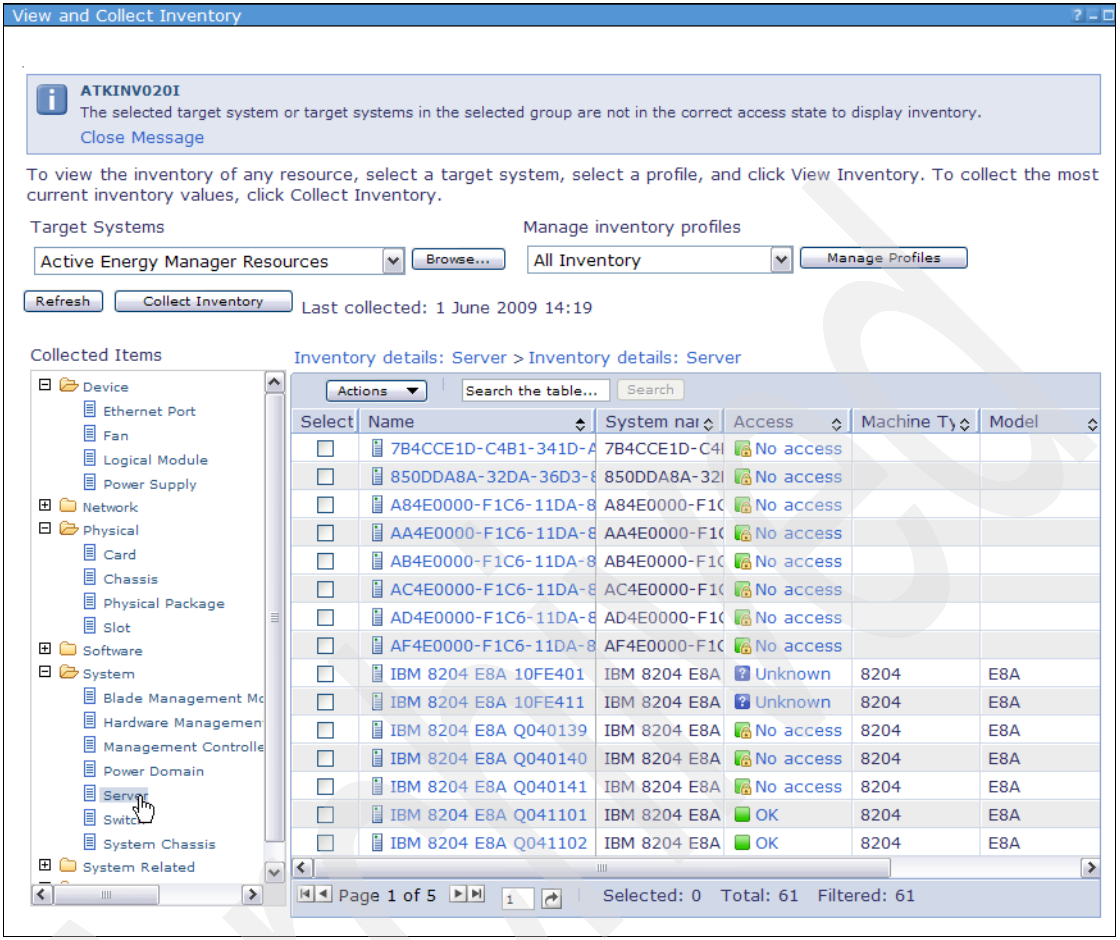


*Figure 6-37 Collecting inventory for the Active Energy Manager group - 2*

## 6.6.3 Monitoring energy usage at the component level

You can use inventory collection to monitor energy usage at a more granular level than is shown in the Active Energy Manager Resources (View Members) display. For example, using the inventory collected for the BladeCenter shown in Figure 6-38, you can view the trend data for a BladeCenter component such as Power Domain 1. This data is not available through the Active Energy Manager Resources (View Members) display.

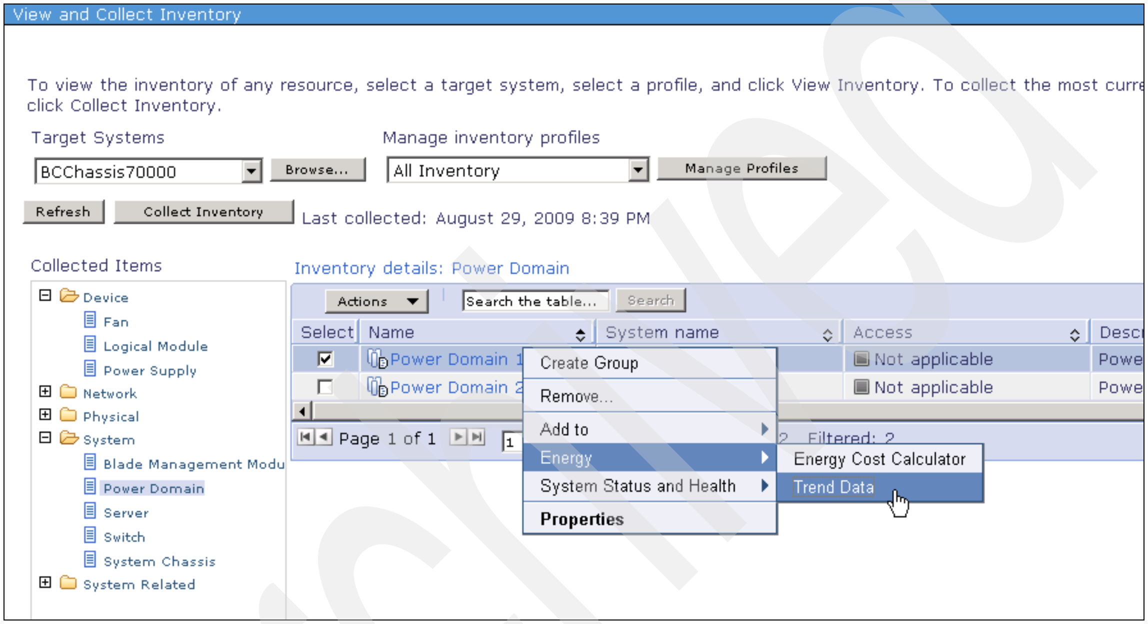


*Figure 6-38 Using inventory collection to monitor energy usage at a component level - 1*

Figure 6-39 shows the actual trend data for Power Domain 1 in this example.

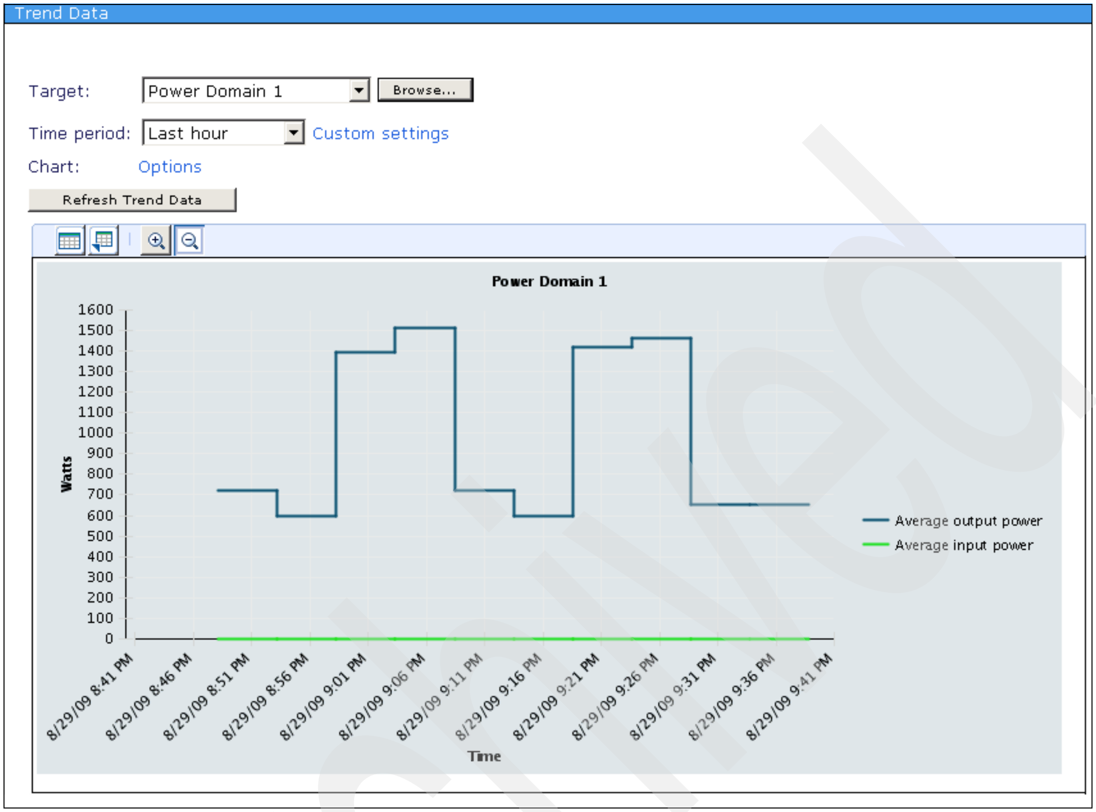


*Figure 6-39 Using inventory collection to monitor energy usage at a component level - 2*

# 6.7 Viewing trend data

We cover the following topics in this section:

- ► Viewing trend data as a chart
- ► Viewing trend data as a table
- ► Viewing current data

The ability to view trending data for a device or group of devices is a powerful tool in the management of power in the data center. Active Energy Manager collects the following types of trend information:

- ► Power consumption
- ► Power cap
- ► Thermal signature

- Humidity and dew point
- CPU speed

The ability to trend humidity and dew point data is new with Active Energy Manager 4.1.1. Humidity and dew point are indicators of how well the air conditioning in a data center is operating. If humidity is too high, then a server becomes more at risk of failing. Using a sensor, users can monitor specific locations within a data center to determine if there is a risk of condensation. In addition, users can set thresholds for humidity and dew point so they can be immediately notified if the value reaches a dangerous level. They can then choose to take appropriate action based on those readings.

If a resource has any associated sensors or is plugged into an intelligent PDU, the externally metered data from these devices can also be viewed.

You can view trend data as a chart (graph) or as a table. We discuss both options in this section.

Note the following points regarding trend data:

- You can set the time interval for viewing trend data from 1 hour up to 1 year, or you can define a custom interval of your choice where the start and end times are both in the past. This option enables you to review historical data.
- Not all Active Energy Manager resources support all types of trending information. To determine which trending functions are supported in each hardware environment, refer to Table 2-1 on page 56. For the most up-to-date list, go to:

  http://publib.boulder.ibm.com/infocenter/director/v6r1x/topic/aem_410/frb0_r_HW_reqs_managed_systems.html
- When viewing trend data as either a chart or as a table, stopping and starting IBM Systems Director server causes gaps in the data. Stopping and starting metering for the resource also causes gaps in the data.
- Although Active Energy Manager displays the current power cap value as trend data, this value is normally fixed and, therefore, does not vary over time unless you specifically change it.
- The minimum and maximum power readings shown in the trend data are representative of measurements made within the power managed object. Some external AC power meters might not be able to capture or display the same peaks as are displayed using the Active Energy Manager power meter. This feature is due to differences such as the averaging periods of the Active Energy Manager power meter, the averaging periods of the external power meter, and the filtering behavior of the power managed object's power supply.

- Although Active Energy Manager only polls a system by the minute, firmware within each power managed object is recording data and statistics hundreds of times per second, including instantaneous minimum and maximum power consumption. This feature allows an operator to view the variation in a given workload, including any sudden peaks in power consumption that might not appear in the average consumption over a polling interval. These instantaneous values can be accessed through the properties of the resource.

## 6.7.1 Viewing trend data as a chart

Being able to monitor trend data as a chart (graph) is a powerful tool for data center administrators because they can easily see how the different metrics are changing over time. The following trend data charts display:

- Power trending (upper) chart
- Environmental trending (lower) chart

### Power trending (upper) chart

The power trending chart provides the following information:

- Power consumption

  Data center administrators can easily see how power consumption is changing over time, and use the information to predict the power consumption of the data center at various times of the day, week, or month. They can also identify anomalies, manage power consumption when electrical demands or costs are high, and determine appropriate power caps if applicable.

- Power cap

  The trending chart also shows the minimum, current, and maximum power cap for the power managed system. Normally, a data center has to provide enough power to cover the rated power consumption (also called label or nameplate power) of all systems connected to the data center's power supply. However, the rated power is usually considerably greater than the actual power consumed by the system. Setting a power cap for each system in the data center to a value which provides adequate power for the system, but is less than the rated power, enables the administrator to calculate the actual power requirements of the data center, instead of using a value based on rated power. Setting a power cap can delay or obviate the need to build additional power supply infrastructure.

### Thermal trending (lower) chart

The thermal trending chart provides the following information:

- Thermal signature

  The inlet (ambient) temperature and outlet (exhaust) temperature of the power managed object over time can be measured and displayed on the chart. This gives data center administrators early warning of potentially dangerous rises in processor temperature and allows them to take corrective action.

- CPU speed

  Active Energy Manager charts the effective CPU speed of the power managed system's processors. For example, when a user enables power savings mode, the effective CPU speed shown on the chart reflects the decrease in the current CPU clock speed from its rated value.

- Humidity and dew point

  Humidity and dew point measurements are not a standard metric that can be measured by power managed systems. However, they can be retrieved from special sensors that have been deployed in the data center and discovered by Active Energy Manager. For a list of supported sensors, refer to:

  http://publib.boulder.ibm.com/infocenter/director/v6r1x/topic/aem_410/frb0_r_HW_reqs_metering_devices.html

  We describe the configuration and discovery of sensors in Chapter 9, "Using sensors with Active Energy Manager" on page 417.

> **Note:** Dew point data is not displayed by default in the trending charts. You need to add this data using the chart options. Figure 6-46 on page 231 shows the Dew point options. See "Setting trend data chart options" on page 230 for details.

## Viewing trend data as a chart

To view trend data as a chart, follow these steps:

1. Click **Active Energy Manager Resources (View Members)** on the Active Energy Manager home page.
2. Right-click a resource, and select **Energy** → **Trend Data** (Figure 6-40).

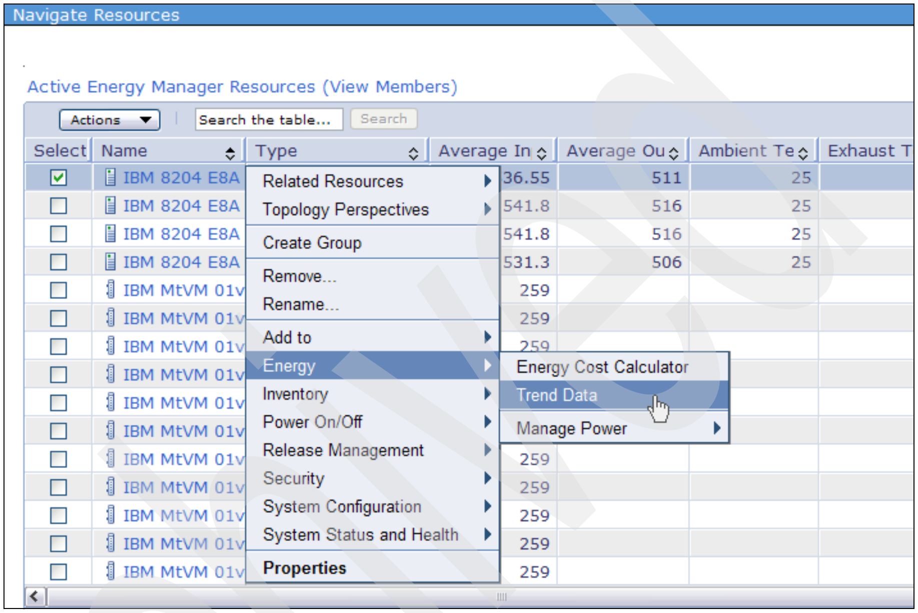


*Figure 6-40 Viewing trend data for an Active Energy Manager resource - 1*

The Trend Data display opens, similar to that shown in Figure 6-41, power trending data shows in the upper chart and thermal trending data in the lower chart.

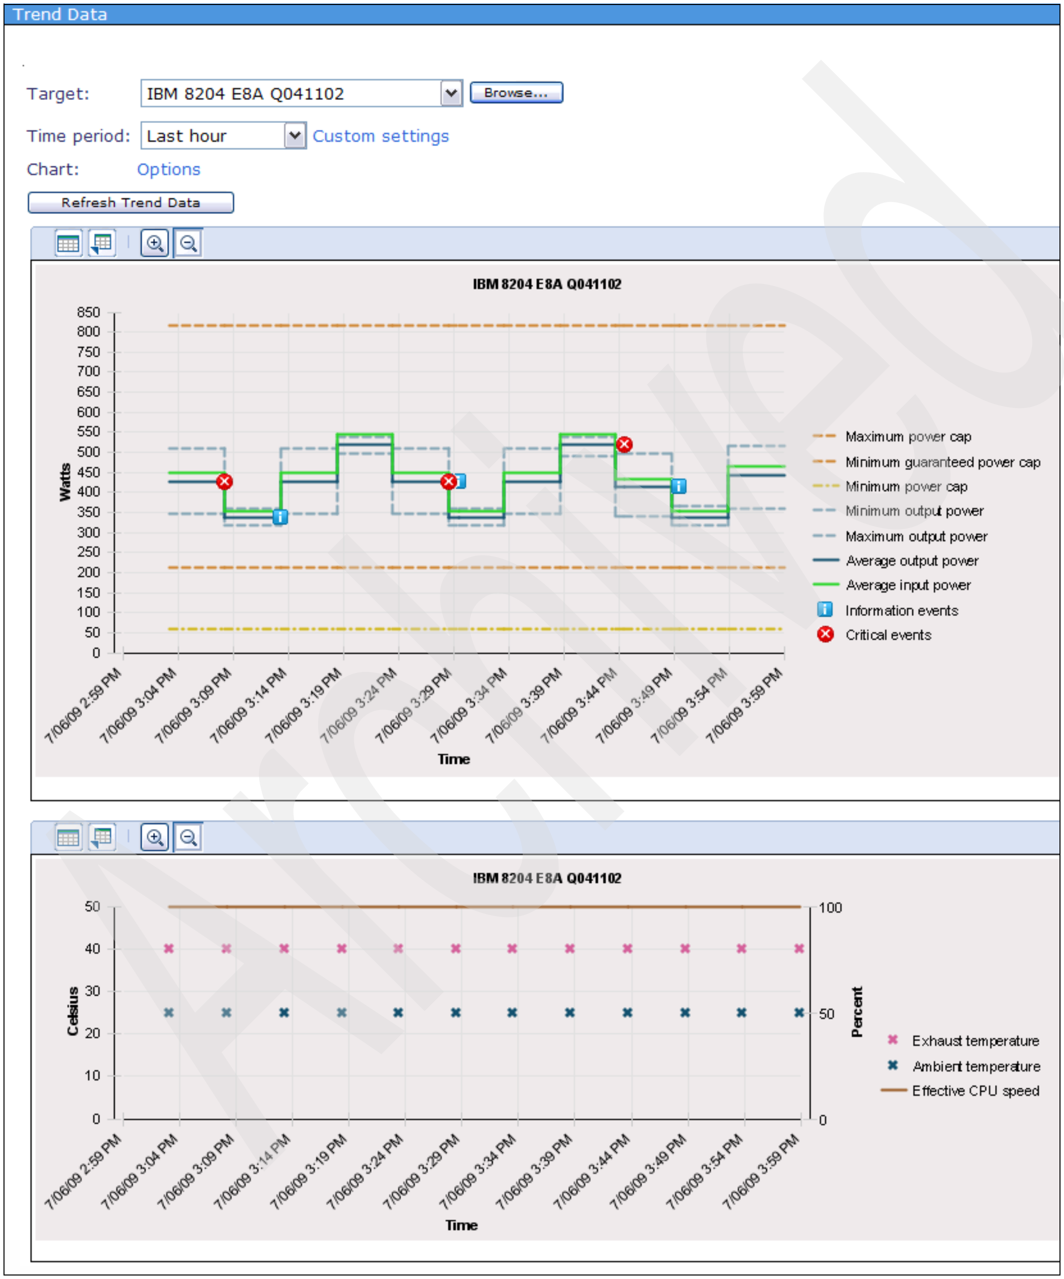




*Figure 6-41 Viewing trend data for an Active Energy Manager resource - 2*

**Note:** For a description of the different colored lines, terms, and event icons shown in Figure 6-41, refer to "Reading the trend data chart legend" on page 231.

In Figure 6-41, notice the following features:

- Graph data
  - Horizontal lines indicate metered values averaged over a polling interval, where the default polling (sampling) interval is 5 minutes.
  - Crosses indicate point in time values. Active Energy Manager does not trend environmental values for a group.
  - After 7 days, values are averaged over an hour and shown as a line chart.
  - When Active Energy Manager trends power values for a group, the values for the members of the group are shown aggregated.
  - The trend data chart refreshes automatically unless you select to show the trend data for a (fixed) period in the past.
  - Not all metrics are shown for all types of Active Energy Manager resources.
  - Both upper and lower charts are always shown even if the resource does not have metered values for one of them.
  - There is a difference between the average input power and the average output power which can be accounted for by power losses in the power supply as the power is converted from AC to DC.
  - There is a difference between the ambient (inlet) temperature and exhaust (outlet) temperature which can be accounted for by heating of the air as it passes through the metered resource.
  - You can only view trend data for one Active Energy Manager resource at a time, unless you have created a group of Active Energy Manager resources, as described in 6.4, "Creating a group of Active Energy Manager resources" on page 195.
  - Trend data for a group shows aggregate data for the devices in the group, but only for input power. (No environmental or CPU data is trended for a group.) Note that all devices in the group must be power managed by Active Energy Manager.
  - If the resource is associated with a metering device, such as a sensor or intelligent PDU, all of the available data from the externally metered device is shown for the selected time period. This is true even if the resource was not associated with the metering device for the entire time period.

- Buttons
  - You can select a different resource to view trend data for by clicking **Browse** and selecting the new object.
  - You can force a refresh of the data by clicking **Refresh Trend Data**. Use this function each time that you change the target resource, time period, or chart options, because the chart or table might not refresh automatically in these cases.
  - Click the magnifying glass icon () to zoom in or zoom out.
- Events
  - Are displayed as red (critical), yellow (warning), and blue (informational) icons, depending on their severity. Refer to “Reading the trend data chart legend” on page 231 for more information.
  - Signify an occurrence of significance to a task or system, such as the completion or failure of an operation.
  - Can be generated either internally or externally
    - Internal: IBM Systems Director
    - External: SNMP traps received data from external sources, such as Emerson-Liebert SiteScan or a SynapSense SNMP agent.
  - By hovering the mouse pointer over an event icon on the graph, you can view a description of the event that occurred. In the trend data table, these events are shown in the Event column.
  - Only one event icon is shown on the chart for a given time and data type, and this event icon is the highest severity message for the time. If there are multiple events generated at the same time, use the table view to see the additional event information.
- Clickable options

  The time intervals specified in the Custom Settings drop-down menu are dynamic. That is, for the time that you specify (except custom interval), the data refreshes according to the data refresh interval being updated (see 6.5.1, “Configuring settings globally” on page 204). For example, if you select **Last hour**, the chart is updated constantly to show the data for the previous hour (Figure 6-42).


*Figure 6-42 Selectable time intervals for showing trend data*

You can choose a different interval or define a custom time interval of your choice where the start and end times are both in the past. To set a custom time interval, click **Custom interval**. (Refer to the next section, “Defining a custom time interval for viewing trend data,” for more information.)

You can change the chart metrics that display (and also for the table view) by clicking **Options**. (Refer to “Setting trend data chart options” on page 230 for more information.)

## Defining a custom time interval for viewing trend data

To define a custom time interval for viewing trend data where the start and end times are both in the past, follow these steps:

1. On the Trend Data display shown in Figure 6-43, click **Custom settings**.

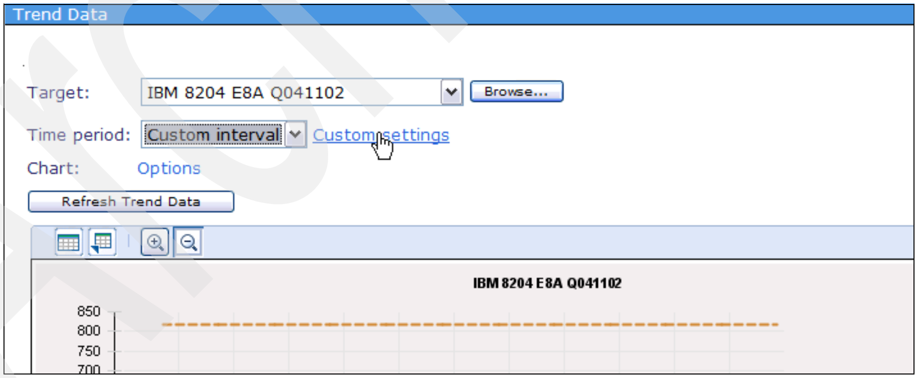


*Figure 6-43 Defining a custom time interval for viewing trend data - 1*

2. In the Custom Time Settings window, shown in Figure 6-44, enter the starting and ending dates and times of the period for which you want to view trend data. Click **OK**.

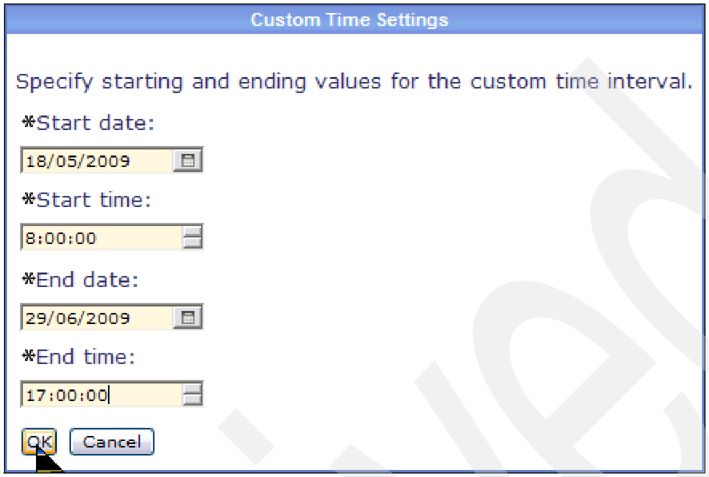


*Figure 6-44 Defining a custom time interval for viewing trend data - 2*

3. Click **Refresh** on the Trend Data display to view trend data for just this period.

## Setting trend data chart options

To select the power and environmental metrics to display on the trend data chart for a resource by changing the chart options, follow these steps:

1. On the Trend Data display shown in Figure 6-45, click **Options**.

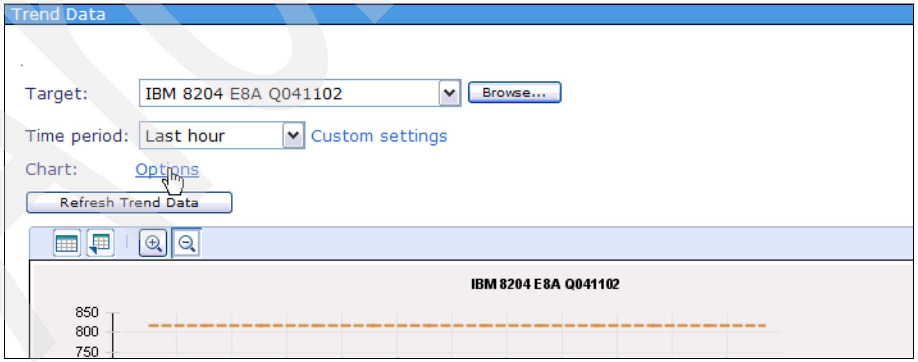


*Figure 6-45 Setting the trend data chart options - 1*

2. Select or clear the chart options that you want to display for this resource (Figure 6-46). Click **OK**.

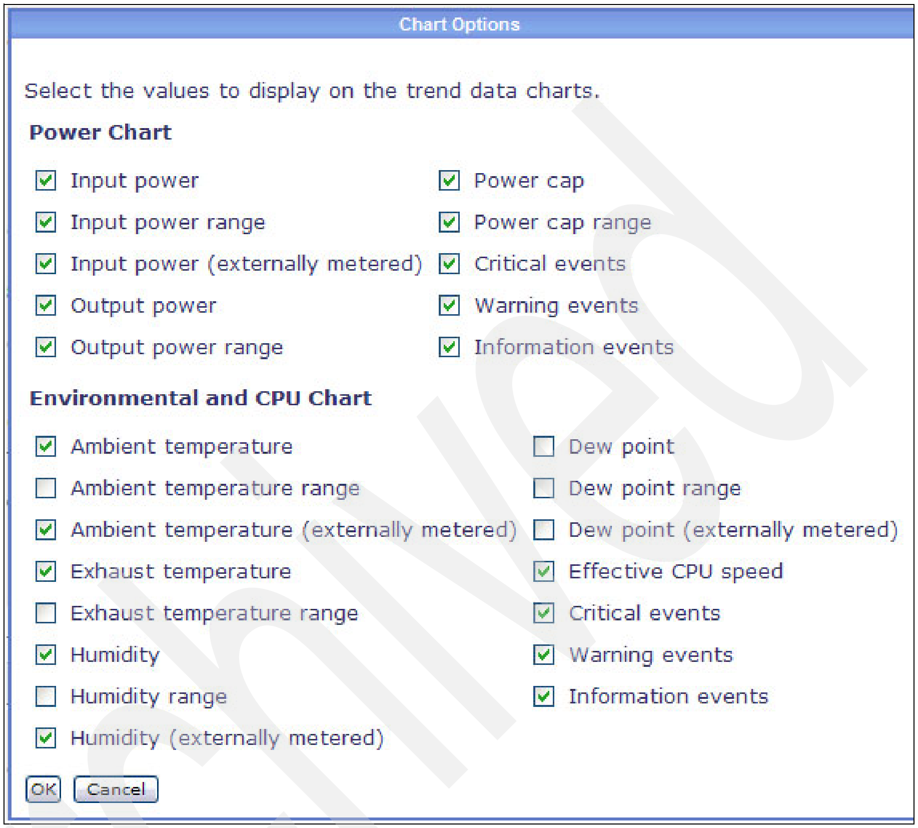


*Figure 6-46 Setting the trend data chart options - 2*

3. Click **Refresh** on the Trend Data display.

## Reading the trend data chart legend

The charts include legend at the side, as shown in Figure 6-41 on page 226. This legend indicates the meaning of each of the lines on the chart. Depending on the resource for which you are viewing trend data, some of the chart metrics might not be available. If a particular resource does not support a certain metric then the metric is not shown.

Table 6-1 lists the complete set of chart metrics that are available, how they translate to the chart options, and what they mean.

*Table 6-1 Trend data chart and table metrics*

| Chart metric | Chart option | What the metric means |
|---|---|---|
| Power (upper chart) | | |
| Average input power | Input power | ► BladeCenter power domain: The power in watts (AC) currently being consumed by the domain.<br>► Internally or externally power meter-enabled managed objects: The power in watts (AC) currently being consumed, as reported by the power meter.<br>► Non-power meter-enabled managed objects: The nameplate watts (AC), also called label power.<br>**Note:** Input power refers to power entering the managed object from the grid. Input power is always alternating current (AC). |
| Average input power (externally metered) | Input power (externally metered) | The average input power that is supplied to a device that does not support power metering natively. The device can either be connected to a PDU+ or to a sensor. |
| Minimum input power<br><br>Maximum input power | Input power range | The same as for average input power but reflects the minimum or maximum AC power consumed during the polling interval rather than the average. |
| Average output power | Output power | ► BladeCenter power domain: The power in watts (DC) currently being consumed by the domain.<br>► Internally or externally power meter-enabled managed objects: The power in watts (DC) currently being consumed, as reported by the power meter.<br>► Non-power meter-enabled managed objects: The nameplate watts (DC), also called label power.<br>**Note:** Output power refers to power exiting the managed object's power supplies which converts alternating current (AC) from the grid to direct current (DC) for use by the object's components such as CPU, memory, fans and disk drives. |
| Minimum output power<br><br>Maximum output power | Output power range | The same as for average output power but reflects the minimum or maximum DC power consumed during the polling interval rather than the average. |
| Minimum power cap | Power cap range | For a managed object that supports power capping, the minimum hard power cap setting in watts. |

| Chart metric | Chart option | What the metric means |
|---|---|---|
| Minimum guaranteed power cap<br><br>Maximum power cap | Power cap range | For a managed object that supports power capping, the minimum or maximum hard power cap setting in watts. |
| Power cap | Power cap | For a managed object that supports power capping, the current power cap setting in watts. |
| Temp/CPU (lower chart) | | |
| Ambient temperature<br><br>Average ambient temperature | Ambient temperature | The current point-in-time or average temperature of air entering the managed object measured in degrees Celsius or Fahrenheit. |
| Minimum ambient temperature<br><br>Maximum ambient temperature | Ambient temperature range | The same as for (average) ambient temperature but reflects the minimum or maximum ambient temperature during the polling interval rather than the average. |
| Ambient temperature (externally metered)<br><br>Average ambient temperature (externally metered) | Ambient temperature (externally metered) | The current point-in-time or average temperature of air entering the managed object which does not support temperature metering natively. The managed object is connected to a sensor. |
| Exhaust temperature<br><br>Average exhaust temperature | Exhaust temperature | An estimate of the current point-in-time or average temperature of air exiting the power managed object measured in degrees Celsius or Fahrenheit and based on the current ambient temperature. The exhaust temperature also depends on the fan speed because the fan speed increases as the temperature increases. The effect is to keep the operating temperature of the ventilated components within safe operating limits. |

| Chart metric | Chart option | What the metric means |
|---|---|---|
| Minimum exhaust temperature<br><br>Maximum exhaust temperature | Exhaust temperature range | The same as for (average) exhaust temperature but reflects the minimum or maximum exhaust temperature during the polling interval rather than the average. |
| Humidity<br><br>Average humidity | Humidity | The point-in-time or average humidity of the resource's surrounding environment measured as a percentage. |
| Minimum humidity<br><br>Maximum humidity | Humidity range | The same as for (average) humidity but reflects the minimum or maximum humidity during the polling interval rather than the average. |
| Humidity (externally metered)<br><br>Average humidity (externally metered) | Humidity (externally metered) | The point-in-time or average humidity of the resource's surrounding environment, metered by an external sensor. |
| Dew point<br><br>Average dew point | Dew point | The point-in-time or average dew point of the resource's surrounding environment measured in degrees Celsius or Fahrenheit.<br><br>The dew point is the temperature to which a given volume of air must be cooled, at constant barometric pressure, for water vapor to condense into water. The dew point is associated with relative humidity. A high relative humidity indicates that the dew point is closer to the current air temperature. Relative humidity of 100% indicates that the dew point is equal to the current temperature, and the air is maximally saturated with water. In this case, there is a risk of droplets forming on solid objects which might damage sensitive equipment. |
| Minimum dew point<br><br>Maximum dew point | Dew point range | The same as for the (average) dew point but reflects the minimum or maximum dew point during the polling interval rather than the average. |

| Chart metric | Chart option | What the metric means |
|---|---|---|
| Dew point (externally metered)<br><br>Average dew point (externally metered) | Dew point (externally metered) | The point-in-time or average dew point of the resource's surrounding environment, metered by an external sensor. |
| Effective CPU Speed | Effective CPU Speed | For a power-metered managed object this is the current CPU speed as a percentage of the rated speed. Values of less than 100% mean that the CPU or CPUs have been throttled back using power savings or power capping. For example, when an administrator enables power savings mode, the effective CPU speed decreases to reflect the decrease in frequency and associated processing power. The CPU trend line shows less than 100% when power savings is switched on, and might also show less than 100% if a power cap is active. |
| Events (either chart) | | |
| Critical | Red circle with a white cross | A severe error indicating that a device has failed or might fail. |
| Warning | Yellow triangle | A warning that a device has suffered an error that might progress to a critical event. |
| Harmless | Blue square | A notification of a change in the environment but does not indicate an error. |

## 6.7.2 Viewing trend data as a table

You can view the data that displays in the charts, shown in Figure 6-41 on page 226, as a table. Click the **Table View** icon on the power trending or thermal trending chart (Figure 6-47).


*Figure 6-47 Accessing trend data in table form using the Table View button*

In the next display, shown in Figure 6-48, notice that the columns shown in the trend data table reflect the metrics chosen to be viewed in the corresponding trend data chart.

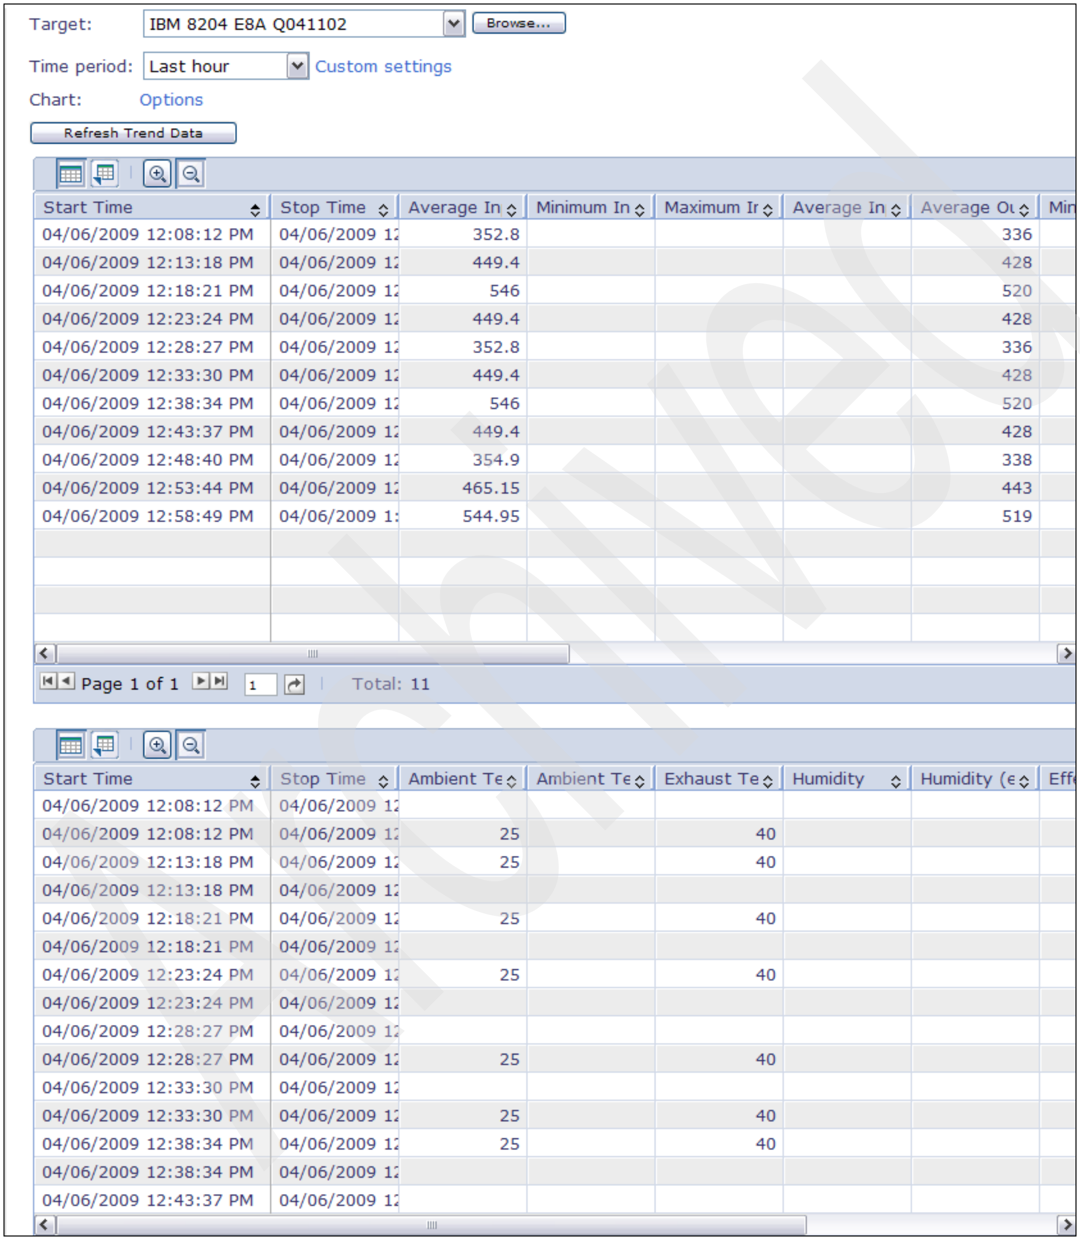


*Figure 6-48 Viewing trend data in tabular form*

In Figure 6-48, notice the following features which are different than the trending data chart:

- The zoom in and zoom out magnifying glass icons have no effect when viewing trend data as a table.
- Events are shown in columns in the table.

### 6.7.3 Viewing current data

You can view point in time data for an Active Energy Manager resource or group of resources. To view the current data for an Active Energy Manager resource or group, simply access the properties as described in 6.5.2, “Configuring settings for an individual resource” on page 208. The metrics that display on the Active Energy tab of the resource’s properties are described in the Active Energy Manager 4.1.1 Information Center:

http://publib.boulder.ibm.com/infocenter/director/v6r1x/topic/aem_410/frb0_t_viewing_aem_properties.html

## 6.8 Exporting trend data

We cover the following topics in this section:

- Exporting chart trend data
- Exporting table trend data

You can export the following types of trend data from the Active Energy Manager database to a file that can be used in an external application:

- *Chart trend data* can be exported in `.BMP` format (Microsoft Internet Explorer) or `.PNG` format (Mozilla Firefox).
- *Table trend data* can be exported in `.CSV` format only.

## 6.8.1 Exporting chart trend data

You can export chart trend data as a `.BMP` file from Microsoft Internet Explorer, and as a `.PNG` file from Mozilla Firefox.

To export a trend data chart, follow these steps:

1. Right-click either the upper or lower trend data charts as required. The context menu displays, shown in Figure 6-49. Select **Save Picture As**.

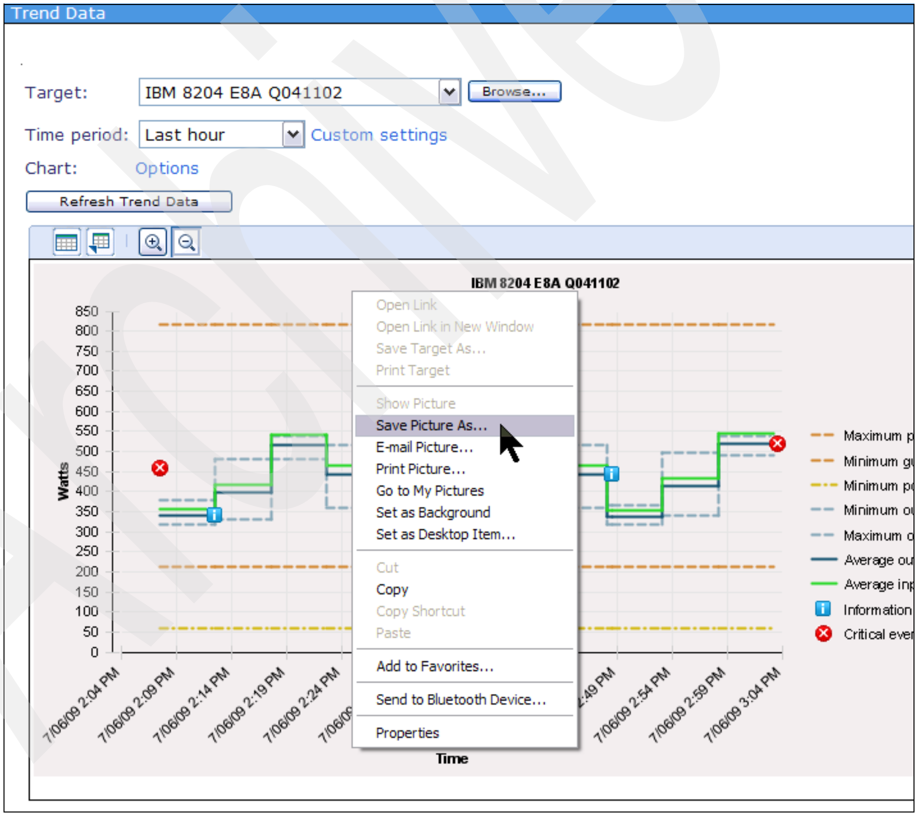


*Figure 6-49 Exporting trend data as a file*

2. Depending on whether you are using Microsoft Internet Explorer or Mozilla Firefox as your browser, slightly different results display as follows:

   – Using Microsoft Internet Explorer, the window shown in Figure 6-50 opens. Enter a name for the file, and choose a location to which to save it. Then, click **Save**.

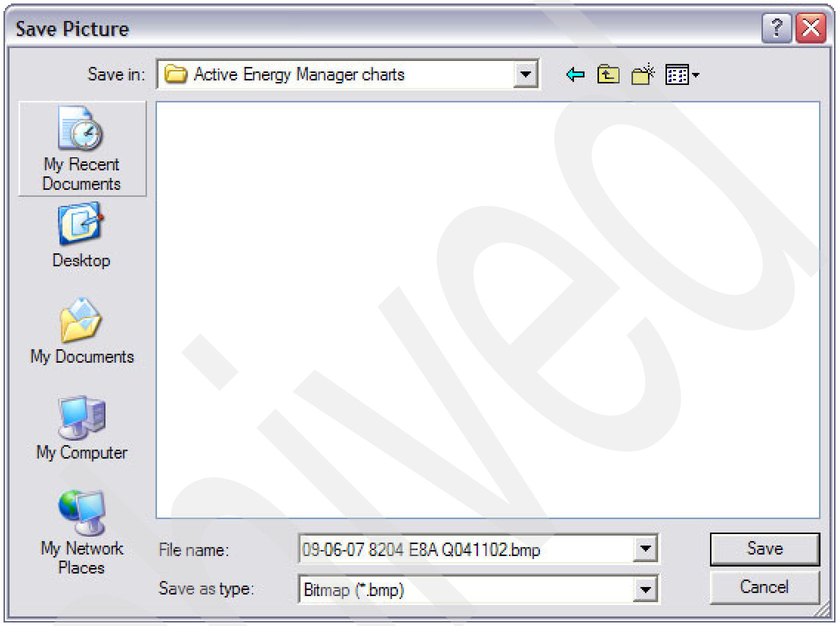


*Figure 6-50 Exporting trend data as a .BMP file*

– Using Mozilla Firefox, the window shown in Figure 6-51 opens. Enter a name for the file, and choose a location to which to save it. Then, click **Save**.

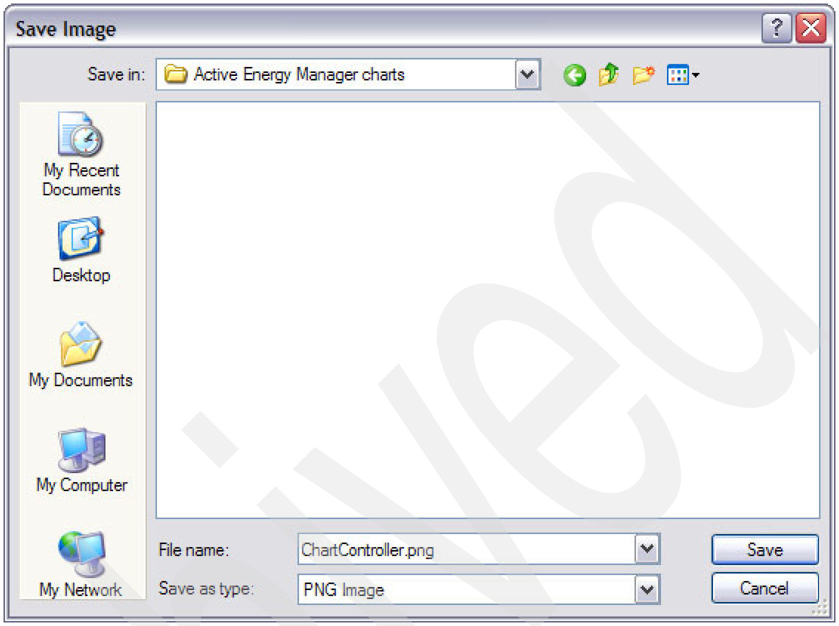


*Figure 6-51 Exporting trend data as a .PNG file*

You can now display the .BMP or .PNG file using any graphics viewer.

## 6.8.2 Exporting table trend data

You can export table trend data as a `.CSV` file from both Microsoft Internet Explorer and Mozilla Firefox. Follow these steps:

1. Click the **Export** icon as shown in Figure 6-52. There is one icon for exporting power trending data and one icon for thermal data. Notice that you can either use the trend data chart or trend data table to export table trend data.

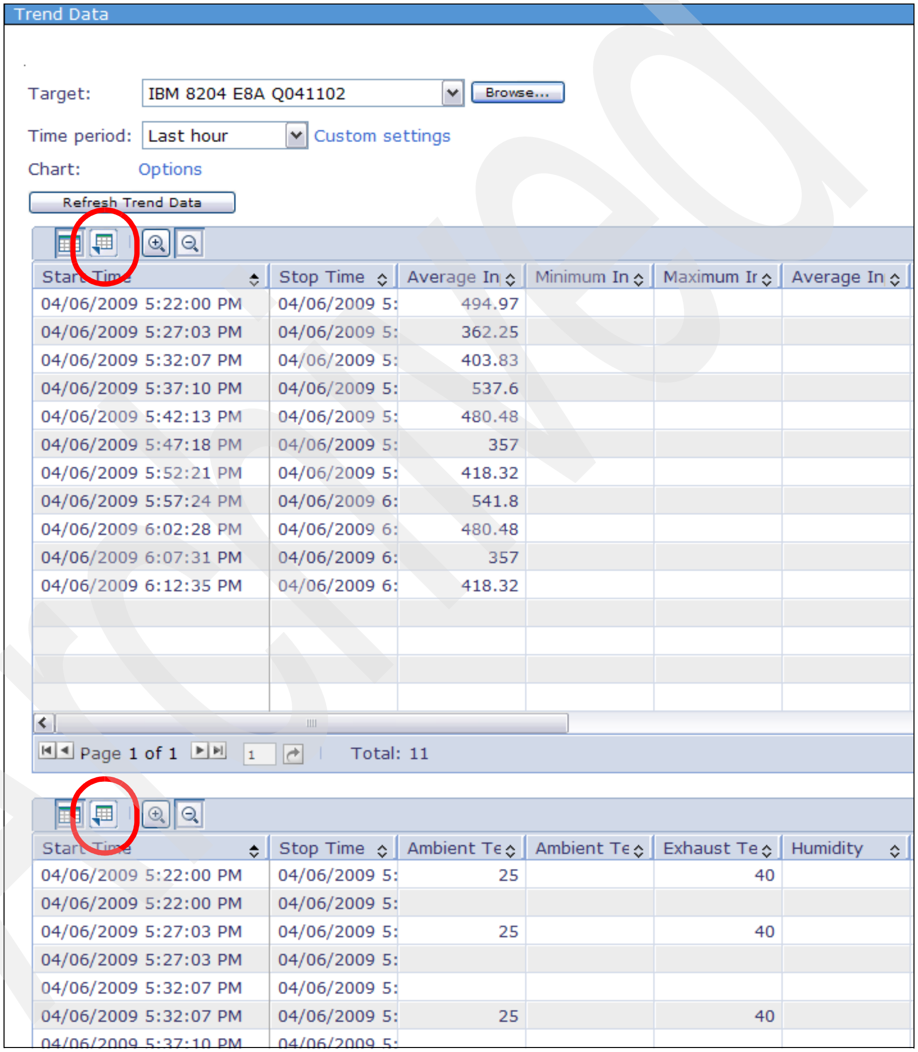


*Figure 6-52 Exporting trend data as a .CSV file*

2. Depending on whether you are using Microsoft Internet Explorer or Mozilla Firefox as your browser, slightly different results displays as follows:

   – For Microsoft Internet Explorer, the window shown in Figure 6-53 opens. You can then choose to either open the `.CSV` file in Microsoft Excel or save the file to a location of your choosing.

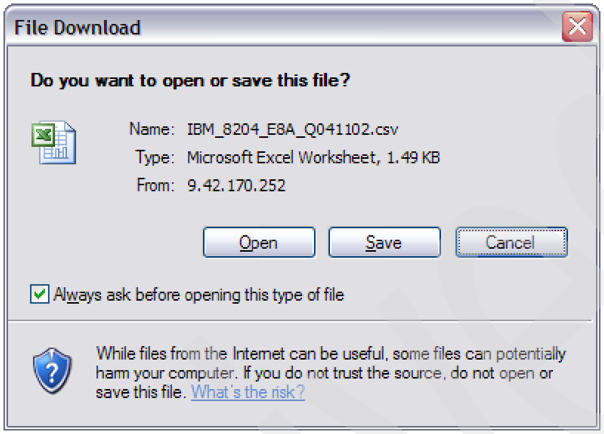


*Figure 6-53 Exporting data in Microsoft Internet Explorer*

   – For Mozilla Firefox, the window shown in Figure 6-54 opens. As for Microsoft Internet Explorer, you can choose to open the `.CSV` file in Microsoft Excel. However, Mozilla Firefox saves the file to your desktop only.

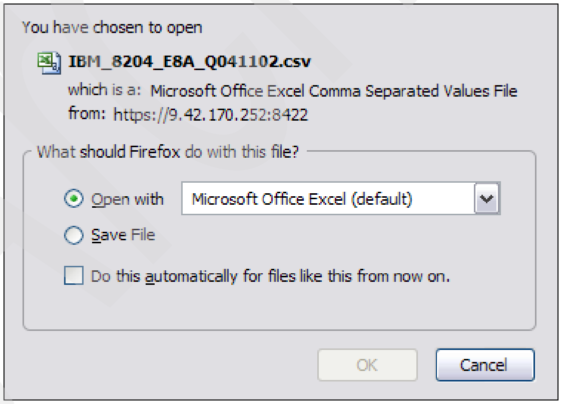


*Figure 6-54 Exporting data in Mozilla Firefox*

## 6.9 Using the Energy Cost Calculator

The Energy Cost Calculator provides a simple but effective way to calculate the total cost of power consumed by an Active Energy Manager resource or group of resources over a specified time period. To use the Energy Cost Calculator on a group, all members of the group must be energy-managed resources, and all members must have polling data for at least half of the selected time interval.

You can use the Energy Cost Calculator on the following resource types:

- BladeCenter chassis
- BladeCenter power domains
- IBM Power Systems
- System x
- System z
- Metered devices
- Intelligent PDUs
- Uninterruptible power supplies

Before using the Energy Cost Calculator, you need to make sure that there is valid energy cost data in the Active Energy Manager settings, especially the cost per kilowatt-hour and the cooling rate factor. Refer to 6.5, "Configuring Active Energy Manager settings" on page 204. Note that you need SMManager authority or higher to change the settings.

The cooling rate multiplier is the amount of power required to cool the managed system divided by the amount of power the managed system consumes. The cost of cooling is therefore equal to the cost of power required to run the managed system multiplied by the cooling rate factor. The default is 1.5. The total cost of power to run a power managed system is the sum of the cost of power to run the managed system plus the cost of power to cool it. To calculate the cost of energy to power the resource only (no cooling), set the cooling rate multiplier to 0.

In addition, the power cost calculator retrieves the value for the nameplate power (also called label power) from the resource and compares it to the actual power being consumed and presents the result as a meter.

Note that you need to have data for at least 50% of the time interval used to calculate the energy cost over, otherwise an error is returned.

The cost is calculated as follows:

```
Total cost = (metered power x time x energy price) + (metered power x
time x energy price x cooling rate multiplier)
```

To use the Energy Cost Calculator, follow these steps:

1. Click **Active Energy Manager Resources (View Members)** on the Active Energy Manager home page.
2. Right-click the Active Energy Manager resource or group for which you want to calculate the energy cost, and select **Energy** → **Energy Cost Calculator** (Figure 6-55).

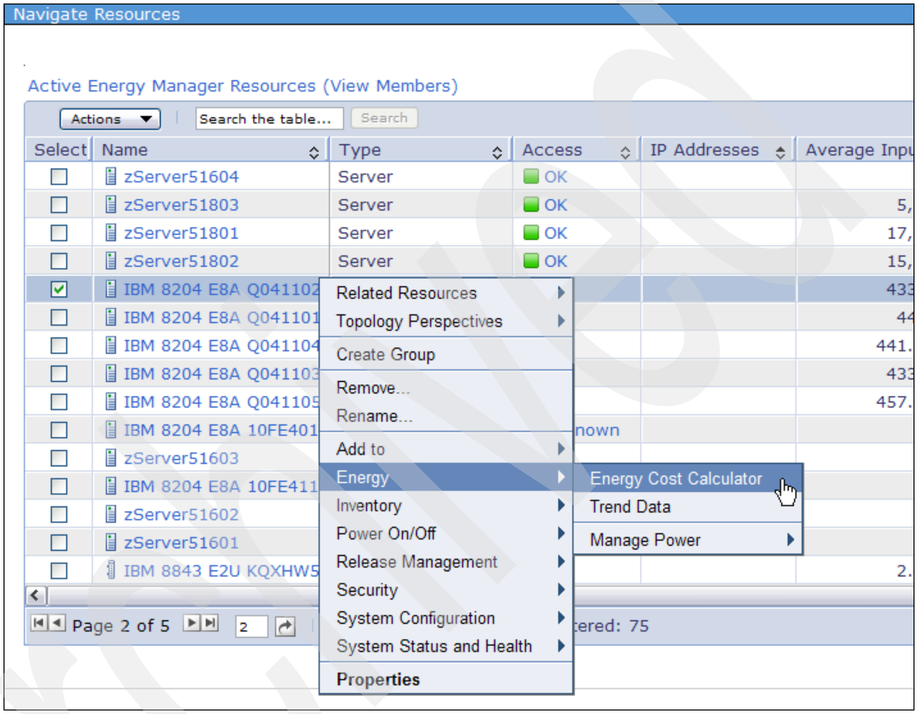


*Figure 6-55 Using the Energy Cost Calculator - 1*

3. You can check or change the energy cost properties by clicking **Cost properties** as shown in Figure 6-56. Click **OK**.

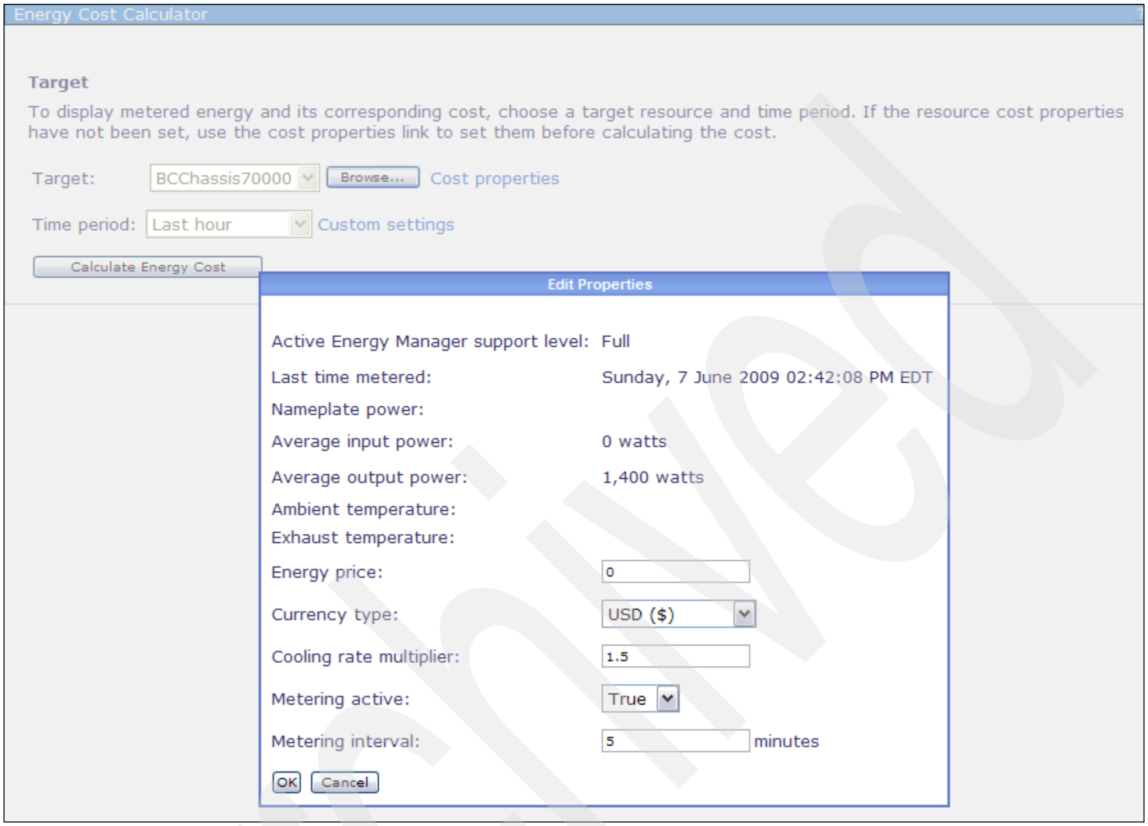


*Figure 6-56 Using the Energy Cost Calculator - 2*

4. Select a time period from the last hour to the last year (or a custom interval by clicking **Custom settings**), and then click **Calculate Energy Cost** (Figure 6-57).

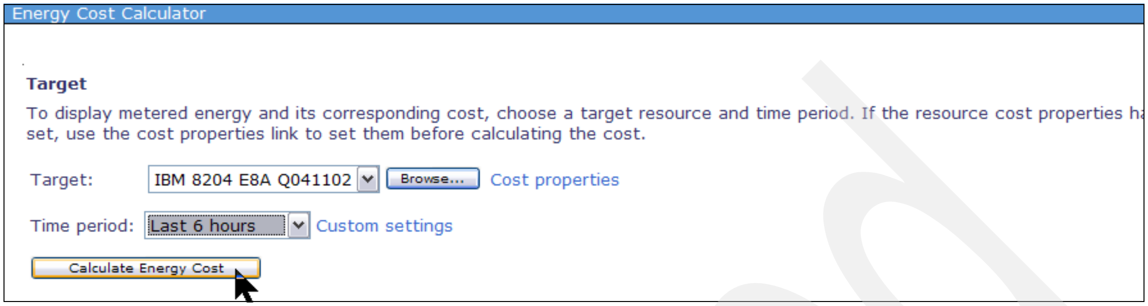


*Figure 6-57 Using the Energy Cost Calculator - 3*

The display shown in Figure 6-58 opens.

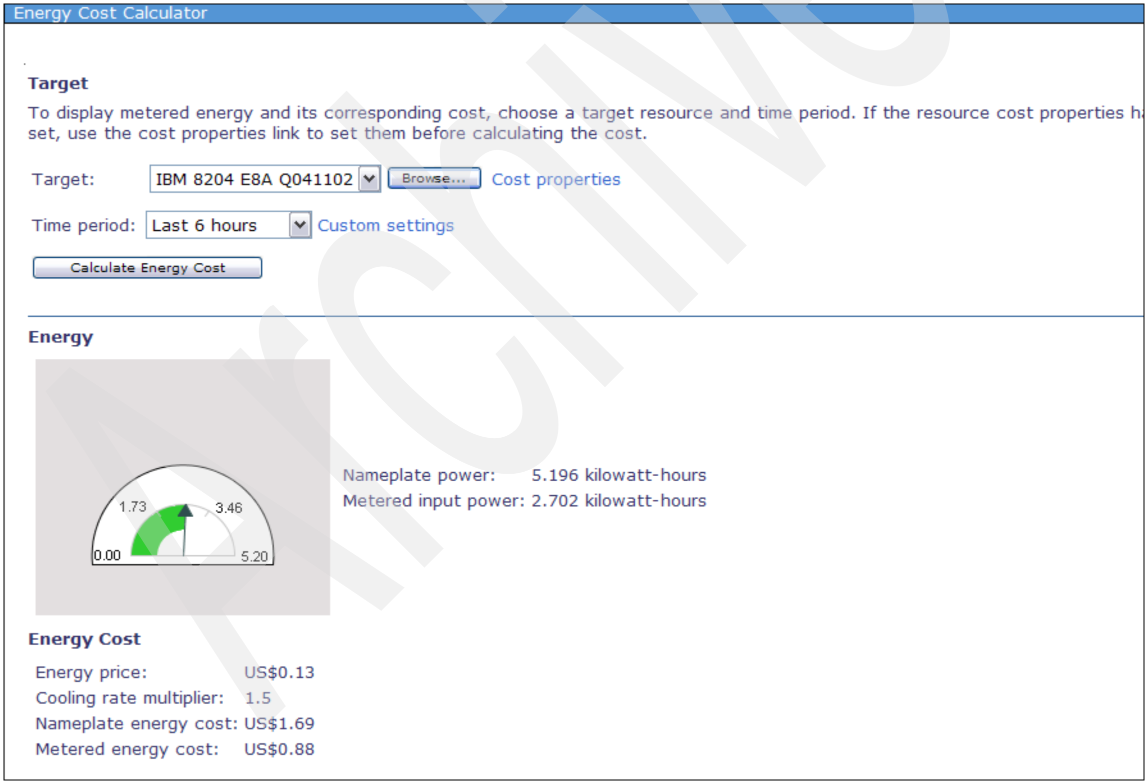


*Figure 6-58 Using the Energy Cost Calculator - 4*

Figure 6-58 on page 247 compares the input power over the period that you specified to the equivalent nameplate (label) power that would have been used over the same period and displays the results on a meter. Notice that the nameplate power is shown only if Active Energy Manager can retrieve it for the resource.

## 6.10 Working with monitors

We cover the following topics in this section:

- Viewing monitors
- Adding a monitor to the dashboard

Active Energy Manager provides the following monitors for resources:

- Average input power
- Average output power
- Ambient temperature
- Exhaust temperature
- Effective CPU speed

These monitors allow the user to retrieve real-time data that can be used to manage the power requirements of data center resources.

Note that not all monitors are available for all Active Energy Manager resources.

## 6.10.1 Viewing monitors

To view the available monitors for an Active Energy Manager resource, follow these steps:

1. From the Active Energy Manager home page, click **Active Energy Manager Resources (View Members)**.
2. On the Navigate Resources display, right-click the Active Energy Manager resource with which whose monitors you want to work, and then select **System Status and Health** → **Monitors** (Figure 6-59).

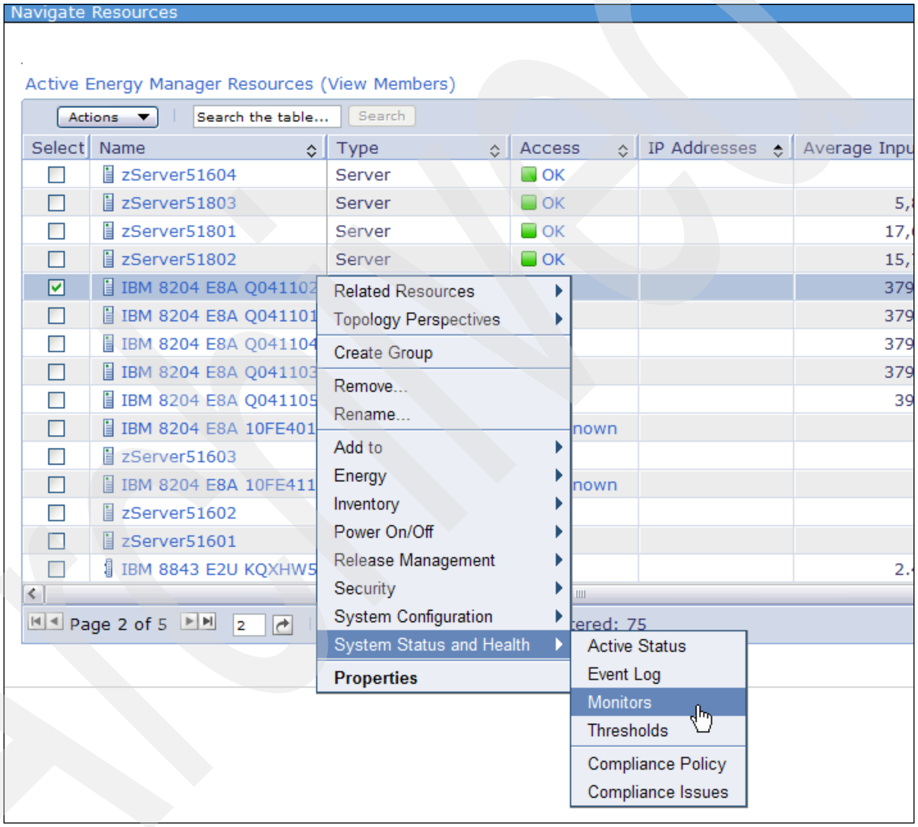


*Figure 6-59 Viewing monitors for an Active Energy Manager resource - 1*

3. In the display shown in Figure 6-60, click **Active Energy Monitors**.

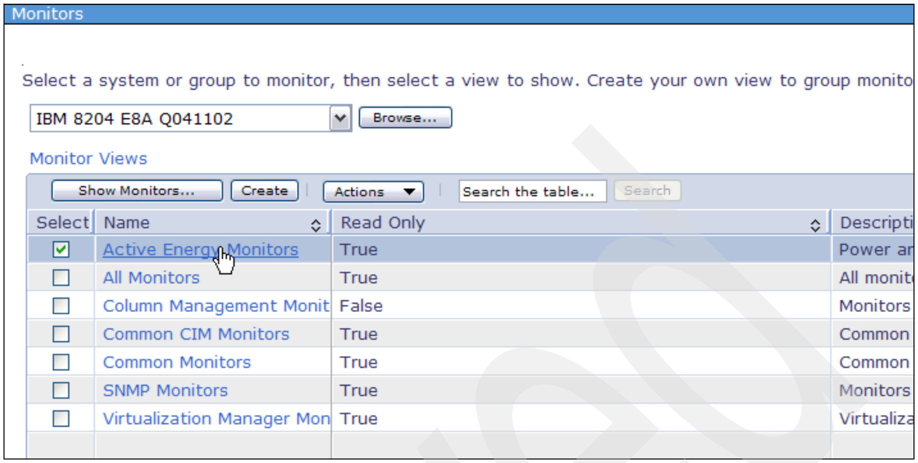


*Figure 6-60 Viewing monitors for an Active Energy Manager resource - 2*

A list of monitors that are available for the resource that you originally selected displays (Figure 6-61).

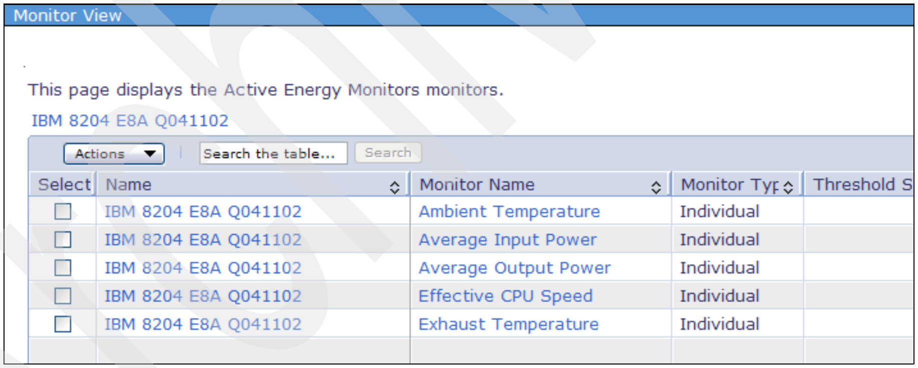


*Figure 6-61 Viewing monitors for an Active Energy Manager resource - 3*

4. Alternatively, you can click **System Status and Health** → **Monitors** to view the display shown in Figure 6-62. Click **Active Energy Monitors** to view a list of monitors for the resource shown in the Browse window.

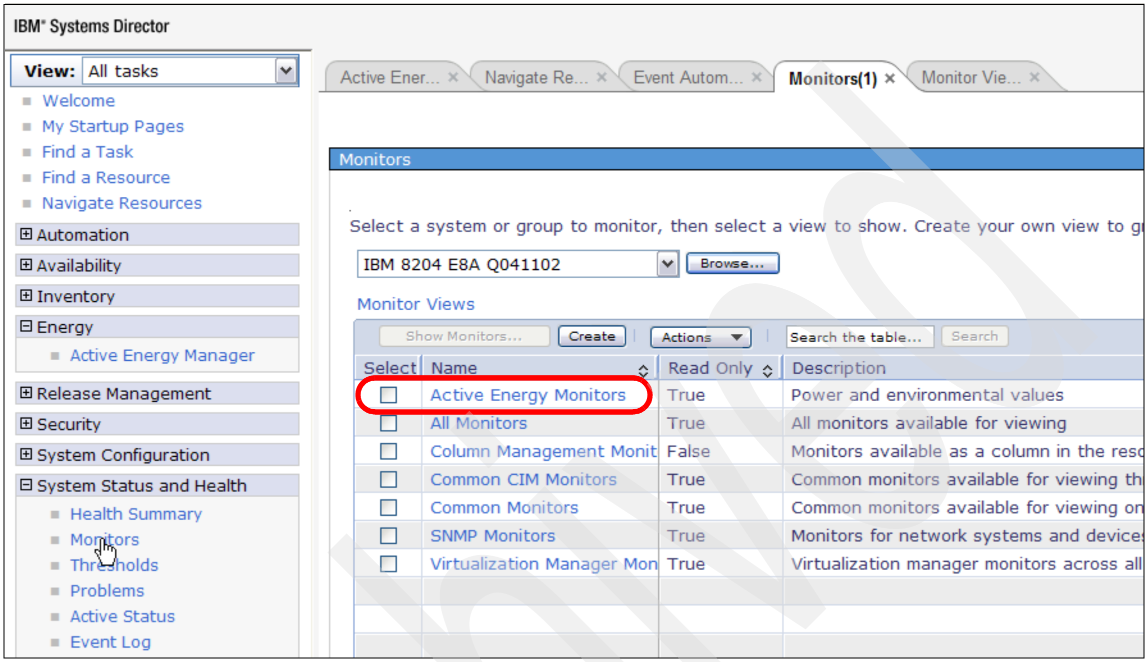


*Figure 6-62 Viewing monitors using System Status and Health*

You can now go ahead and create thresholds for a monitor for this resource, as described in 6.11, “Working with thresholds” on page 254.

## 6.10.2 Adding a monitor to the dashboard

To add a monitor to the System Status and Health dashboard, follow these steps:

1. Right-click the monitor, and select **Add To Dashboard** (Figure 6-63). Notice also the option to Graph the monitor.

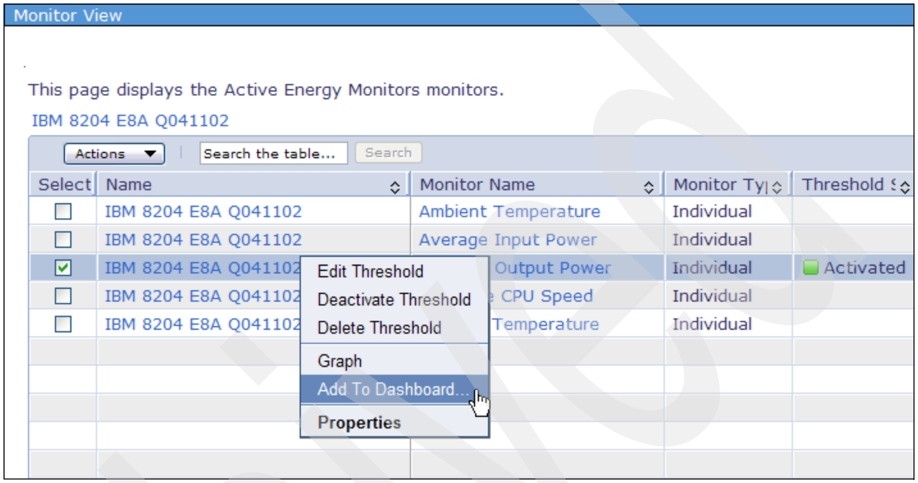


*Figure 6-63 Adding a monitor to the dashboard*

2. You are prompted to enter the name of the monitor and the type of chart that you want to use to display the monitor on the dashboard (Figure 6-64).

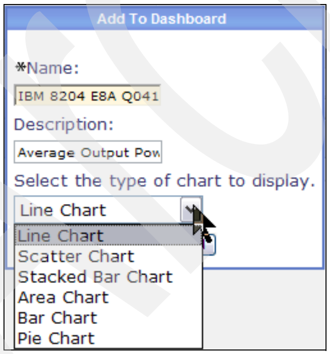


*Figure 6-64 Add to Dashboard dialog box*

3. To view the dashboard, expand **System Status and Health**, and then click **Health Summary** (Figure 6-65).

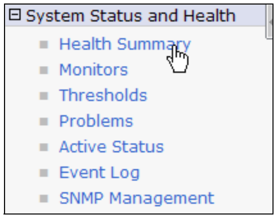


*Figure 6-65 Viewing the dashboard - 1*

A display similar to that shown in Figure 6-66 opens. Notice that you can right-click the gauge that you created and select from options to Edit, Remove, Export Data, and Graph the monitor data.

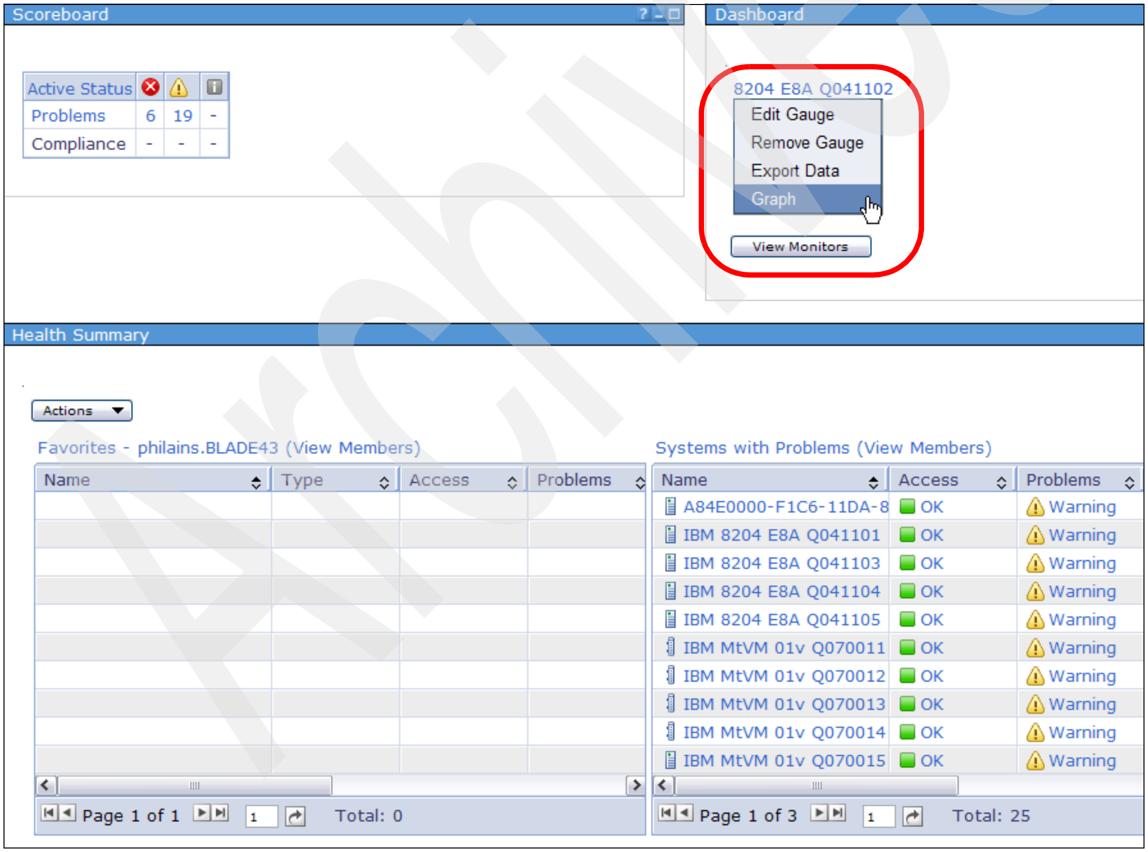


*Figure 6-66 Viewing the dashboard - 2*

# 6.11 Working with thresholds

We cover the following topics in this section:

- Activating a threshold
- Editing, deactivating, and deleting a threshold

For any given monitor, you can set a *threshold*. When triggered, the threshold causes an event to be created and logged. For example, you can create a threshold to cause a critical event to be created and logged when the exhaust temperature of a server rises above 70 °C. IBM Systems Director allows you to send alerts by pager or e-mail. (Refer to *Implementing IBM Systems Director 6.1*, SG24-7694.)

IBM Systems Director also provides tools for automation plans to be activated when a threshold has been reached. For example, when the exhaust temperature of a server rises above 70 °C, a threshold can be triggered, a critical event created, and an automation plan invoked that causes Active Energy Manager to throttle the CPU back to reduce the exhaust temperature. Active Energy Manager monitors and thresholds are integrated into the IBM Director System Status and Health function.

## 6.11.1 Activating a threshold

To activate a threshold for a particular monitor for a power managed system, follow these steps:

1. Access the available monitors for an Active Energy Manager resource by following the steps described in 6.10.1, “Viewing monitors” on page 249
2. Right-click the monitor for which you want to activate a threshold, and select **Edit Threshold** (Figure 6-67).

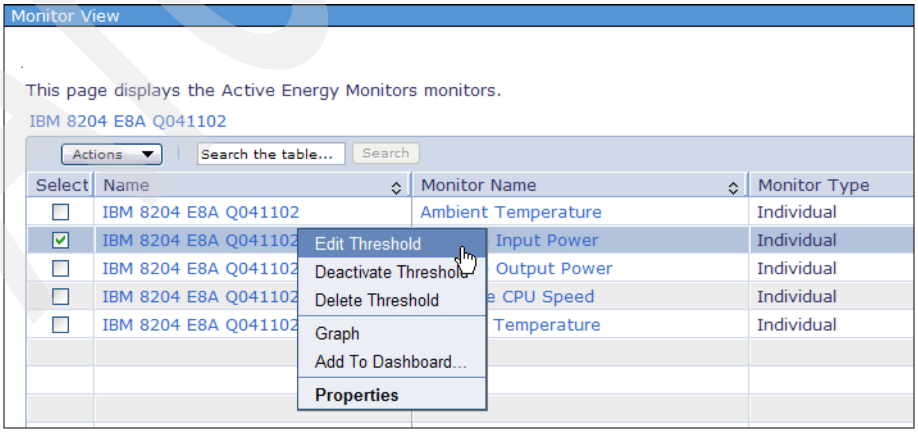


*Figure 6-67 Activating a threshold on an Active Energy Manager resource - 1*

3. On the Threshold display, specify the monitor values for which you want to set thresholds, and select the appropriate level (the **Critical** or **Warning** check boxes) as shown in Figure 6-68. Note that you can specify multiple thresholds for the same monitor. Then, click **OK**.

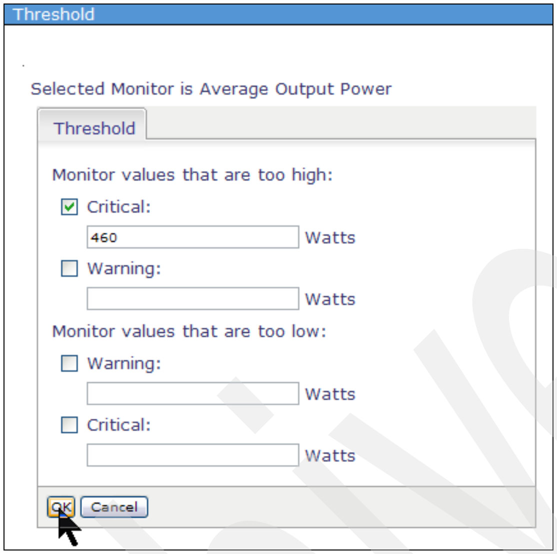


*Figure 6-68 Activating a threshold on an Active Energy Manager resource - 2*

> **Tip:** If you create a threshold on a temperature resource, be sure to use the "Monitor values that are too low" field to ensure that you are not wasting energy by cooling the computer room too much.

The threshold is activated for the power managed object's selected monitor (Figure 6-69).

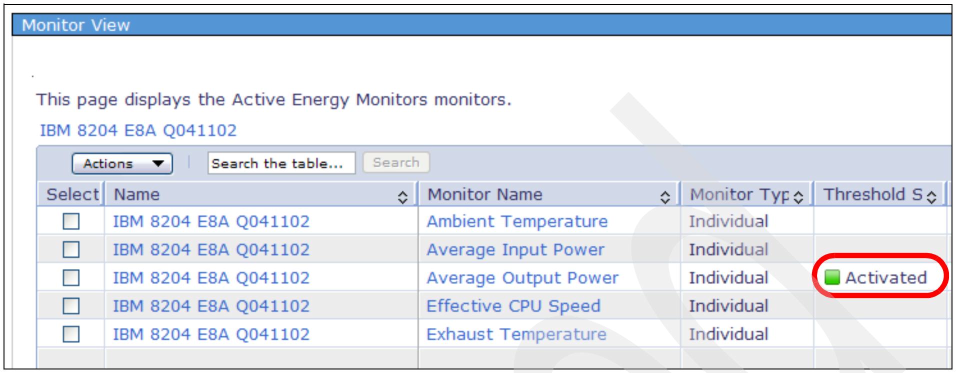


*Figure 6-69   Activating a threshold on an Active Energy Manager resource - 3*

If the threshold is reached, an event is created in the Active Energy Manager event log, as described in 6.12, "Working with events" on page 257.

4. To view a list of thresholds, click **System Status and Health** → **Thresholds** (Figure 6-70).

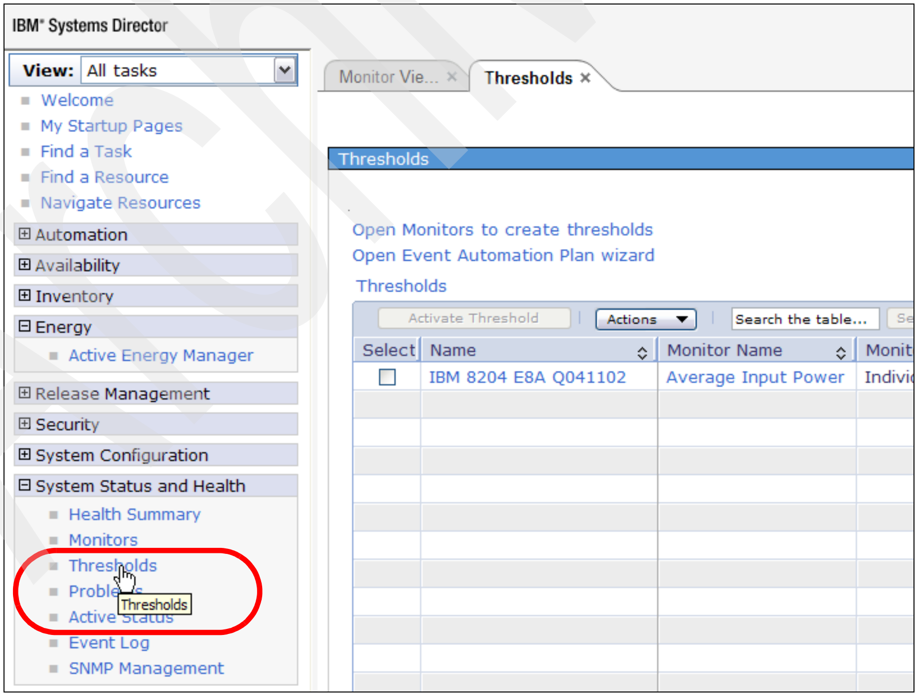


*Figure 6-70   Viewing thresholds using System Status and Health*

## 6.11.2 Editing, deactivating, and deleting a threshold

You can edit, deactivate (without deleting), or delete a threshold by right-clicking the associated monitor and selecting the appropriate menu action, as shown in Figure 6-71.

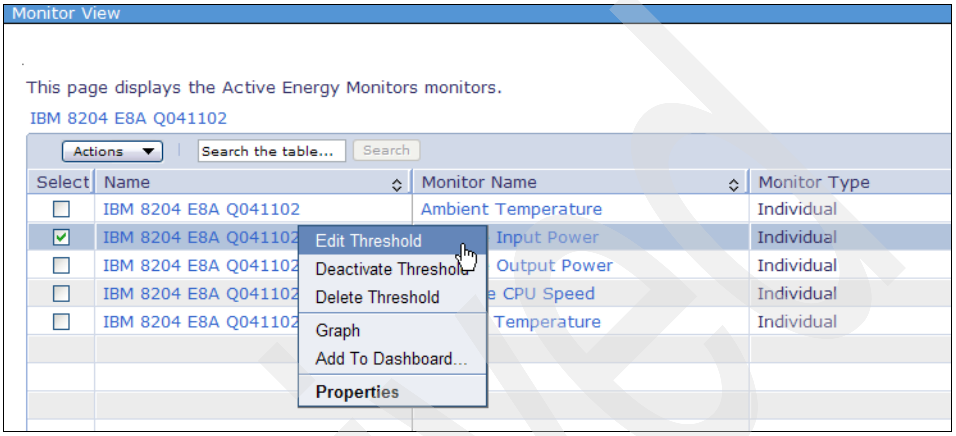


*Figure 6-71 Editing, deactivating, or deleting a threshold*

# 6.12 Working with events

We cover the following topics in this section:

- Viewing events on the trend data charts and tables
- Viewing events in the event log
- Viewing events for a resource

An *event* is an occurrence of significance to either a task or a system, such as the completion or failure of an operation. Events can be generated either internally by Active Energy Manager or externally by third-party energy management solutions capable of sending SNMP traps, such as Emerson-Liebert SiteScan Web or SynapSense sensor networks.

Active Energy Manager logs system events and also events that are triggered when a threshold is reached. Events can appear on any of the trending graphs and can be used to generate actions automatically in response.

The following types of events are logged in the event log:

- *Critical* events, shown as a red circle with a white cross, are severe errors indicating that a device has failed or might fail.
- *Warning* events, shown as a yellow triangle, are warnings that a device has suffered an error that might progress to a critical event.
- *Harmless* events, shown as a blue square, are notifications of a change in the environment. Harmless events do not indicate an error.

## 6.12.1 Viewing events on the trend data charts and tables

When the threshold is reached for a monitor, an event displays on the trend data chart. To view the event, hover the mouse over the event icon as shown in Figure 6-72.

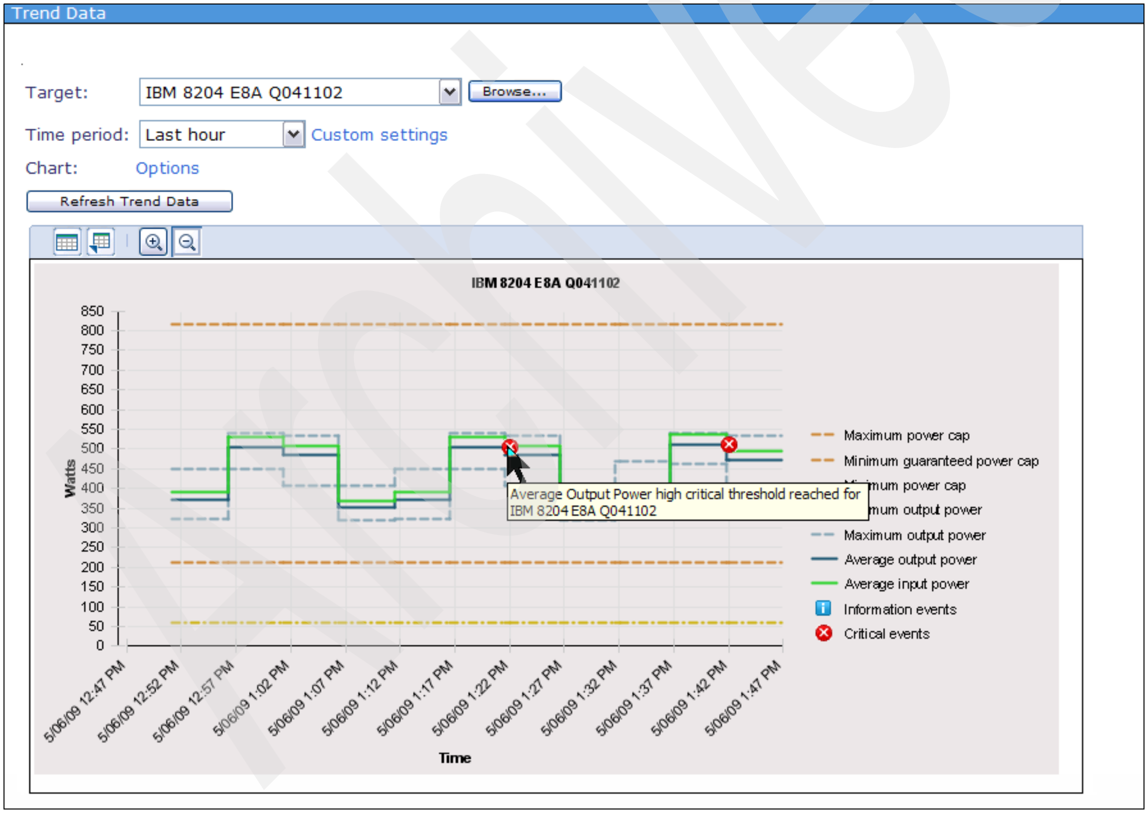


*Figure 6-72 Viewing events on the trend data chart*

The same event is shown in the trend data table (Figure 6-73).

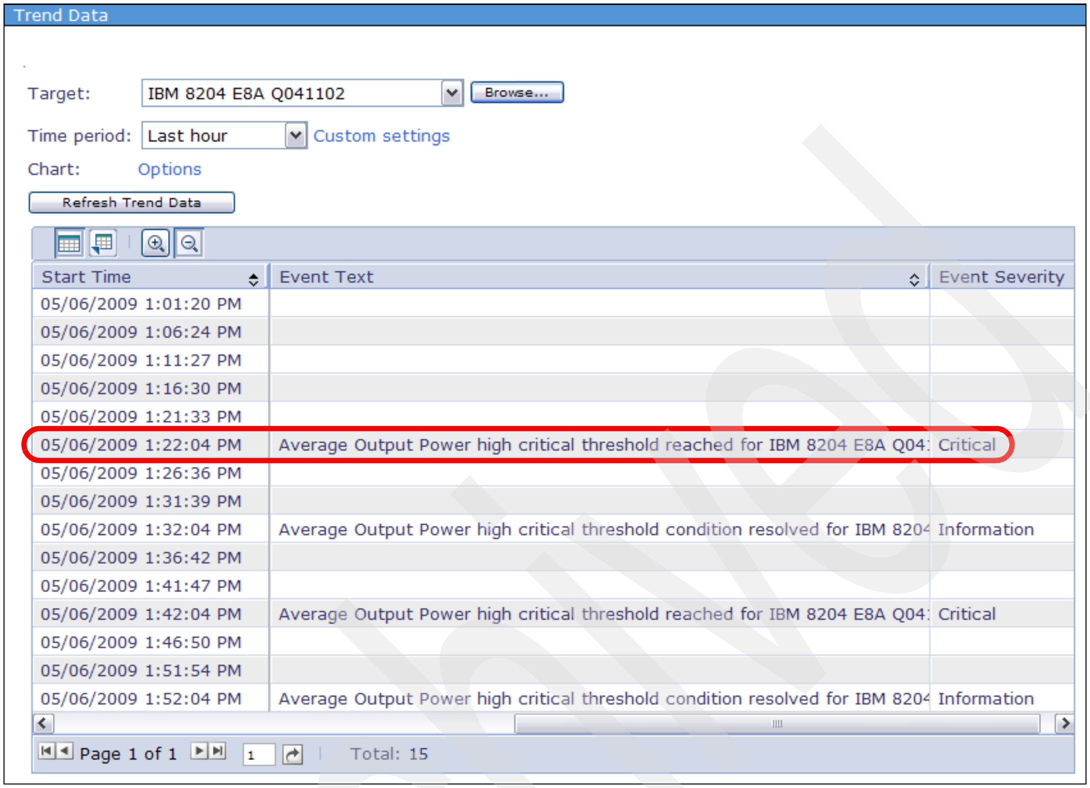


*Figure 6-73 Viewing events on the trend data table*

## 6.12.2 Viewing events in the event log

To view Active Energy Manager events in the IBM Systems Director event log, follow these steps:

1. Access the event log by clicking **Event Log** under **System Status and Health** in the IBM systems Director navigation pane (Figure 6-74).


*Figure 6-74 Navigating to the IBM Systems Director event log*

2. In the Event Log display, shown in Figure 6-75 on page 261, in the Event Filter drop-down menu, there are two filters for Active Energy Manager events:

   – Active Energy Critical Events
   – Active Energy Events

   Select the most appropriate filter, depending on the messages that you want to view.

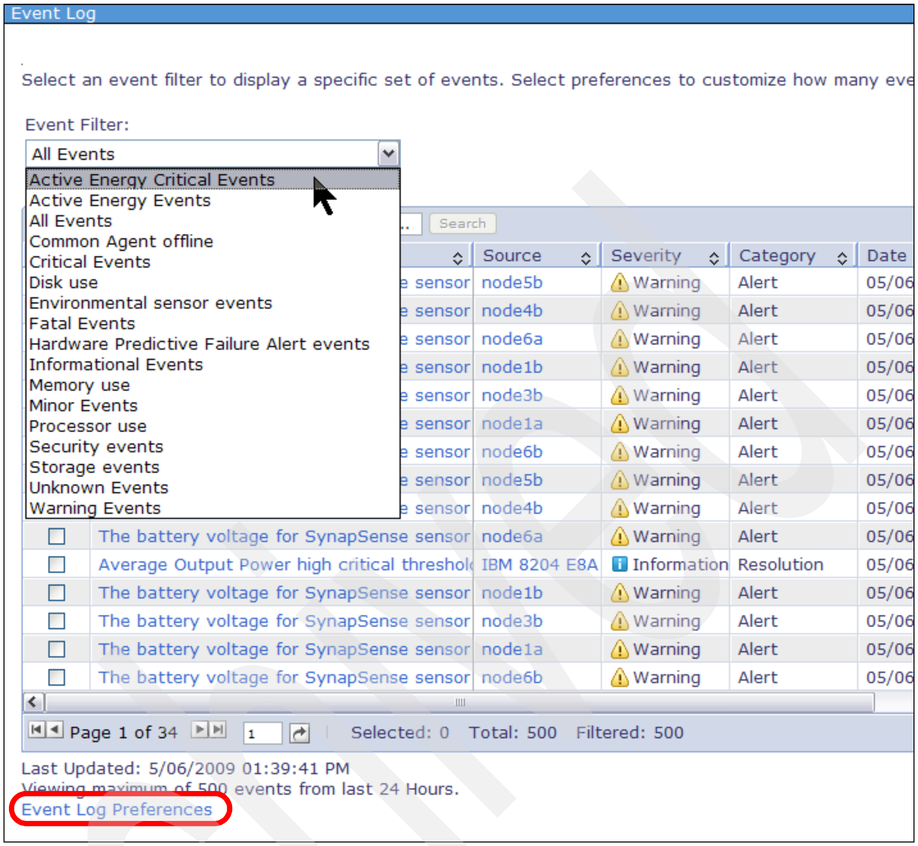


*Figure 6-75 Subsetting the Active Energy Manager events*

3. Click the Event Log Preferences shown in Figure 6-75 to specify a Time Range to display events, and the maximum number of entries to retrieve (Figure 6-76).

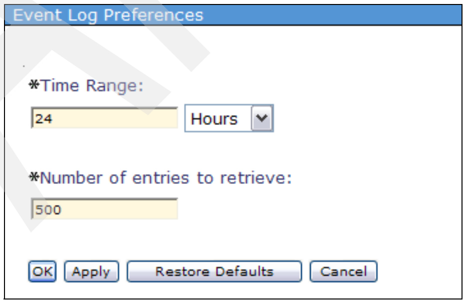


*Figure 6-76 Changing the event log preferences*

4. If you select the **Active Energy Critical Events** filter in Figure 6-75 on page 261, the Active Energy Manager critical events display (Figure 6-77).

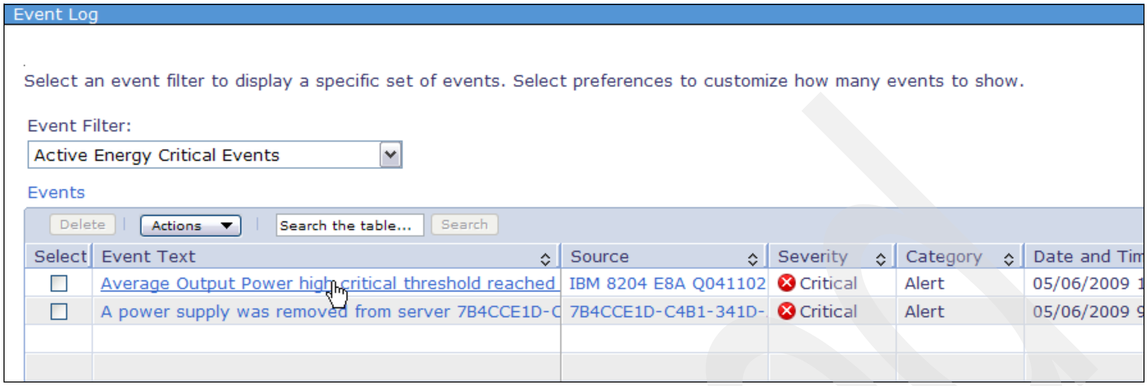


Figure 6-77 Active Energy Manager critical events

5. If you click an event in Figure 6-77, the properties for that event display, as shown in Figure 6-78.

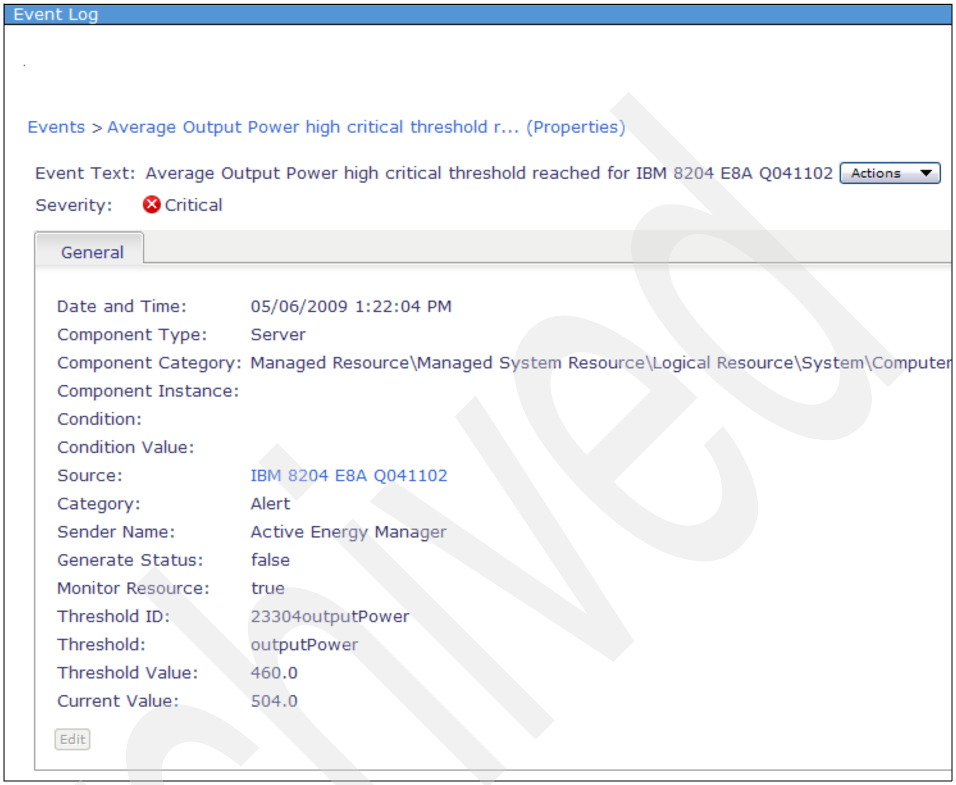


*Figure 6-78 Event properties*

## 6.12.3 Viewing events for a resource

You can also view events for a single Active Energy Manager resource as follows:

1. On the Active Energy Manager home page, click **Active Energy Manager Resources (View Members)**.
2. Right-click the resource for which you want to view events. Select **System Status and Health** → **Event Log** (Figure 6-79).

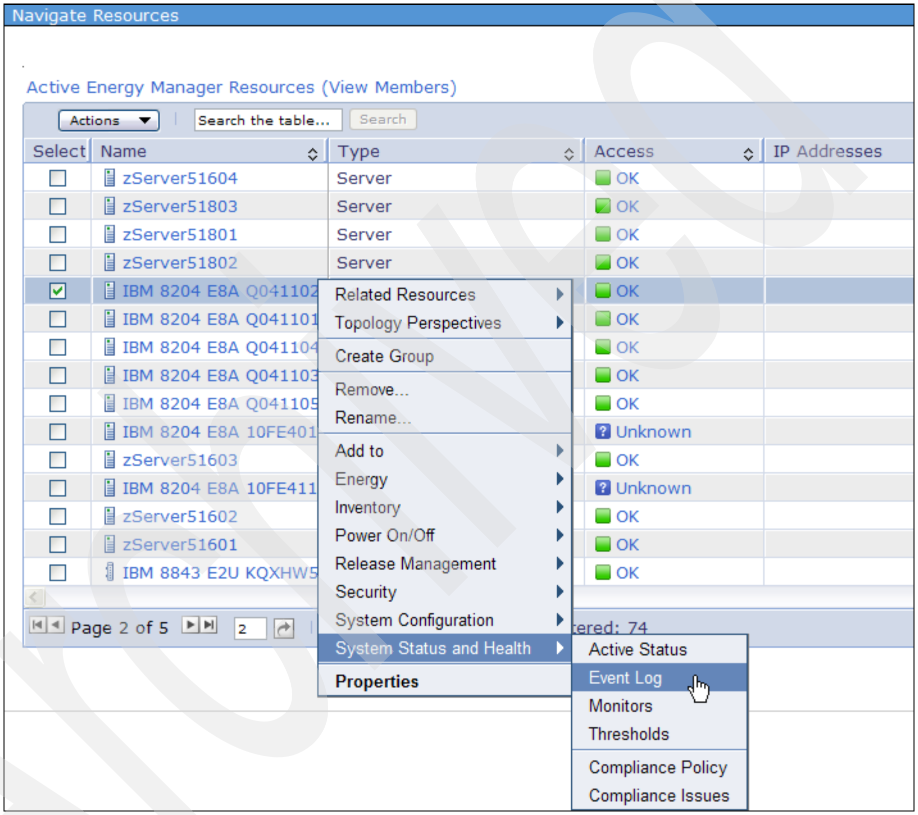


*Figure 6-79 Viewing events for a single Active Energy Manager resource - 1*

3. A display similar to that shown in Figure 6-80 opens. You can use the Event Filter to view only those events in which you are interested.

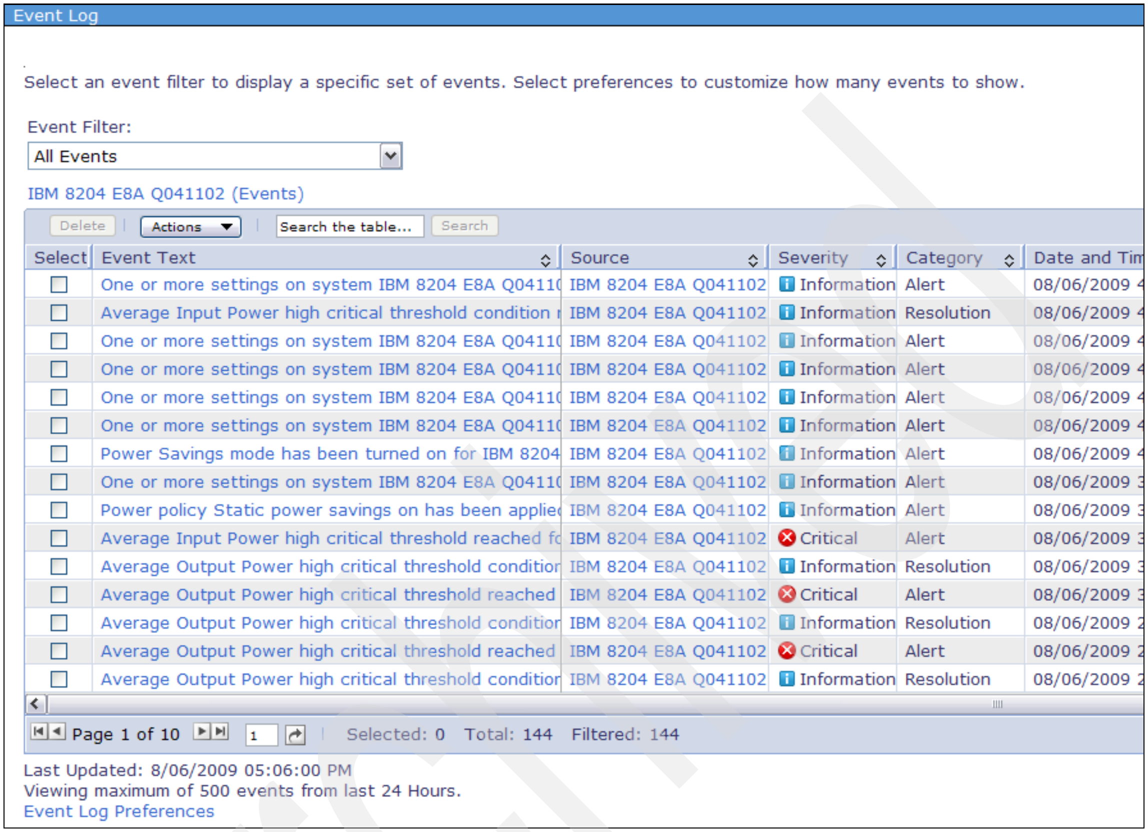


*Figure 6-80 Viewing events for a single Active Energy Manager resource - 2*

## 6.13 Working with automation plans

We cover the following topics in this section:

- Introducing automation plans
- Creating an automation plan
- Working with automation plans

Automation plans automate power-related tasks. You can create an event automation plan to define power event criteria (filters) to trigger power-related event actions. For example, you could create an event automation plan to trigger an event action when the exhaust temperature for a resource is greater than 70°C. You then define what the triggered event action is when this threshold is reached. For example, you can define the event action to be an e-mail sent to the administrator or to have Active Energy Manager throttle back the CPU on that resource.

## 6.13.1 Introducing automation plans

Before you create an automation plan, you need to assemble the various components of the plan as follows:

1. Define your objective. Some example objectives include:
   - I want to receive an e-mail when the ambient or exhaust temperatures rises above a certain value.
   - I want to throttle back the processors on my database server if the exhaust temperature rises above 70°C, and then up throttle the processor when the exhaust temperature falls below 60°C.
   - I want to activate power savings mode on my Power 550 system if the input power goes above 900 watts.
   - I want to impose a hard cap on power to my virtual server environment so that aggregate input power does not rise above 5000 watts.
   - I want to activate power savings on my IBM Power systems if the ambient temperature rises above 30°C, and then disable power savings when the ambient temperature falls below 25°C.
2. Create a threshold for the monitor that you want to use from the following list:
   - Average input power
   - Average output power
   - Ambient temperature
   - Exhaust temperature
   - Effective CPU speed

   Refer to 6.11.1, “Activating a threshold” on page 254 for a description about how to create a threshold.
3. Create a policy that enables the required action. You can choose from the following policy types:
   - Enable system power savings
     - Static
     - Dynamic that favor power savings
     - Dynamic that favor performance
   - Enable system power capping
   - Enable group power capping

   Refer to 7.5.2, “Creating a policy” on page 327 for a description about how to create a policy.

4. Use the Create an event automation plan wizard to create your automation plan, as described in the next section.

Note that it might be helpful to set up a table that contains all the information for the automation plans that you want to implement. Table 6-2 shows the requirements for the example objectives that we listed previously.

*Table 6-2 Collecting automation plan requirements*

| Objective | Monitor threshold | Policy - event action |
|---|---|---|
| Activate power savings mode on my Power 550 system if the input power goes above 900 watts. | Average input power, critical above 900 watts | Enable system power savings, dynamic power savings (favor power savings) |
| Impose a hard cap on power to my virtual server environment so that aggregate input power does not rise above 5000 watts. | Average input power, critical above 5000 watts | Enable group power capping, hard power cap |
| Throttle back the processors on my database server if the exhaust temperature rises above 70°C, then throttle up the processors when the exhaust temperature falls below 60°C. | Exhaust temperature, critical above 70°C, warning below 60°C | Enable system power savings, dynamic power savings (favor power savings)<br><br>Disable system power savings, dynamic power savings (favor power savings) |
| Activate power savings on all my IBM Power systems if the ambient temperature rises above 30°C, and then disable power savings when the ambient temperature falls below 25°C. | Ambient temperature, critical above 30°C, warning below 25°C | Enable system power savings, static power savings<br><br>Disable system power savings, static power savings |

## 6.13.2 Creating an automation plan

In this section, we create a simple automation plan with the following requirements:

- Objective

  Activate static power savings on IBM Power system 8204 E8A Q041102 if the average input power rises above 400 watts.

- Monitor threshold

  Set a critical threshold of 400 watts for the *average input power* monitor for power managed system 8204 E8A Q041102 as described in 6.11.1, “Activating a threshold” on page 254.

Note that the power values in watts for the monitor threshold were arrived at by observing the trend data for the chosen system.

- Policy - event action

  Create a policy to apply static system power savings as described in 7.5.2, “Creating a policy” on page 327.

> **Note:** Before you create the plan, make sure that you have created monitor thresholds and event action policies.

To create an event automation plan, follow these steps:

1. If you have not already, create a policy to apply static system power savings as described in 7.5.2, “Creating a policy” on page 327.
2. From the Active Energy Manager home page, click **Create Automation Plans** under the Automate heading (Figure 6-81).

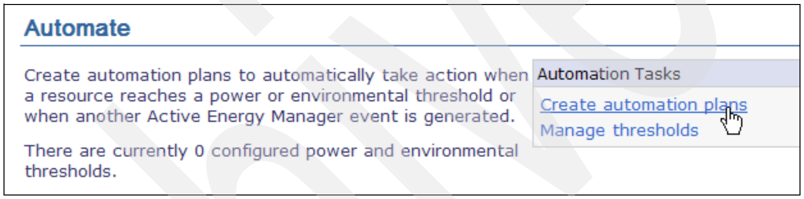


*Figure 6-81 Creating an event automation plan - 1*

3. On the Event Automation Plans display, click **Create** (Figure 6-82).

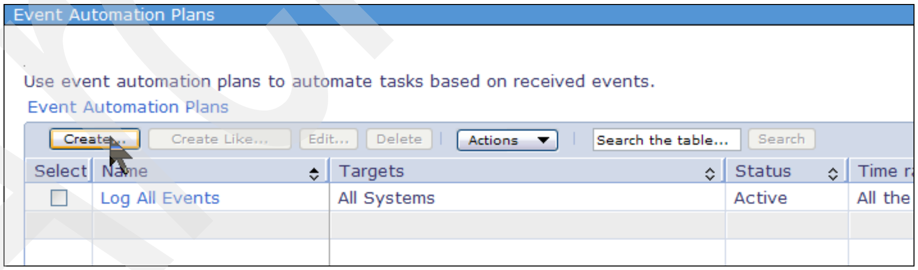


*Figure 6-82 Creating an event automation plan - 2*

4. On the Welcome panel, click **Next**. Then, enter a name and description for the automation plan, as shown in Figure 6-83. Click **Next**.

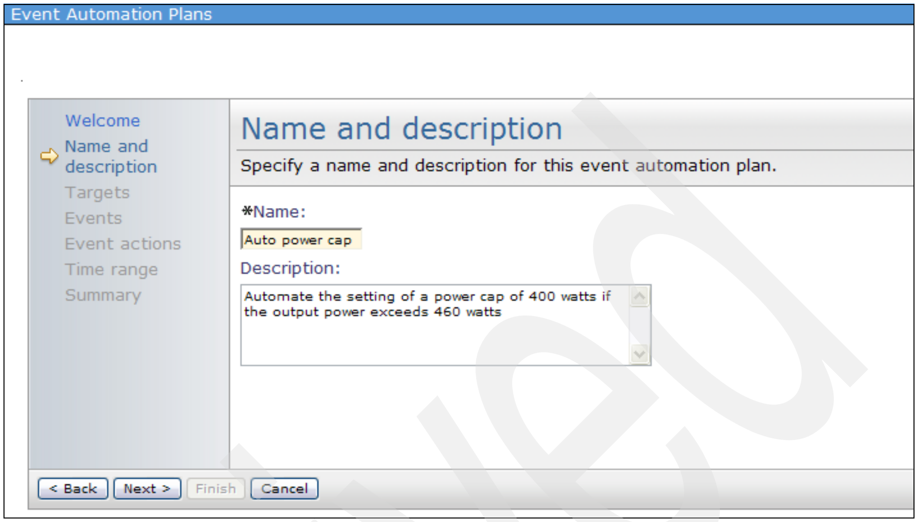


*Figure 6-83 Creating an event automation plan - 3*

5. Select the target Active Energy Manager resource or group.

   The default view of Targets is IBM Systems Director groups. You might need to drill down to find your desired group. Notice that you can create a group of Active Energy Manager resources by referring to 6.4, “Creating a group of Active Energy Manager resources” on page 195.

To select a single resource, click the **Active Energy Manager Resources** group (Figure 6-84).

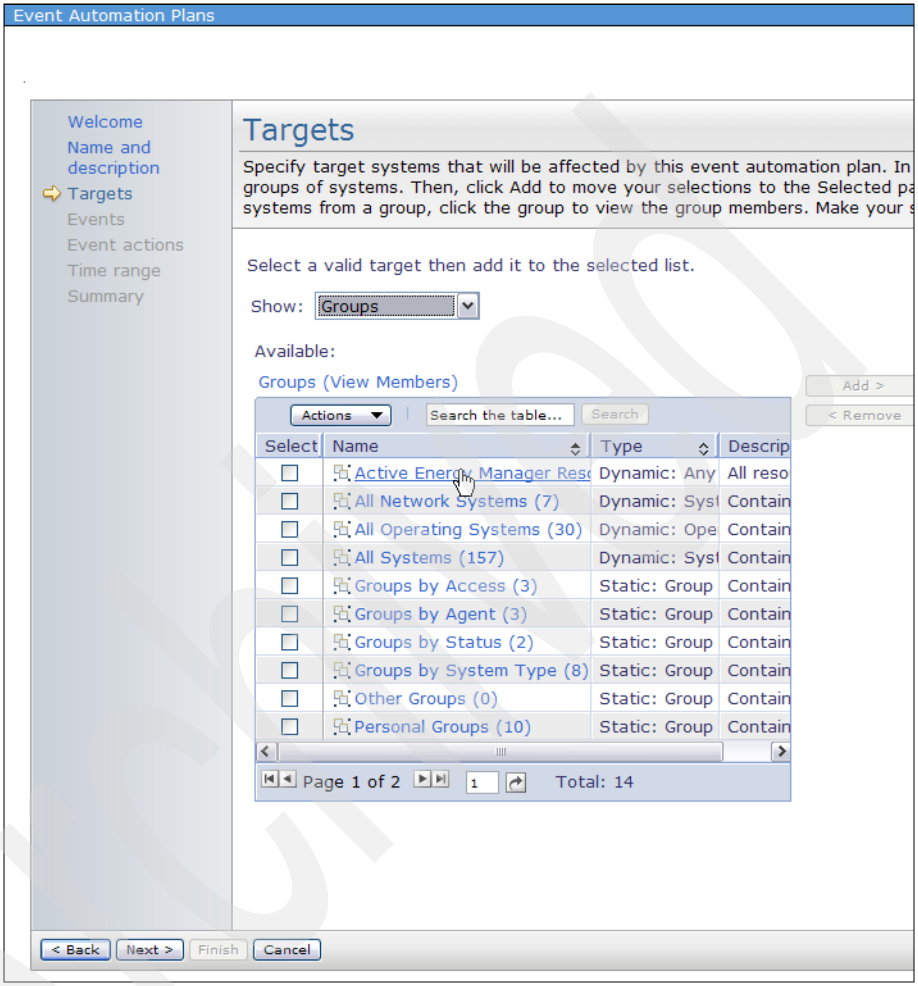


*Figure 6-84 Creating an event automation plan - 4*

6. A display similar to that shown in Figure 6-85 opens. Select the resource or resources for which you want to create the automation plan, and then click **Add** to move them to the Selected box.

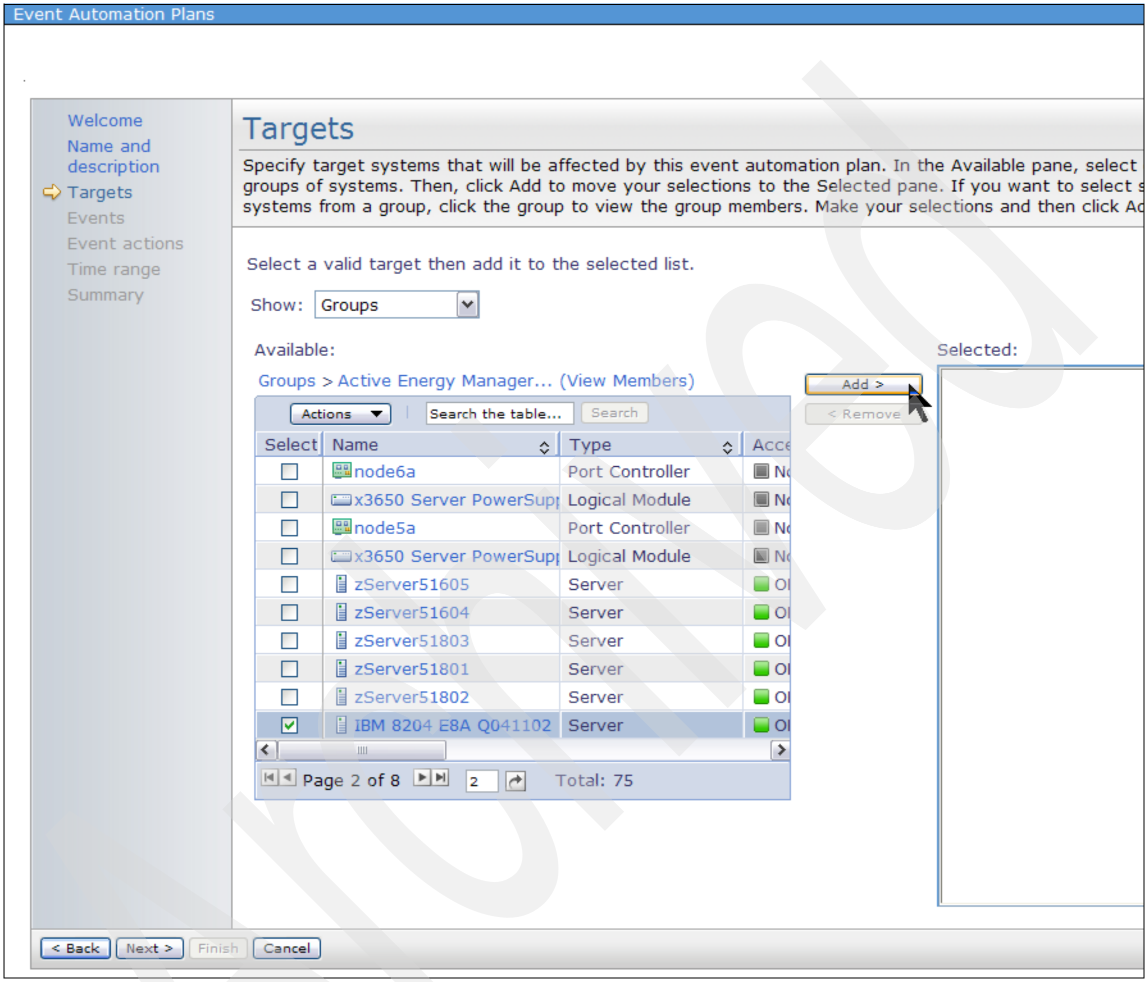


*Figure 6-85 Creating an event automation plan - 5*

7. The selected resource appears in the Selected box (Figure 6-86). Click **Next**.

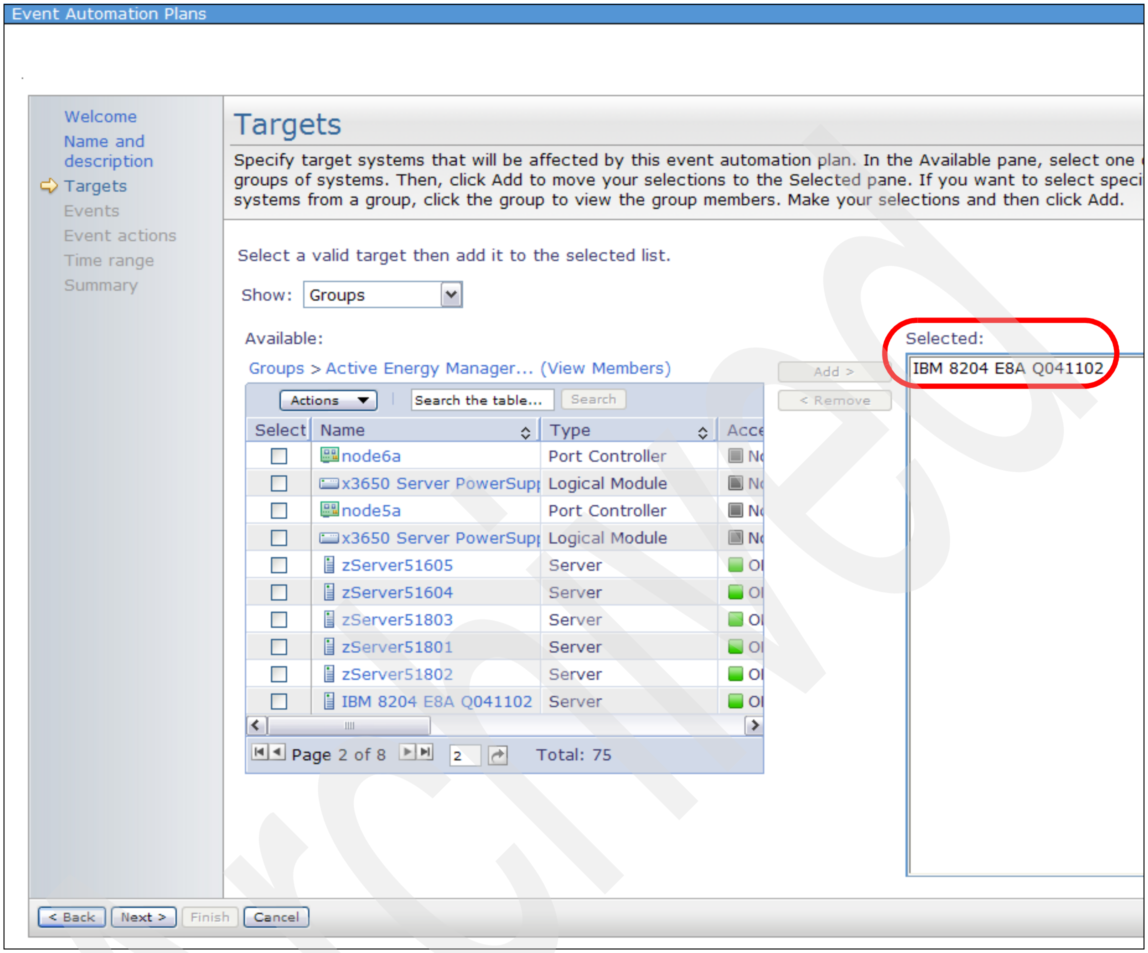


*Figure 6-86 Creating an event automation plan - 6*

8. Select the monitor that will be used to trigger the threshold (Figure 6-87). In this case, Active Energy Manager detects that we have already created a threshold for the selected monitor and inserts the threshold value. Click **Next**.

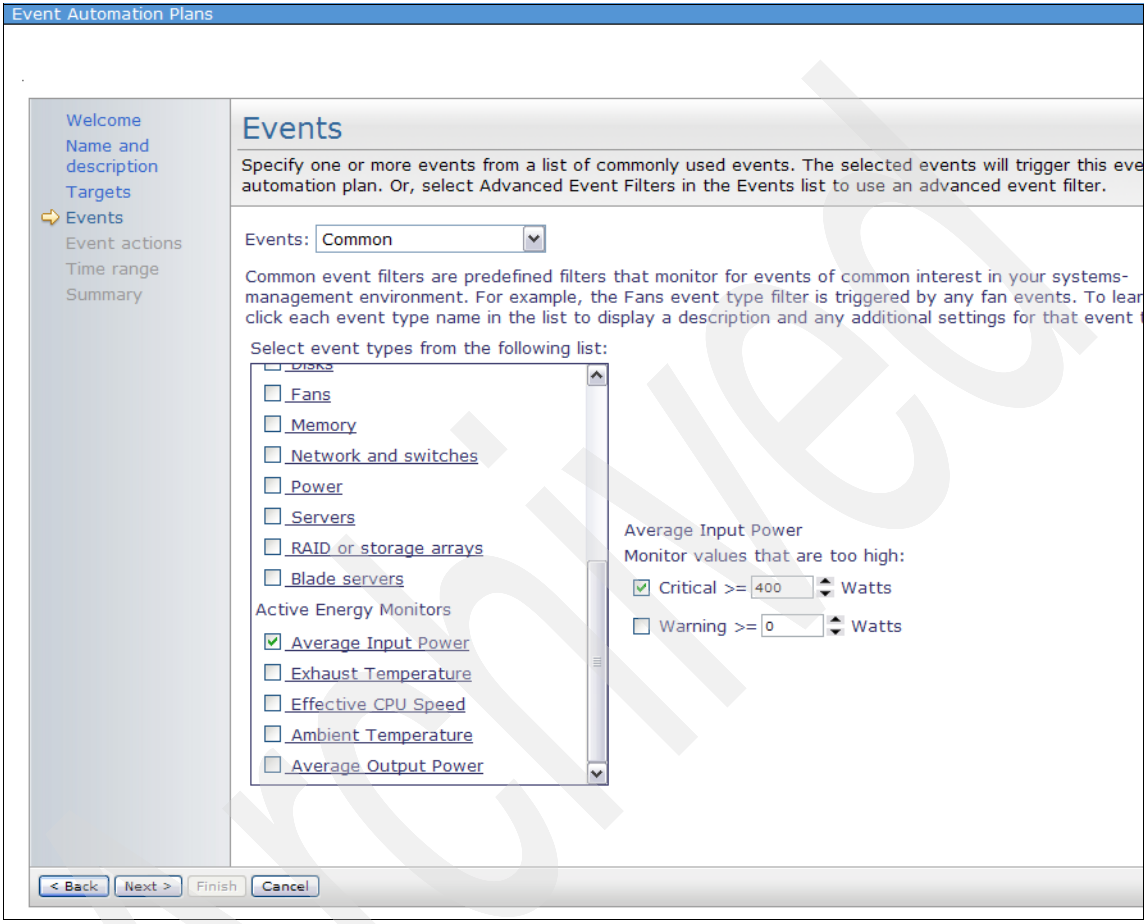


*Figure 6-87 Creating an event automation plan - 7*

9. A display similar to that shown in Figure 6-88 opens. Click **Create** to create an event action.

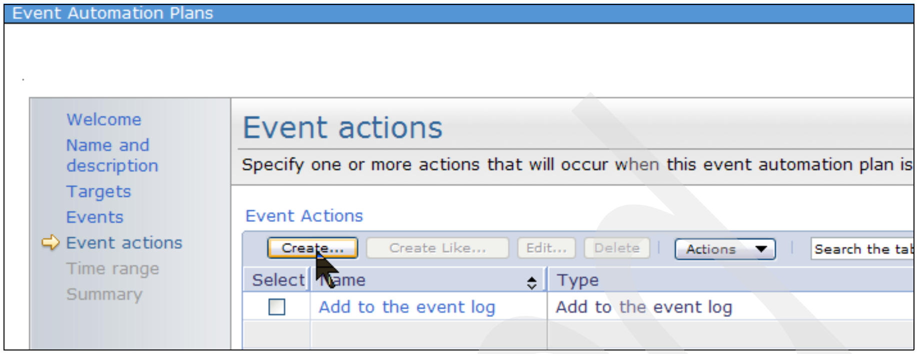


*Figure 6-88 Creating an event automation plan - 8*

10. The first page of the standard action types that IBM Systems Director supports (Figure 6-89) displays. Go to page 2, and select the "Start a task on a system that generated the event" action type. This action type is the most appropriate for our requirements. Click **OK**.


*Figure 6-89 Creating an event automation plan - 9*

11. Give the action a name, and select the policy that you created for this automation plan from the **Select a task to run** drop-down menu (Figure 6-90). In our example, we created a policy called *Static power savings policy* to turn on static power savings (refer to “Creating a system power savings policy” on page 327).

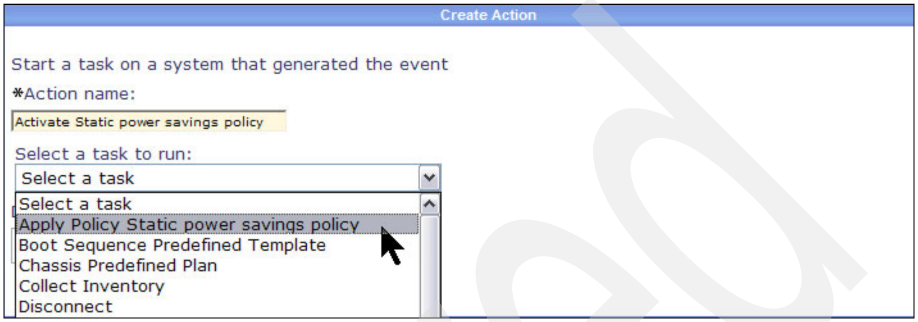


*Figure 6-90 Creating an event automation plan - 10*

12. Click **OK** to create the action (Figure 6-91).

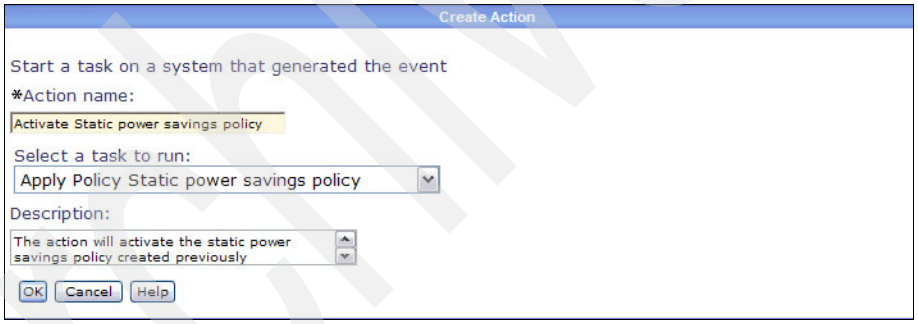


*Figure 6-91 Creating an event automation plan - 11*

13.The event action is added to the list as shown in Figure 6-92. Select the event action that you just created. Click **Next**.

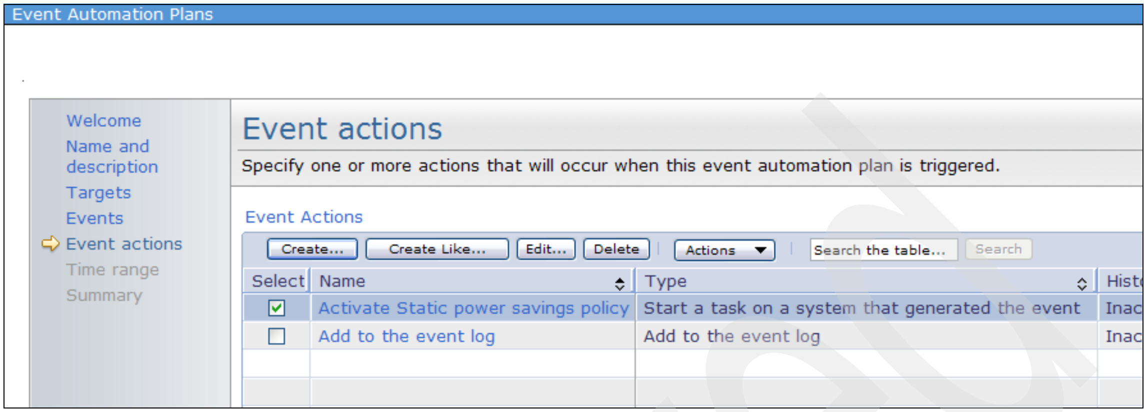


Figure 6-92 Creating an event automation plan - 12

14.Select a time range over which you want the automation plan to be active. You can specify that you want the plan to be activated all the time (Figure 6-93), or during a specific time range or ranges (Figure 6-94 on page 277).

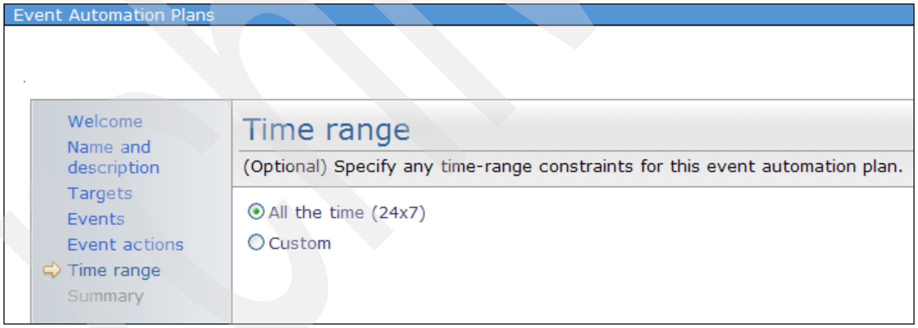


Figure 6-93 Creating an event automation plan - 13

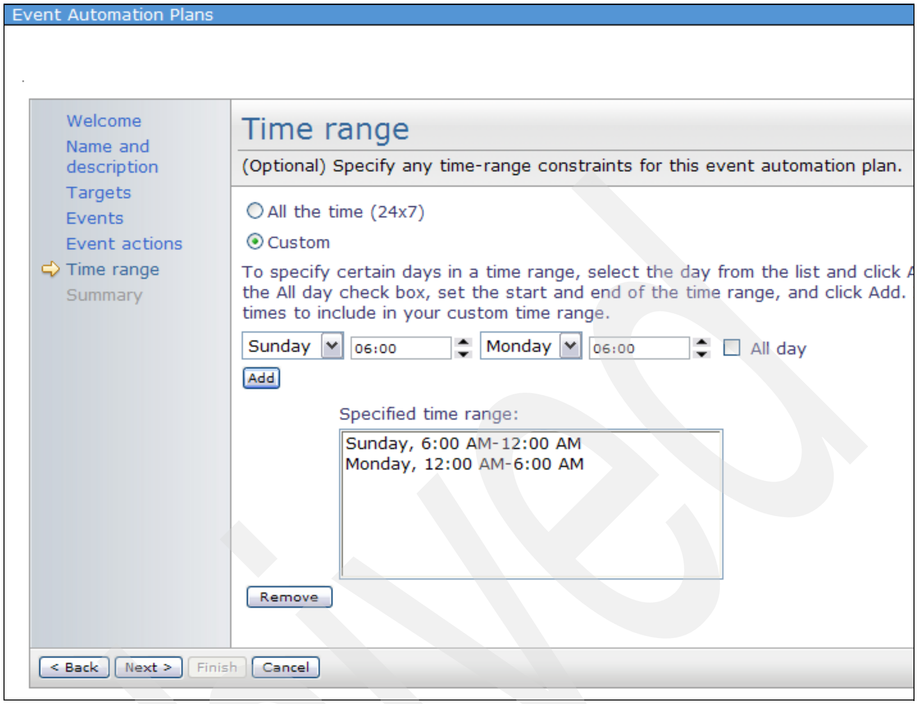


*Figure 6-94 Creating an event automation plan - 14*

15. Review your automation plan on the Summary page, and click **Finish** if you are satisfied with it. Notice that you can activate your plan when you click Finish by selecting the option **Apply this event automation plan when I click Finish** (Figure 6-95). This option is the default.

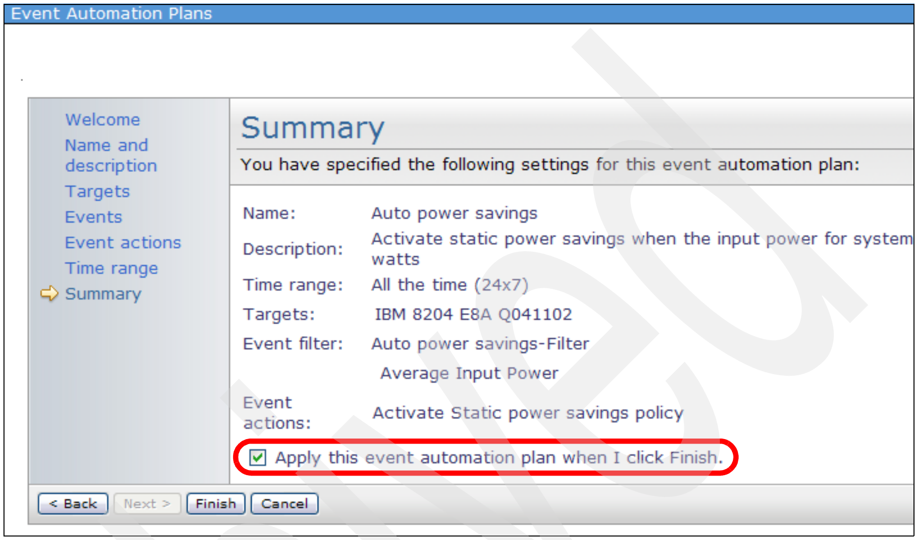


*Figure 6-95 Creating an event automation plan - 15*

16. You see a list of automation plans and their status (Figure 6-96).

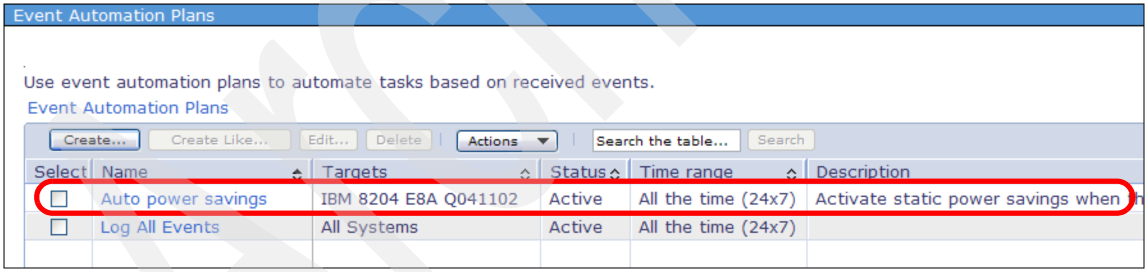


*Figure 6-96 Creating an event automation plan - 16*

17.Now view the trend data for the resource or resources for which you activated the automation plan to confirm that the actions are as you expect.

Figure 6-97 shows the trend data chart for the resource for which we have just activated the automation plan. As you can see, a critical event has been triggered because the average input power was in excess of the threshold when the automation plan was activated.

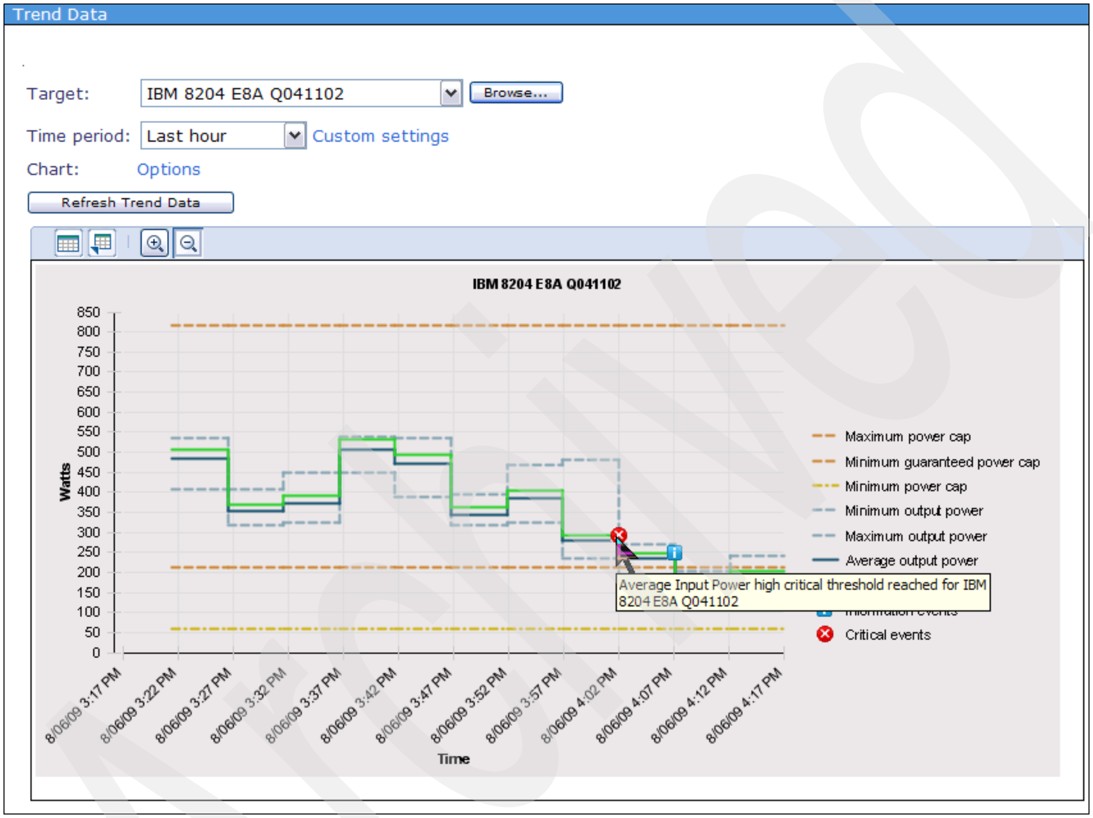


*Figure 6-97 Creating an event automation plan - 17*

The informational event icon in Figure 6-98 confirms that (static) power savings mode is turned on. Notice that the input and output power graphs have trended lower in response.

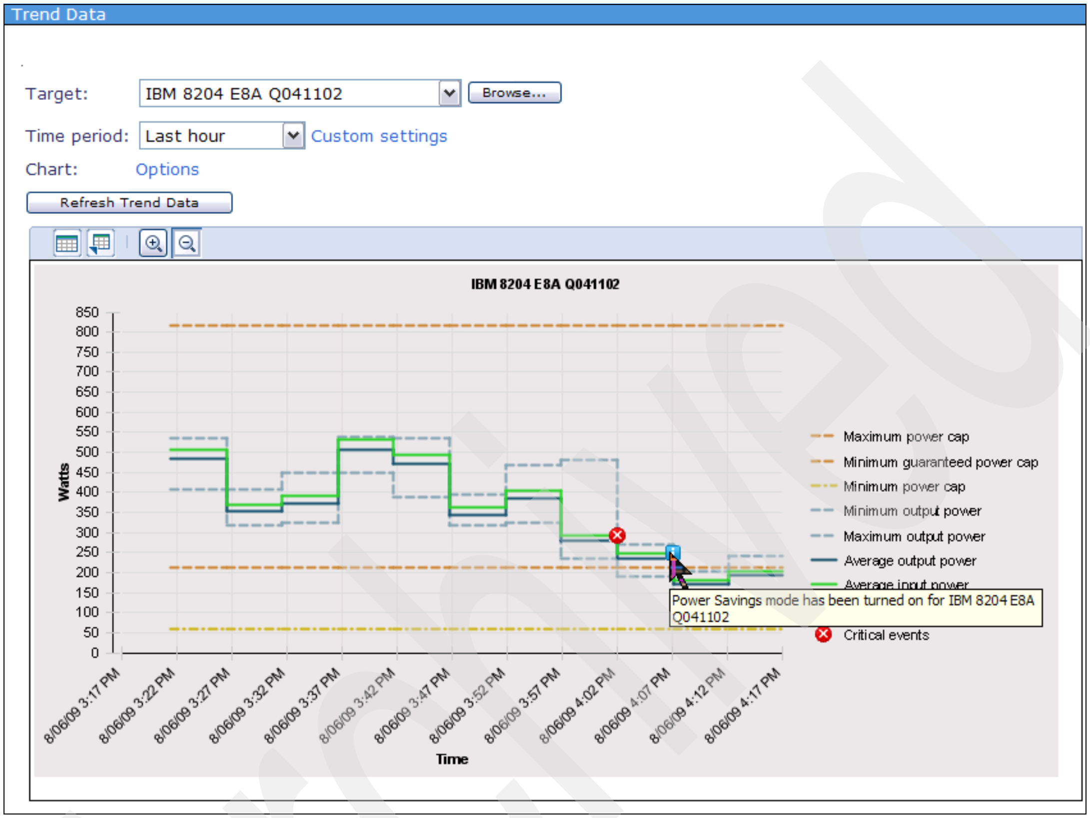


*Figure 6-98 Creating an event automation plan - 18*

This is a simple example, but it shows that the combination of setting thresholds on monitors and linking the threshold to a power management policy through an automation plan is a very powerful tool to help manage energy consumption in the data center.

### 6.13.3 Working with automation plans

If you chose not to activate your event automation plan when you finished creating the plan (clear the option in Figure 6-95 on page 278), you can enable it interactively by right-clicking the automation plan and selecting **Activate** from the context menu (Figure 6-99).

Notice also that you can create a new plan based on an existing one, edit a plan, deactivate a plan (without deleting it), delete a plan, and view the properties for a plan from this context menu.

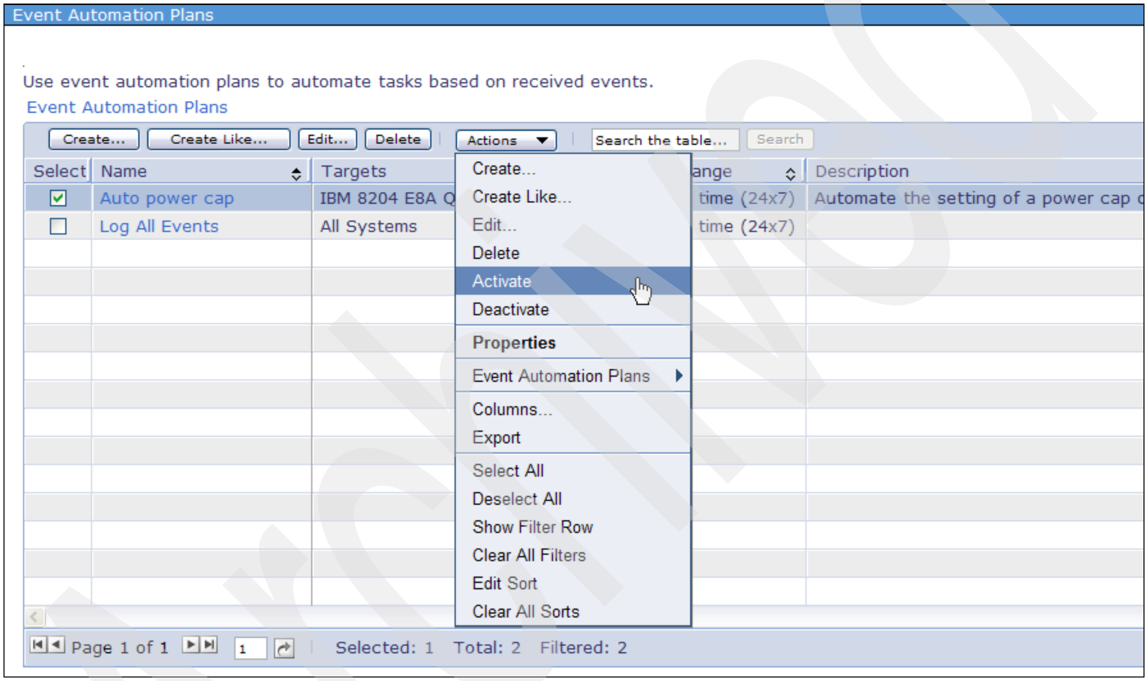


*Figure 6-99 Activating an automation plan interactively*

# 7

# Managing Active Energy Manager resources

In this chapter we describe how to manage Active Energy Manager resources. It includes the following topics:



> **Note:** These managing functions of Active Energy Manager are available for you to use free of charge for 60 days. After that, you need a license. Refer to 2.5, "Licensing" on page 64 for a description of licensing.

# 7.1 Introduction to the Active Energy Manager management tasks

This section provides a high-level introduction to the Active Energy Manager management tasks.

> **Note:** For a list of power managed systems that support Active Energy Manager managing functions, refer to 2.3.2, “Hardware requirements for managed systems” on page 56. For the most up-to-date list, go to:
>
> `http://publib.boulder.ibm.com/infocenter/director/v6r1x/topic/aem_410/frb0_r_HW_reqs_managed_systems.html`

The following list provides a brief description of each Active Energy Manager managing task and points to a detailed description for each task later in the chapter:

- Setting power savings

  *Power savings mode* throttles back the POWER6 processor by reducing the frequency (GHz) at which the processor is driven and the voltage supplied to the processor. Power savings mode has the effect of lowering the processor’s power consumption as well as its processing capacity.

  There are three different settings for power savings:

  – Static power savings
  – Dynamic power savings (favor power savings)
  – Dynamic power savings (favor performance)

  You can set power savings on Active Energy Manager resources using the following methods:

  – Using the Active Energy Manager displays, CLI or APIs
  – Using ASMI through the power managed system’s FSP
  – Using the HMC which is managing the power managed system

  We do not discuss how to set power savings using ASMI or the HMC in this book. If you want to use either of these techniques, refer to *Going Green with IBM Systems Director Active Energy Manager 3.1*, REDP-4361.

  Refer to 7.2, “Setting a power savings mode” on page 287 for a description about how to set power savings mode.

- Setting a power cap

  A *power cap* aims to set a fixed limit on the amount of power that a power managed system or group of systems can consume.

  Normally, the data center manager calculates the power requirements of the center by adding up the individual requirements of each device based on its label power. Label power, which is a rating set by an independent organization, is the maximum possible power that the device can consume. In most cases, this value is far in excess of the actual power normally consumed.

  You can set two types of power cap:

  – A *hard* power cap
  – A *soft* power cap

  If you set a hard power cap, the power consumption is guaranteed not to rise above this value. Setting a soft power cap provides some of the benefits of hard power capping, but the setting is not guaranteed.

  Refer to 7.3, “Setting a power cap” on page 306 for a description about how to set a power cap.

- Working with policies

  A *policy* is a stand-alone Active Energy Manager construct that you can use to change and enforce power savings and power capping settings. The ways that you can use policies with Active Energy Manager is greatly enhanced in Active Energy Manager 4.1.

  The main difference between settings and policies is that when you create a policy, you do not specify any target systems or groups. After you create a policy, you can then apply it to target systems and groups. A policy does not take effect until you apply it to a target. When you apply a policy, it is enforced at each poll and, therefore, cannot be overridden by changing a power savings or power capping setting from within Active Energy Manager.

  Refer to 7.5, “Working with policies” on page 322 for a description about how to use policies.

- Using the Active Energy Manager command-line interface

  There is a command-line interface (CLI) available for Active Energy Manager functions from both the Windows and Linux interfaces. The CLI provides almost identical function to the IBM Systems Director graphical user interface.

  In this chapter, we provide a summary of the commands and common uses for the CLI. For a full description of all the supported commands, refer to:

  http://publib.boulder.ibm.com/infocenter/director/v6r1x/topic/aem_410/frb0_r_ref_commands.html

  Otherwise, refer to 7.6, "Using the command-line interface" on page 362 for a summary about how to use the CLI.

## 7.1.1 Using settings and policies

You can enable power savings and power capping using settings and policies as we describe in this section.

### Using settings

You can set power savings and power capping interactively for an individual power managed system or for a group of systems. In either case, the action of setting a power savings mode or a power cap requires either an Active Energy Manager resource or group of resources.

Refer to 7.2, "Setting a power savings mode" on page 287 and 7.3, "Setting a power cap" on page 306 for more information.

### Using policies

Policies are different than settings. A *policy* is a stand-alone entity that defines an action, such as activate static power savings on a system or activate a group power cap of 4000 watts on a group of systems, but does not specify an actual system or group of systems. When created, you can *apply* a policy to an Active Energy Manager resource or group of resources to activate the setting that is defined in the policy.

You can create the following types of policies:

- System power savings
- System power capping
- Group power capping

Refer to 7.5, "Working with policies" on page 322 for more information.

# 7.2 Setting a power savings mode

We cover the following topics in this section:

- Setting static power savings
- Setting dynamic power savings (favor power savings)
- Setting dynamic power savings (favor performance)
- Setting power savings for a group of power managed systems
- Power savings examples

Power savings mode throttles back the POWER6 processor by reducing the frequency (GHz) at which the processor is driven and the voltage supplied to the processor. Power savings mode has the effect of reducing the processor's power consumption by up to 25%, but also its processing capacity.

Dynamic power savings mode is new with Active Energy Manager 4.1. It balances power usage and processor performance automatically. The net effect is to save energy while minimizing the impact on performance. For more information about dynamic power savings, refer to "Dynamic power savings" on page 18.

You can set or schedule power savings for periods of low processor utilization, which can be determined by monitoring the CPU trend data as described in 6.7, "Viewing trend data" on page 221. You can also turn on power savings mode automatically using an automation plan. For example, you might want to turn power savings mode on during periods of low processor utilization. You can manage these settings using an automation plan, as described in 6.13, "Working with automation plans" on page 265.

Active Energy Manager can detect if power savings is enabled or disabled on x86 architecture machines but cannot change the setting. On x86 architecture servers, the implementation of power savings is similar but not the same as on IBM POWER6 processor-based systems. Power savings on x86 systems is set in the BIOS and is controlled by the installed operating system.

Note the following important points about power savings:

- SMManager authority or higher is required to set a power savings mode.
- You can have both a power savings mode and a power cap set at the same time as they are independent.
- You can schedule power savings by the use of policies, as discussed in 7.5.2, “Creating a policy” on page 327.
- A power savings mode setting is persistent across system power cycles.
- In the case of a BladeCenter server, you can also set power savings directly on the blade through the management module. Refer to 7.4, “Managing power on a BladeCenter server” on page 320 for more information.
- Power savings is available on selected IBM systems as shown in Table 2-1 on page 56. For the most up-to-date list, go to:

  http://publib.boulder.ibm.com/infocenter/director/v6r1x/topic/aem_410/frb0_r_HW_reqs_managed_systems.html

> **Note:** For a technical discussion of the implementation of EnergyScale power savings on the POWER6 processor, refer to 1.3.3, “Power savings” on page 14.

Assuming that your power managed system supports power savings, you can choose from the following alternatives to set it:

- Static power savings
- Dynamic power savings (favor power savings)
- Dynamic power savings (favor performance)
- Setting power savings for a group of power managed systems

We describe how to set each of these alternatives in the following sections.

## 7.2.1 Setting static power savings

To set static power savings, follow these steps:

1. Click **Active Energy Manager Resources (View Members)** on the Active Energy Manager home page.
2. Right-click a resource, and select **Energy** → **Manage Power** → **Power Savings** (Figure 7-1).

> **Note:** If the context menu does not show the Power Savings option, then the resource does not support it.

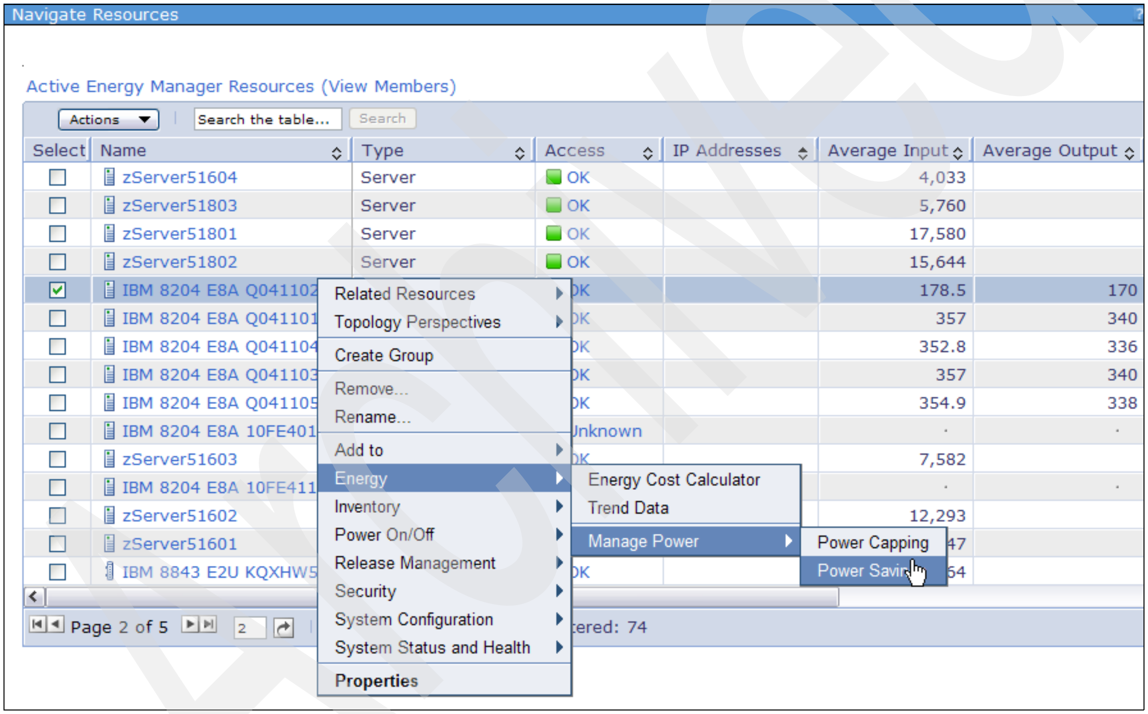


*Figure 7-1 Setting static power savings - 1*

3. The display shown in Figure 7-2 shows the available power savings modes for the selected resource. You can set static power savings or dynamic power savings. Select **Static power savings**, and then click **Save**.

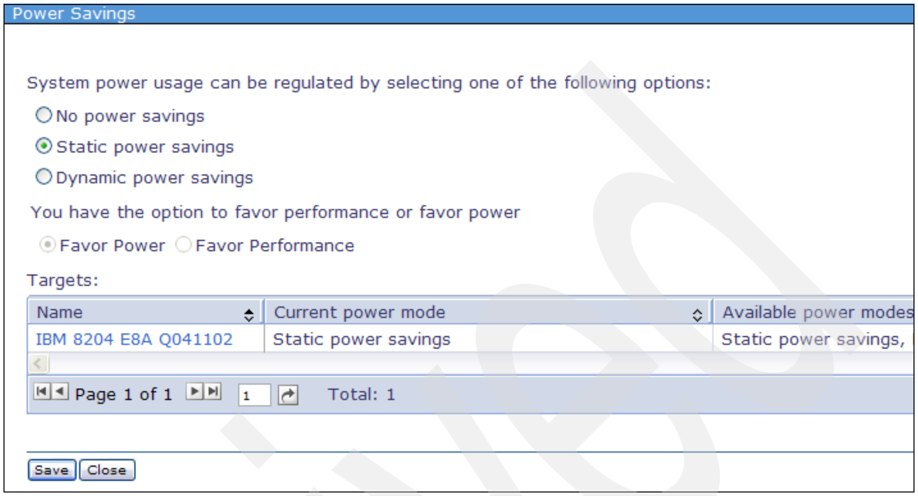


*Figure 7-2 Setting static power savings - 2*

On the Trend Data display you see an informational message indicating that power savings mode has been turned on for this system. Figure 7-3 shows an example.

It might take some time (a polling interval) for the change in power savings mode to be visible in the trend data. The informational message on the trend data chart indicates that power savings are turned on for the resource and that the average input and output power drop suddenly.

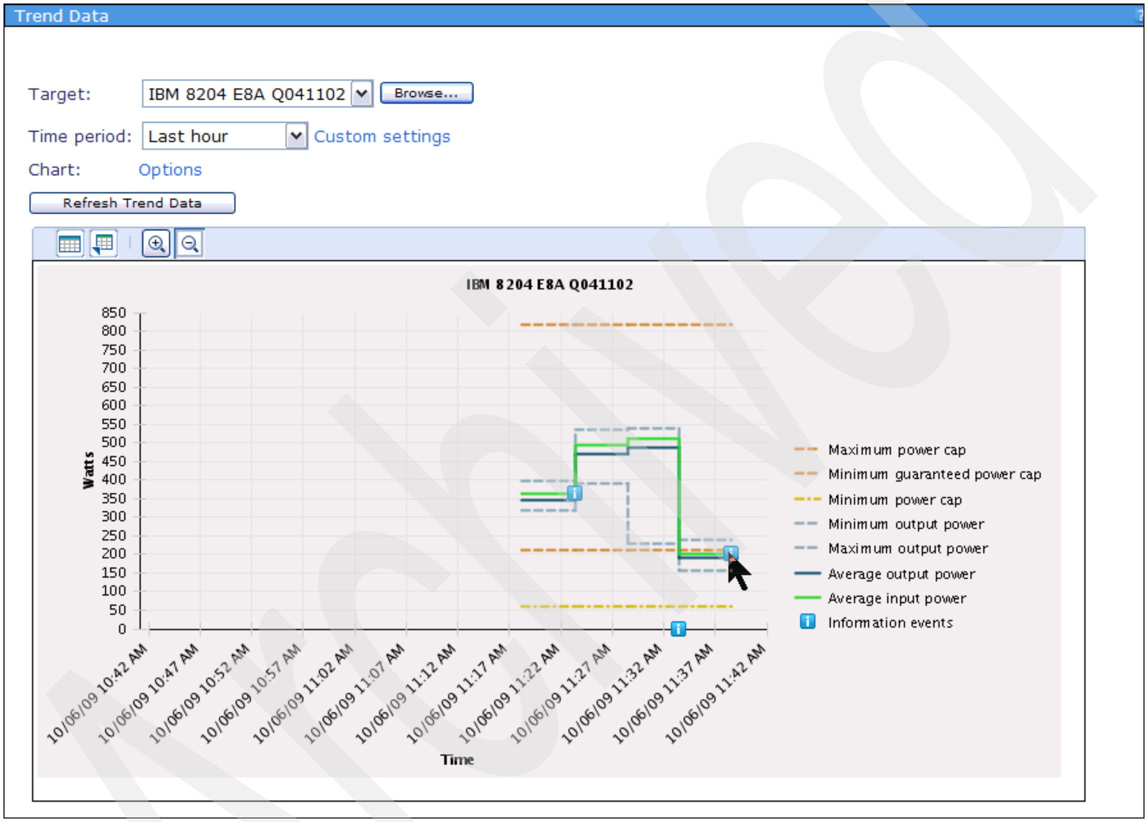


*Figure 7-3 Setting static power savings - 3*

4. Over time, you see a consistent lowering of the power requirements for the managed system when static power savings is turned on (Figure 7-4).

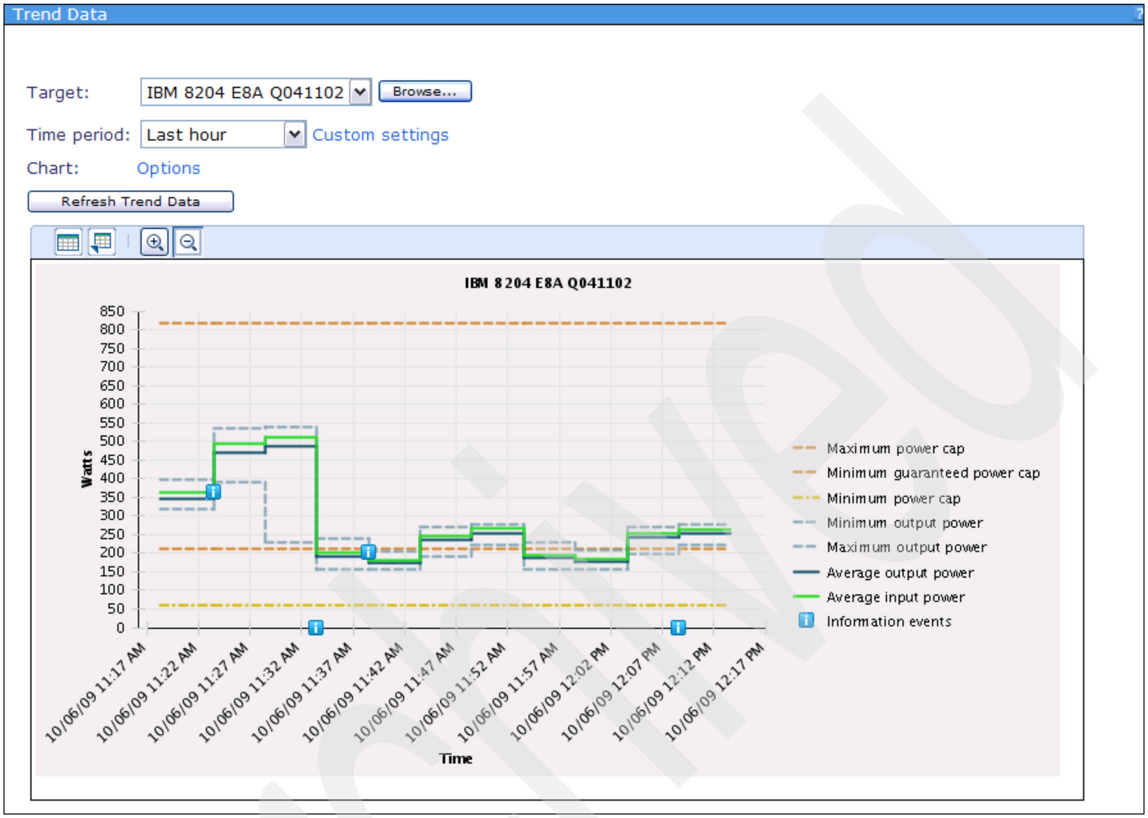


*Figure 7-4 Setting static power savings - 4*

## 7.2.2 Setting dynamic power savings (favor power savings)

To set dynamic power savings (favor power savings), follow these steps:

1. Click **Active Energy Manager Resources (View Members)** on the Active Energy Manager home page.
2. Right-click a resource and select **Energy** → **Manage Power** → **Power Savings** (Figure 7-5).

> **Note:** If the context menu does not show the Power Savings option, then the resource does not support it.

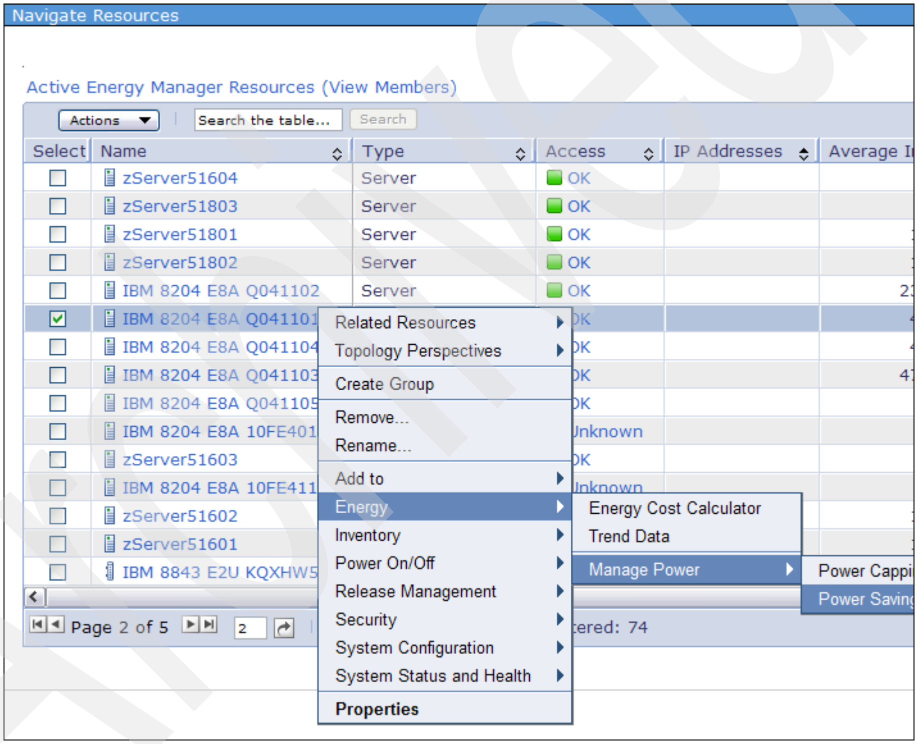


*Figure 7-5 Setting dynamic power savings (favor power savings) - 1*

3. In the display shown in Figure 7-6, the available power savings modes display for the selected resource. The dynamic power savings options are shown only if the resource supports these options. Select **Dynamic power savings**, and then **Favor Power** (savings). Click **Save**.

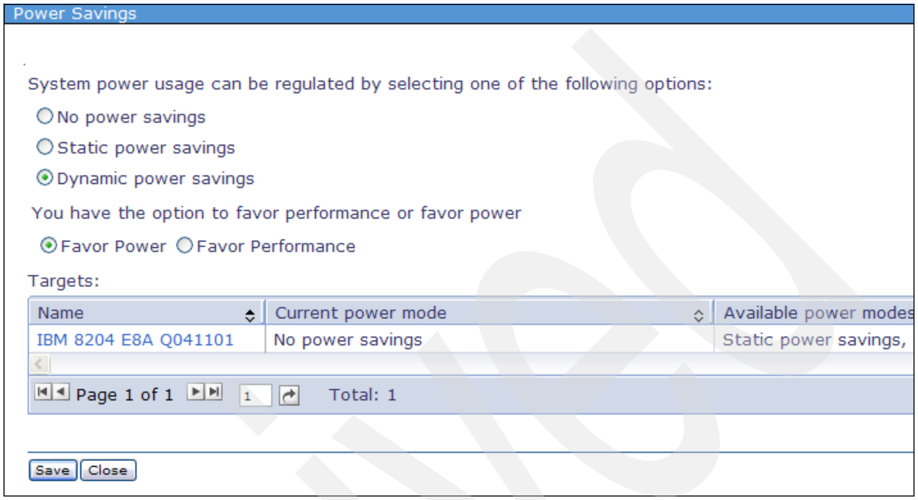


*Figure 7-6 Setting dynamic power savings (favor power savings) - 2*

The Trend Data display shows an informational message indicating that dynamic power savings mode is turned on for this system. Figure 7-7 shows an example. It might take some time (a polling interval) for the change in power savings mode to be visible in the trend data. The informational message on the trend data chart indicates that power savings is turned on for the resource and that the average input and output power drop as a result.

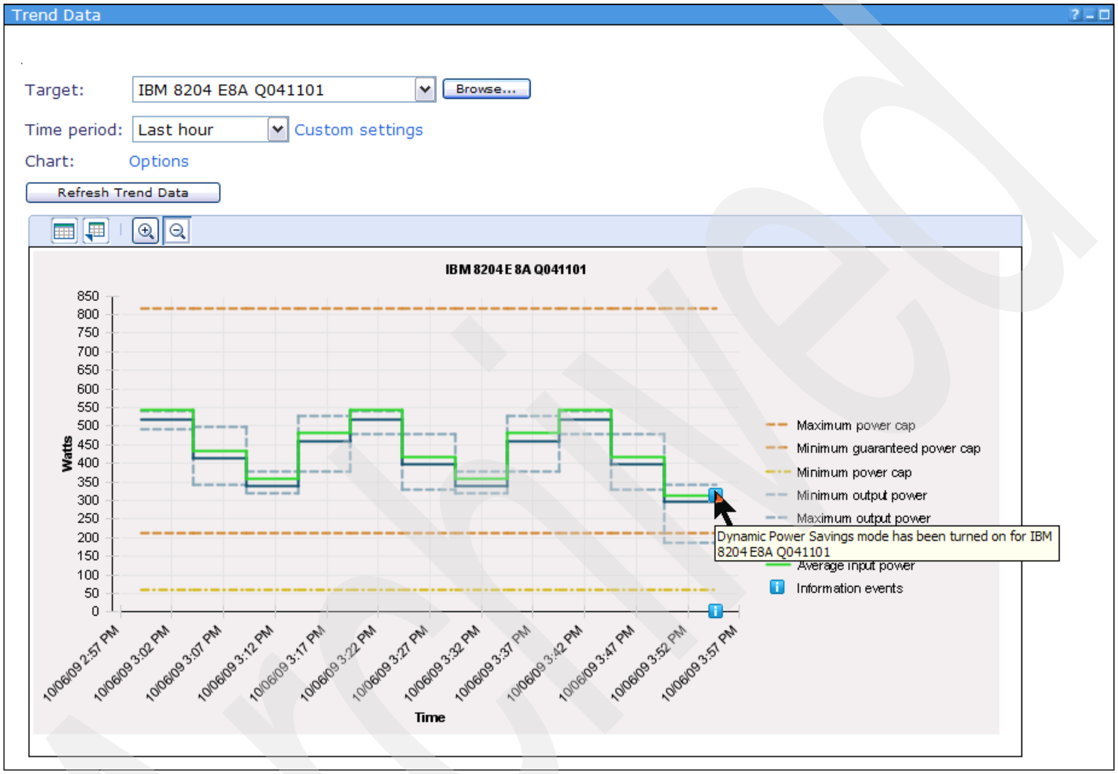


*Figure 7-7 Setting dynamic power savings (favor power savings) - 3*

Over time, you see a consistent lowering of the power requirements for the managed system when static power savings (favor power savings) is turned on (Figure 7-8).

Notice that the behavior is the same as for static power savings when the workload is light. You only notice a difference when the system is busy. In this case, you might see the power trending lines rise to reflect the increased requirement for processing capacity. This power trending occurs because dynamic power savings provides additional processing capacity over and above static power savings when the workload requires it but only up to 95% of the processor's nominal capacity. Compare Figure 7-8 with Figure 7-4 on page 292 to see the difference.

For a more detailed explanation about how dynamic power savings works, refer to "Dynamic power savings" on page 18.

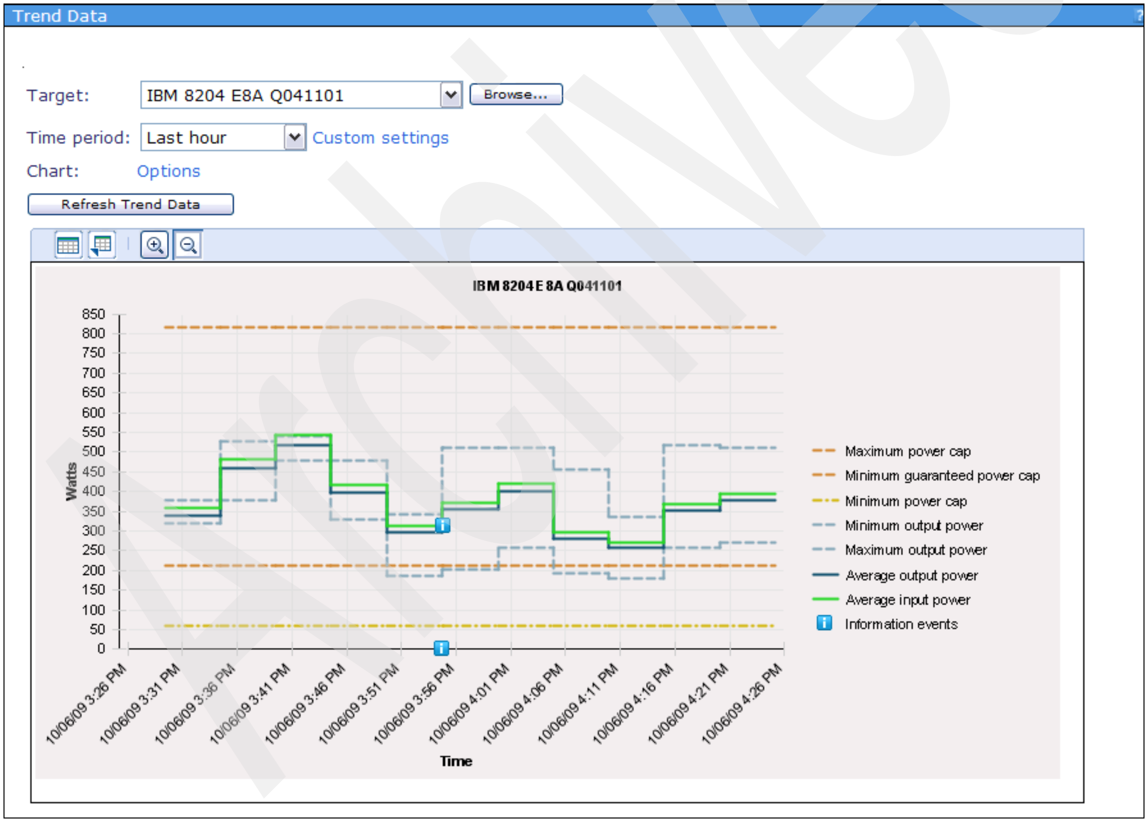


Figure 7-8 Setting dynamic power savings (favor power savings) - 4

## 7.2.3 Setting dynamic power savings (favor performance)

To set dynamic power savings (favor performance), follow these steps:

1. Click **Active Energy Manager Resources (View Members)** on the Active Energy Manager home page.
2. Right-click a resource, and select **Energy** → **Manage Power** → **Power Savings** (Figure 7-9).

> **Note:** If the context menu does not show the Power Savings option, then the resource does not support it.

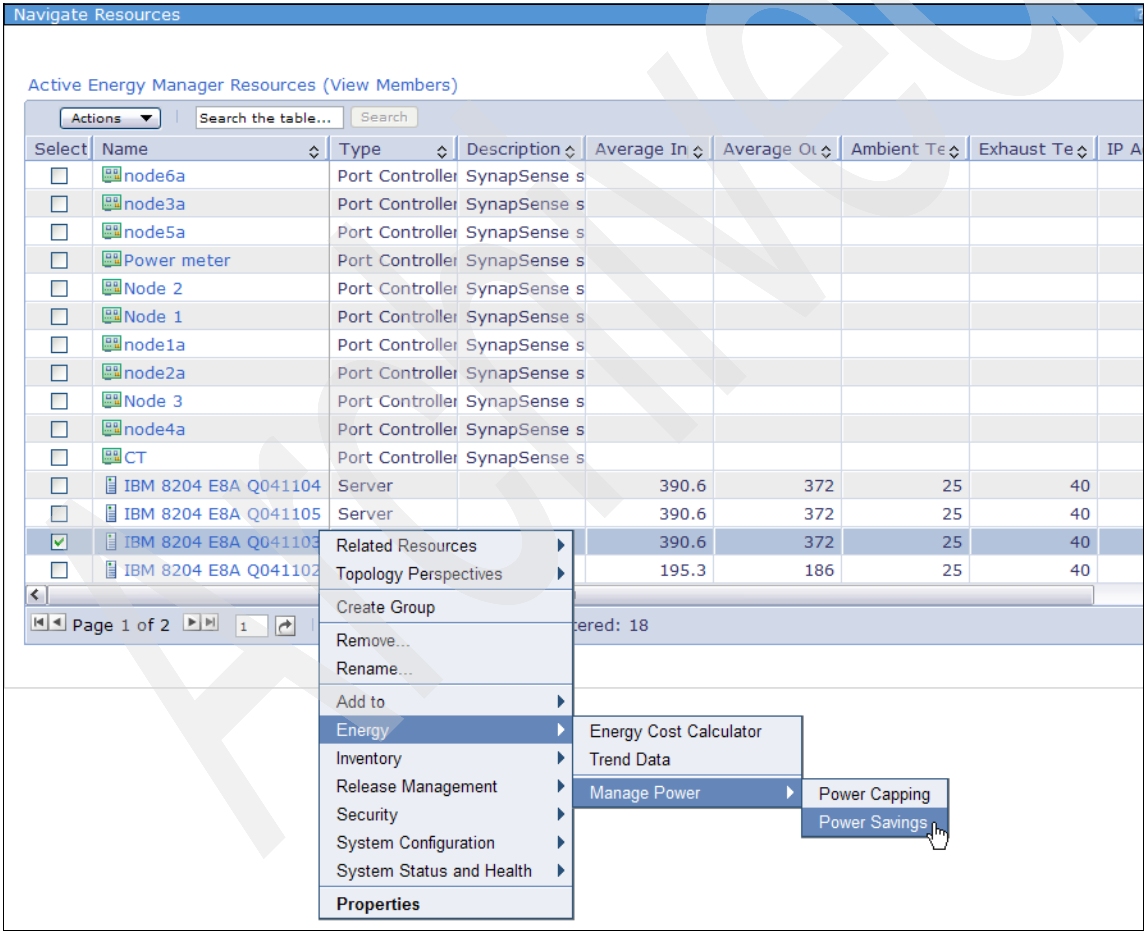


*Figure 7-9 Setting dynamic power savings (favor performance) - 1*

3. In the display shown in Figure 7-10, the available power savings modes display for the selected resource. The dynamic power savings options are only shown if the resource supports these options. Select **Dynamic power savings**, and then **Favor Performance**. Click **Save**.

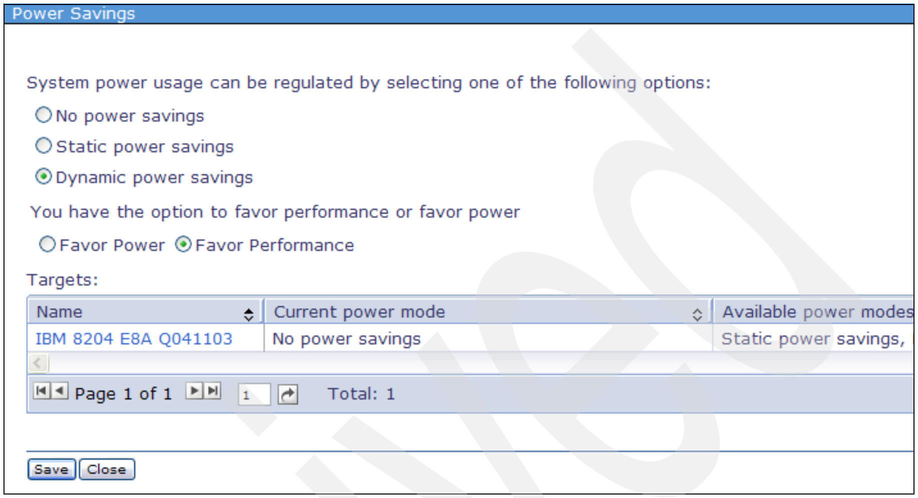


*Figure 7-10 Setting dynamic power savings (favor performance) - 2*

The Trend Data display shows an informational message indicating that dynamic power savings mode is turned on for this system. Figure 7-11 shows an example. It might take some time (a polling interval) for the change in power savings mode to be visible in the trend data. The informational message on the trend data chart indicates that power savings is turned on for the resource and that the average input and output power drop as a result.

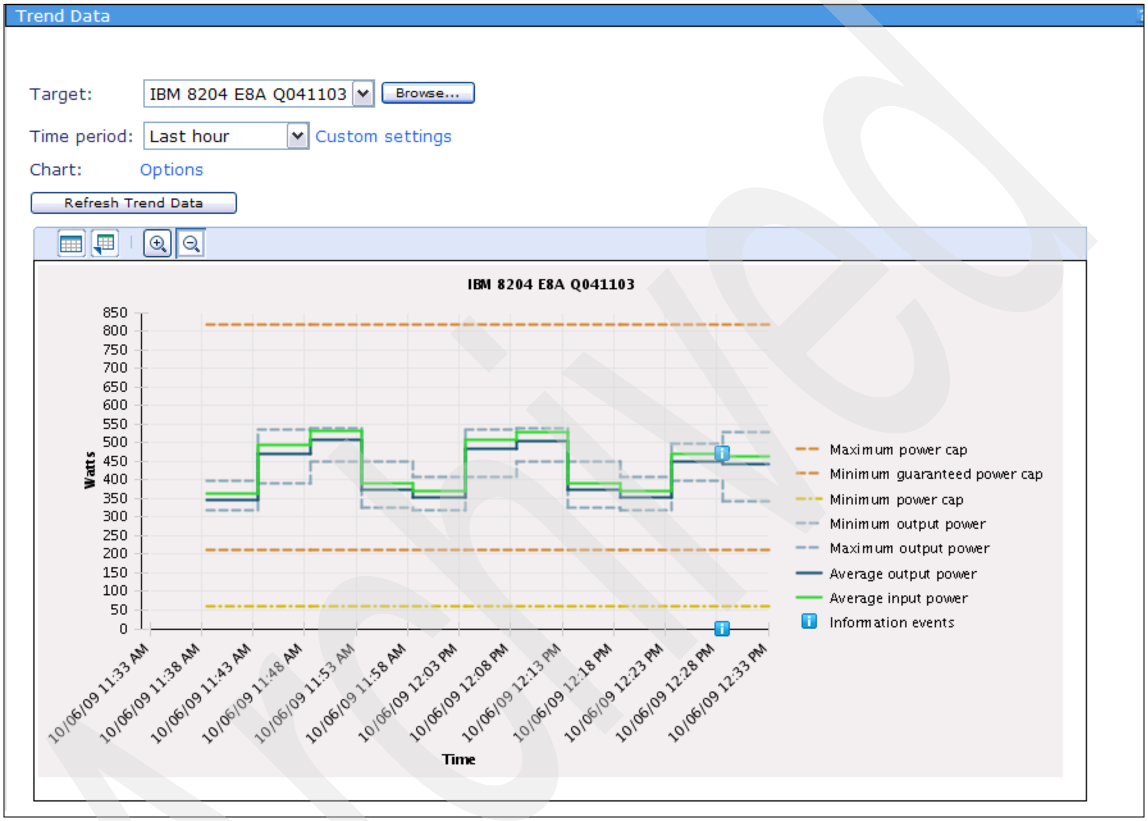


*Figure 7-11 Setting dynamic power savings (favor performance) - 3*

Over time, you see a consistent lowering of the power requirements for the managed system when static power savings (favor performance) is turned on (Figure 7-12 on page 300). Notice that the behavior is the same as for static power savings when the workload is light. You notice a difference only when the system is busy. In this case, you might see the power trending lines rise to reflect the increased requirement for processing capacity. This power trending occurs because dynamic power savings provides additional processing capacity over and above static power savings when the workload requires it.

If the system is really busy, you might even see the power graphs rise *above* previous levels (before you turned on power savings) if the Power system model supports overclocking. Compare Figure 7-12 with Figure 7-4 on page 292 and Figure 7-8 on page 296. Behavior is very similar to that shown for dynamic power savings (favor power savings), because the power curves shown in these examples are taken from a Power system that does not support overclocking.

For a more detailed explanation about how dynamic power savings works, refer to “Dynamic power savings” on page 18.

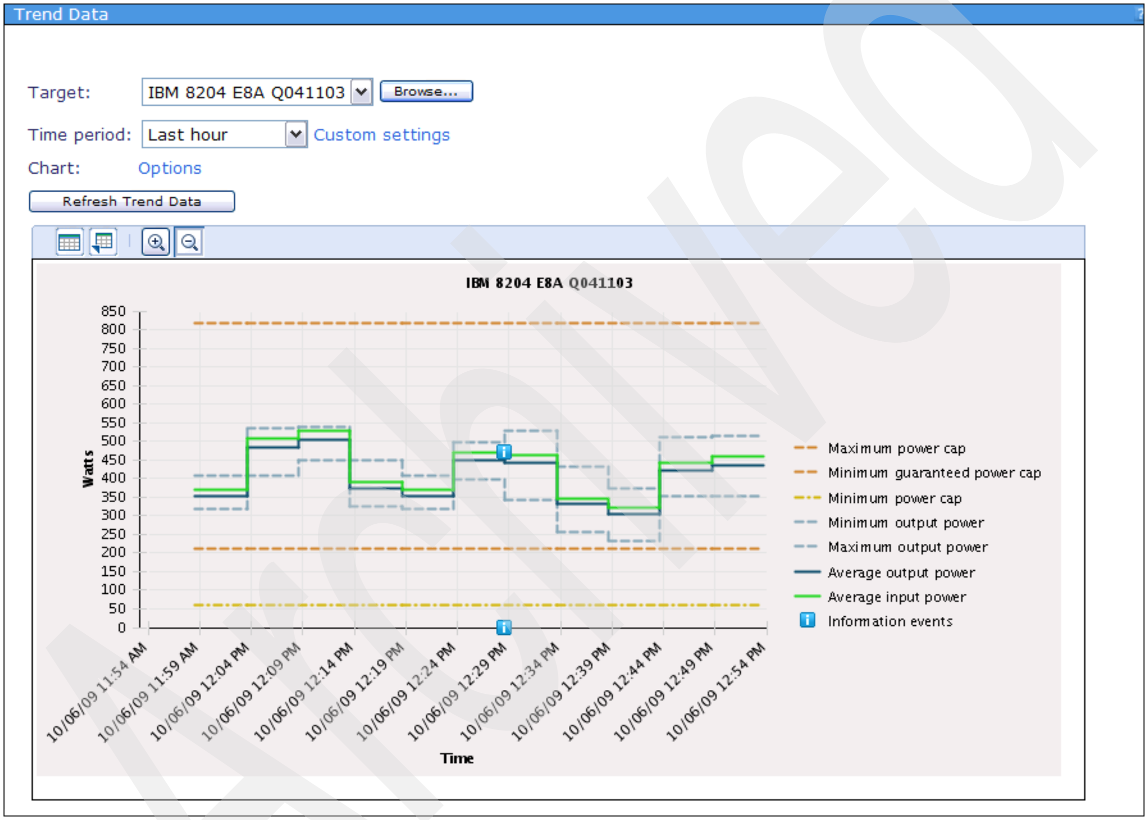


*Figure 7-12 Setting dynamic power savings (favor performance) - 4*

## 7.2.4 Setting power savings for a group of power managed systems

When you set power savings for a group of power managed systems, the power savings setting that you choose is activated on each member of the group.

> **Note:** To activate power savings mode on a group of power managed systems, all members of the group must support it.

To set power savings mode for a group of power managed systems, follow these steps:

1. Close any open Navigate Resources tabs, and then click **Navigate Resources** in the left hand navigation pane.
2. Right-click a group of Active Energy Manager resources, and select **Energy** → **Manage Power** → **Power Savings** (Figure 7-13).

   > **Note:** If the context menu does not show the Power Savings option, then not all members of the group support it.

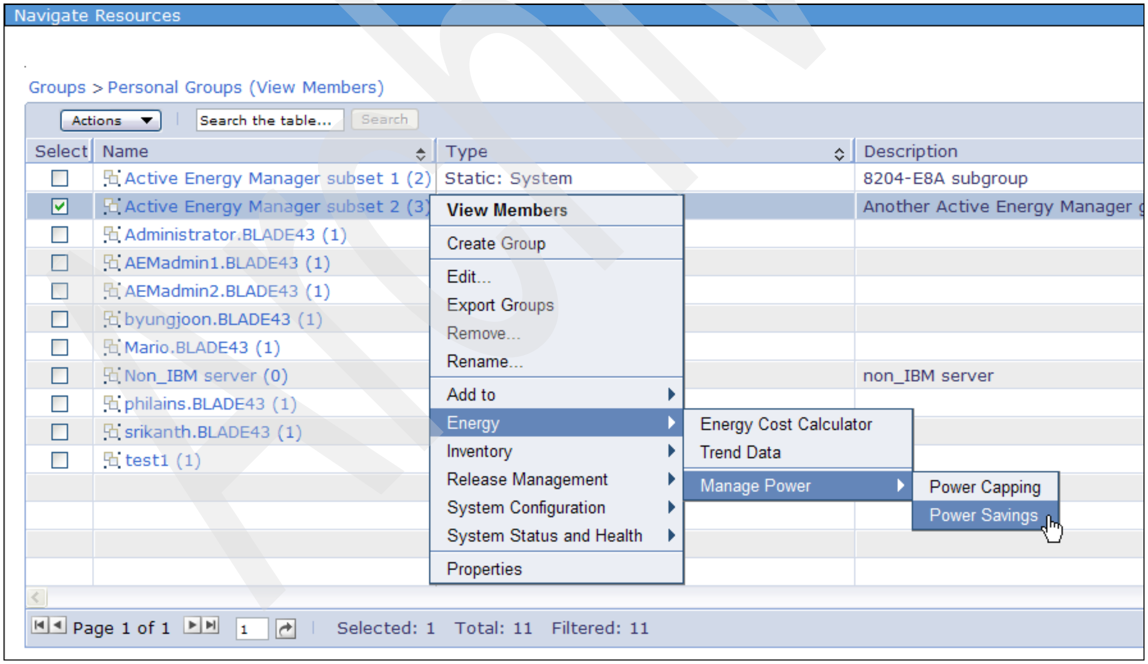


*Figure 7-13 Setting power savings for a group of power managed systems - 1*

The display shown in Figure 7-14 shows all the power savings modes supported by members of the group. If some resources do not support the selected power savings mode, no change is applied to those resources. Set the power savings mode of your choice, and then click **Save**.

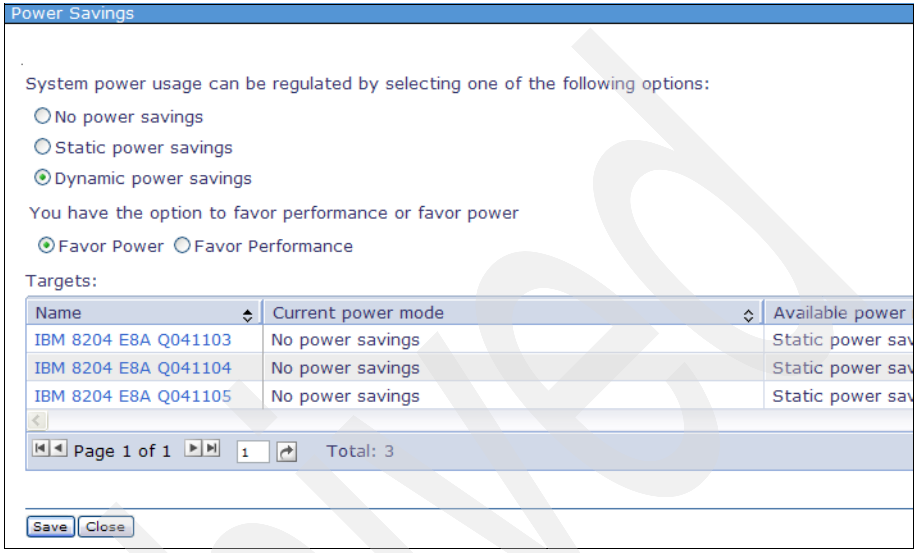


*Figure 7-14 Setting power savings for a group of power managed systems - 2*

3. After a few minutes, the status changes in the “Current power mode” column (Figure 7-15) to the setting that you chose. If you check the trend data for each of the power managed systems in the group, you see that power savings are activated on each one.

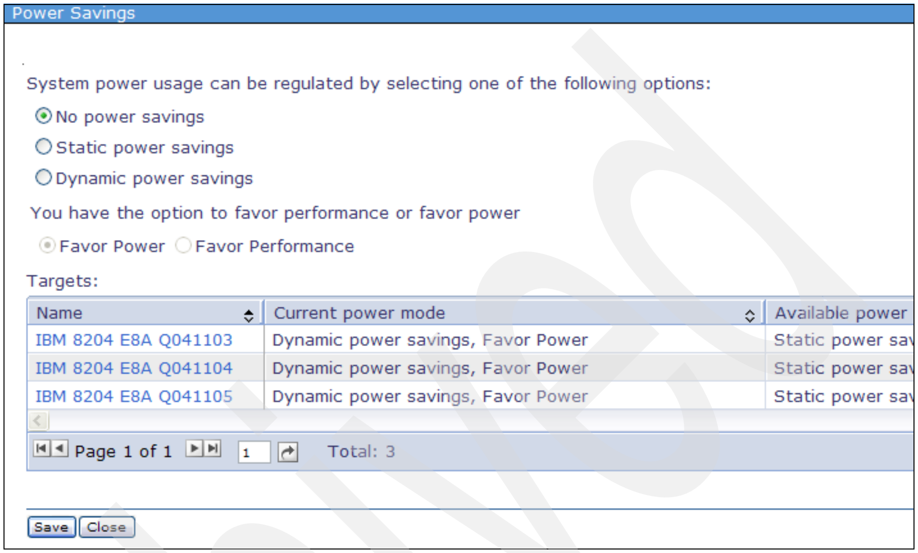


*Figure 7-15 Setting power savings for a group of power managed systems - 3*

## 7.2.5 Power savings examples

This section provides two examples of how to use power savings.

## Power savings example 1

This example sets static power savings on a system and observes the results. Figure 7-16 shows that power capping and power savings are initially disabled on the system.

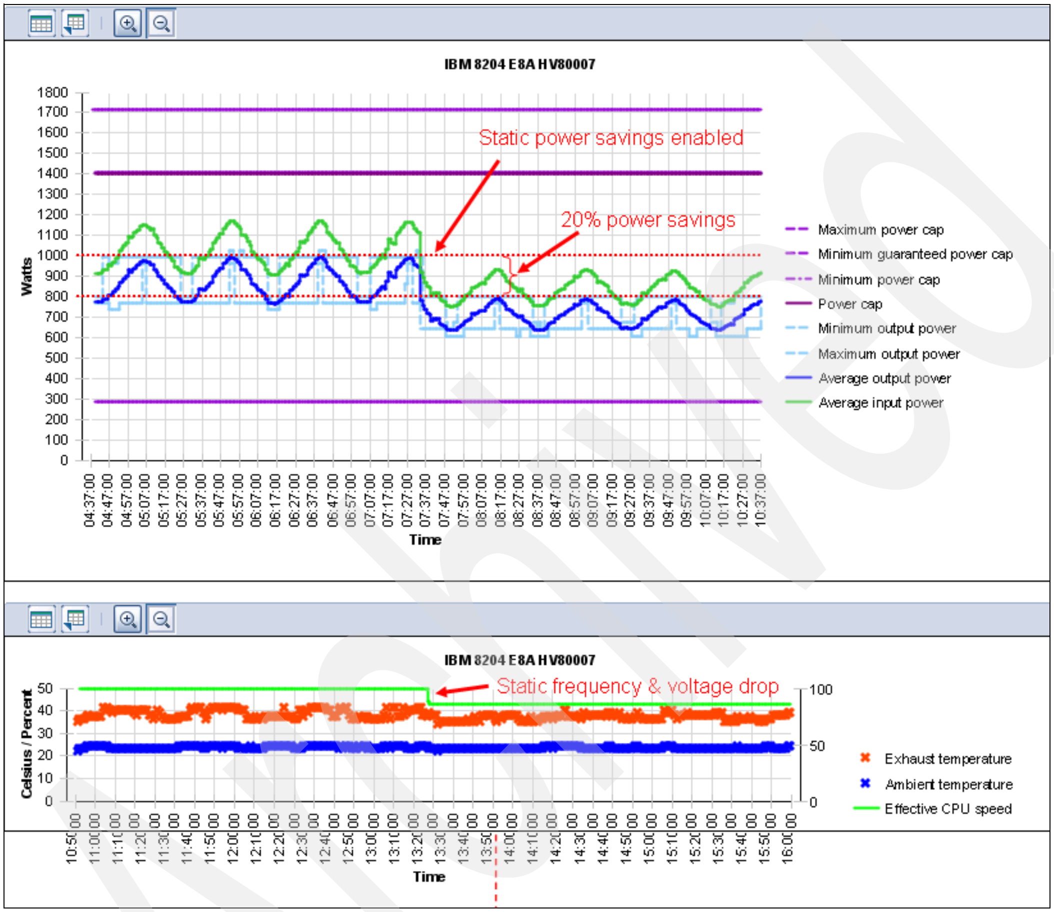


*Figure 7-16 Static power savings for example 1*

The workload shows a cyclical behavior, with output power (DC) varying between approximately 750 watts when idle and 1,000 watts under full load. Shortly after 07:30:00, an administrator enables static power savings mode. The system firmware immediately lowers the processor frequency and voltage to their static low power values, resulting in immediate and significant power savings over the (default) high performance mode. The workload continues to run, however, until peak performance is reduced because the processor is now running at a lower frequency. As shown in Figure 7-16, the static power savings setting results in approximately 20% power savings.

### Power savings example 2

This example sets dynamic power savings on a system and observes the results.

Dynamic power savings delivers nearly full workload performance while still saving power. This example disables power savings and power capping, as shown in Figure 7-17, and the target system is idle.

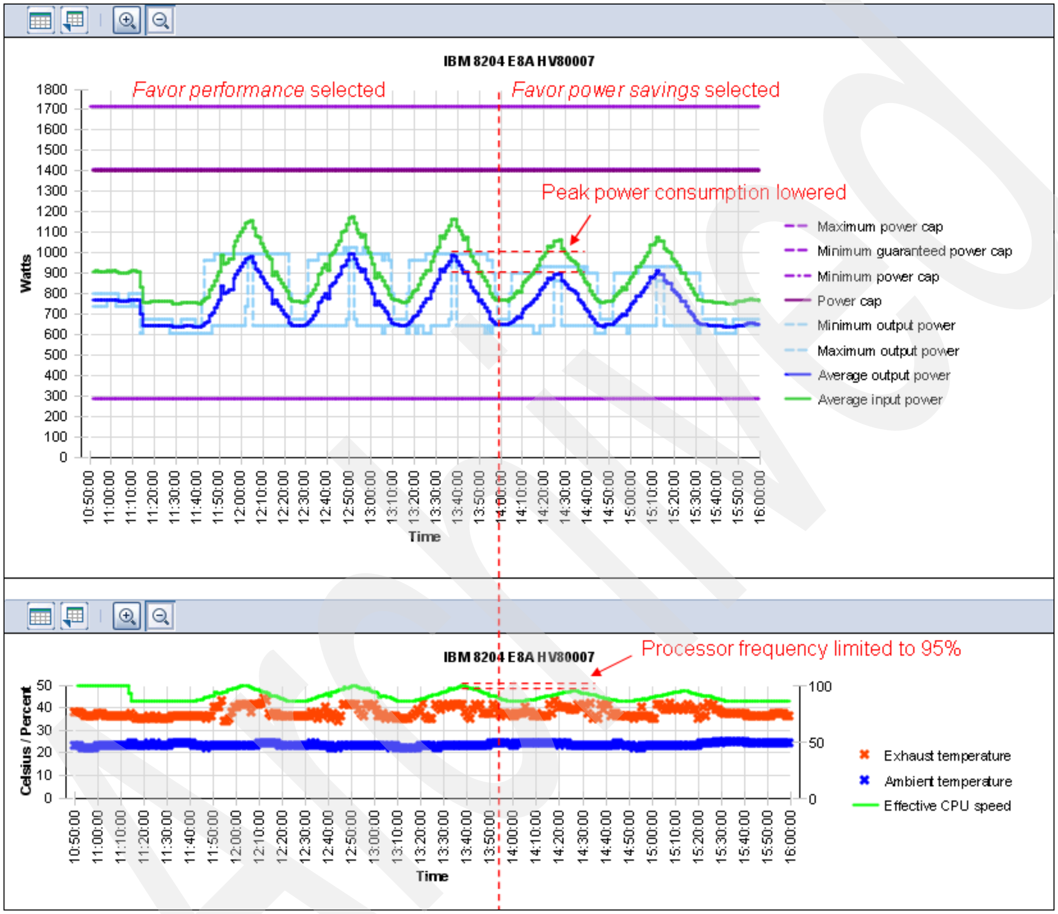


*Figure 7-17 Dynamic power savings for example 2*

At 11:15:00, dynamic power savings is enabled with the *favor performance* option selected. The system firmware reads the idle utilization of the POWER6 processors, and lowers the frequency to the minimum supported frequency for the system. As the workload increases, the system firmware automatically increases the processor frequency and voltage to supply adequate processing power, while minimizing any performance impact. On supported systems the peak frequency can be greater than 100% when the *favor performance* option is selected. At 14:00:00, the dynamic power savings setting is modified to *favor power savings*. This results in additional power savings as the system processors are limited to 95% of their nominal frequency.

## 7.3 Setting a power cap

We cover the following topics in this section:

- Setting a power cap for a single power managed system
- Setting a power cap for a group of systems
- Power capping examples

Hard power capping sets a fixed limit (called a *cap*) on the amount of power that a power managed system can consume. If the cap is reached, the service processor firmware turns down the processor's voltage and clocking rate to reduce power consumption to keep the power used by the managed system under the cap.

Normally, the data center manager calculates the power requirements of the center by adding up the individual requirements of each device based on the device's label power. Label power, which is a rating set by an independent organization, is the maximum possible power that the device can consume. In most cases, this value is far in excess of the actual power consumed. The idea of limiting the power consumption of a managed system by setting a power cap enables you to accurately calculate the actual power requirements of the data center by setting a cap on the maximum power consumption of the systems in the data center. You can then use this information to plan your power supply infrastructure and budget for power that will be consumed.

In many cases the power consumption of the managed system will never reach the power cap. If, for example, the data center suffers the failure of a cooling unit, the fans in the managed systems will need to work harder to cool the systems, thus increasing the power consumption of the system as a whole. In this case, power capping reduces the power consumption of the processor, thereby reducing the heat produced and the fan speed. Therefore, power consumption is reduced at the processor level, fan level and cooling equipment level.

**Note:** Hard power cap settings are guaranteed not to be exceeded.

Starting with Active Energy Manager 4.1, the range of power values in which power capping can be set is divided into a hard power cap range and a soft power cap range. Soft power capping is only supported on IBM POWER6 systems. We explain these terms in the following sections.

**Note:** Soft power cap settings are *not* guaranteed and can be exceeded under certain circumstances.

Note the following important points about power capping:

- Power capping works with output power (DC), not input power (AC).
- SMManager authority or higher is required to set a power cap.
- You can have both a power savings mode and a power cap set at the same time as they are independent.
- A power cap setting is persistent across restarts of a power managed system, but in the case of a blade server, the setting is lost if the blade server is removed from the chassis.
- The Active Energy Manager 3.1 option to set a power cap based on measured trend data is not available with version 4.1.1.
- On a BladeCenter, you can do power capping at the slot *or* blade level but not both.
- Power capping can be scheduled by the use of policies as discussed in 7.5.2, "Creating a policy" on page 327.
- In the case of a BladeCenter server, you can also set a power cap directly on the blade through the management module. Refer to 7.4, "Managing power on a BladeCenter server" on page 320 for a description of how to do this.
- Power capping is available on selected IBM systems as shown in Table 2-1 on page 56. For the most up-to-date list, go to:

  http://publib.boulder.ibm.com/infocenter/director/v6r1x/topic/aem_410/frb0_r_HW_reqs_managed_systems.html

**Note:** For a technical discussion of the implementation of EnergyScale power capping on the POWER6 processor, refer to 1.3.4, "Power capping" on page 23.

Using Active Energy Manager, you can set a power cap for a single power managed system or a group of power managed systems. We cover both cases in the following sections.

## 7.3.1 Setting a power cap for a single power managed system

To set a power cap for a single power managed system using Active Energy Manager, follow these steps:

1. Click **Active Energy Manager Resources (View Members)** on the Active Energy Manager home page.
2. Right-click a resource and select **Energy** → **Manage Power** → **Power Capping** (Figure 7-18).

> **Note:** If the context menu does not show the Power Capping option, then the resource does not support it.

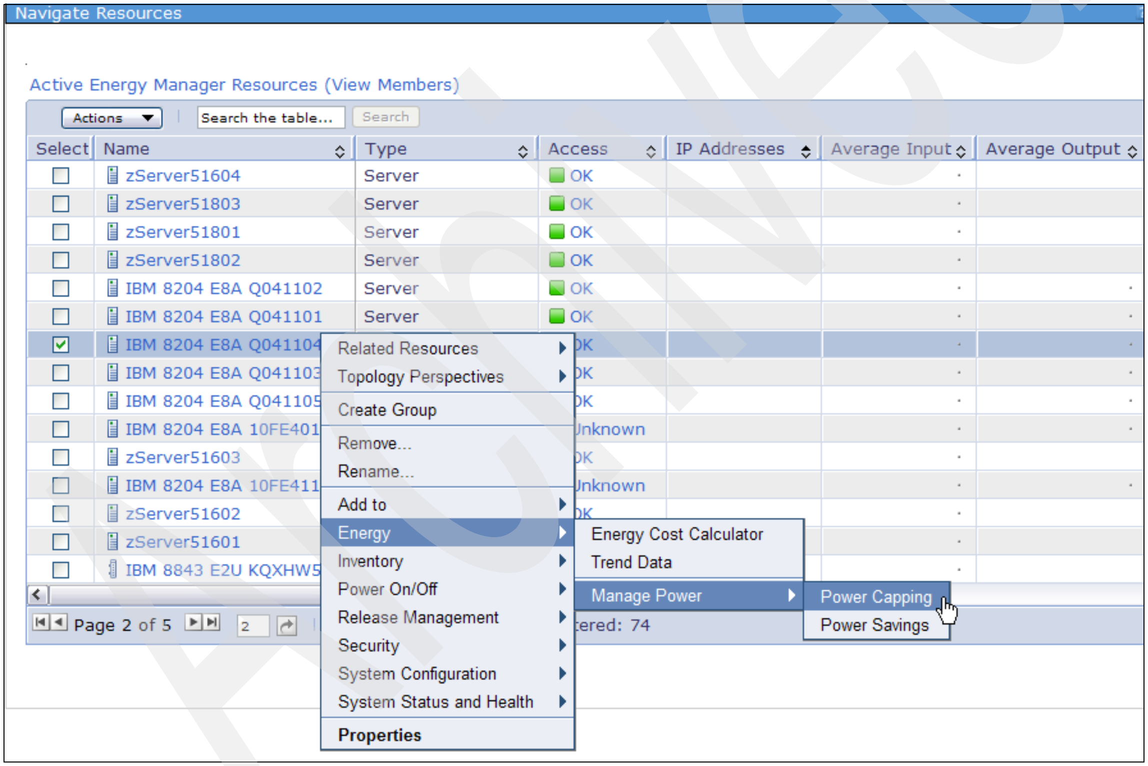


Figure 7-18 Setting a power cap for a single power managed system - 1

The display shown in Figure 7-19 opens.

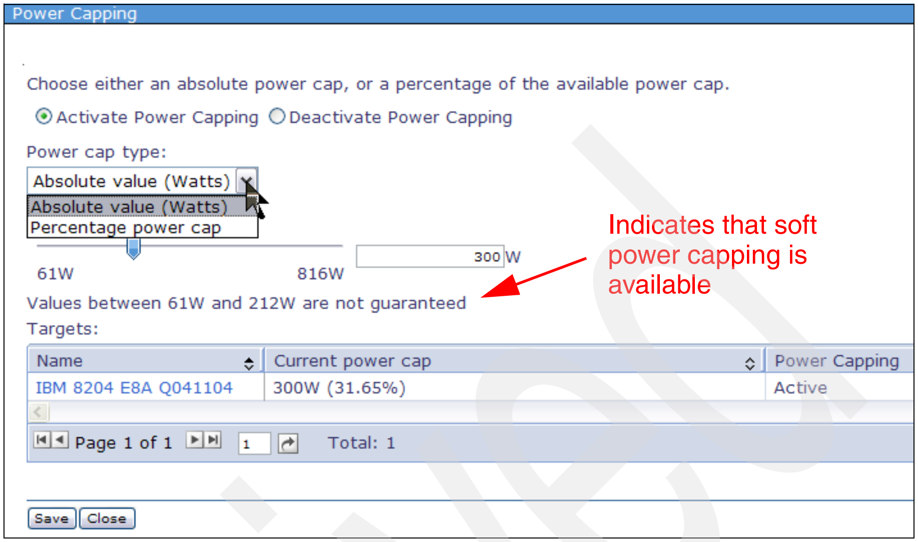


*Figure 7-19 Setting a power cap for a single power managed system - 2*

3. Perform the following tasks:

   a. Activate or deactivate a power cap.

      Select **Activate Power Capping** if you are setting a power cap, or select **Deactivate Power Capping** if you are disabling a previously set power cap.

   b. Choose a hard or soft power cap value.

      Figure 7-19 shows that there are two values in watts shown (61 W and 816 W) between which you can set the power cap value. These two values are calculated by the system firmware and depend on the hardware configuration. On IBM Power systems, the range between the two values can be divided into two regions:

      - The *hard power cap* option is available on all systems that support power capping, not just IBM Power systems. A setting in the hard power cap region is guaranteed.
      - The *soft power cap* option is available *only* on IBM Power systems. A setting in the soft power cap region is not guaranteed.

        If your system supports soft power capping, you see a statement such as the one indicated in Figure 7-19 on page 309:

        ```
        Values between 61W and 212W are not guaranteed.
        ```

For more information about hard and soft capping and how the guarantee works, refer to 1.3.4, “Power capping” on page 23.

> **Note:** On x86 power managed systems, the power capping range is for hard capping only.

By observing historical power trend data as described in 6.7, “Viewing trend data” on page 221, you can determine an appropriate value for the power cap that is appropriate for your requirements.

As shown in Figure 7-19 on page 309, enter the number of watts in the box, or move the slider to the desired value. We entered a value of 300 watts, which in this example is a hard power cap.

c. Set the power cap as either an absolute value or as a percentage.

After you have decided on a power cap value, you need to decide whether to set the power cap as an absolute value or as a percentage, as shown in in the Power cap type drop-down menu in Figure 7-19 on page 309:

- Setting the power cap as an absolute value in watts

  The simplest way to set a power cap is to choose a value in watts somewhere between the values shown in Figure 7-19 on page 309. The range of settings varies between 61 W and 816 W in this example.

- Setting the power cap as a percentage

  The advantage of setting a power cap as a percentage between the minimum and maximum power cap values is that should the hardware configuration of the managed system change, so will the minimum and maximum values as determined by the power capping algorithm. In this case, because you have set the power cap as a percentage, the power cap self adjusts to the new hardware configuration. Therefore, you do not need to worry about having to manually change power cap values every time there is a hardware configuration change to a power managed system to ensure that the power cap value you set does not fall below the minimum hard (or soft) values.

  A second advantage of setting a power cap as a percentage is that when using a policy to set power caps across multiple servers that have different capping ranges, using the percentage option allows this to be done, without having to worry about whether the cap is within the range of all the servers in the group.

  For example, say you turn on additional processors on a supported POWER6 processor-based system, the service processor firmware recalculates the minimum and maximum power cap values, and automatically resets the power cap in watts to a new (higher in this case) value to allow for the additional power consumed. Note that the percentage you set doesn't change, but the underlying setting in watts must rise to allow for the increase in power consumption. In this case you should review the power trending charts to determine the additional power consumed and adjust your power planning and budgetary assumptions accordingly.

d. After you have set the power cap to either an absolute value in watts or a percentage, click **Save**.

4. Figure 7-20 on page 312 shows an example of the trend data for a power managed system that has had a power cap set. It might take some time (a polling interval) for the change in power capping to be visible in the trend data. Notice the following features of the trend data chart:

   – There is an informational message that indicates that power capping is active.

   – A new solid brown line that represents the power cap value appears on the chart.

   – The average output power is reduced to the power cap value because in the example shown Figure 7-19 on page 309, we set the power cap in the hard region, which is guaranteed.

   > **Note:** Power capping works with the output power.

   – The hard power capping region is between the Maximum power cap and Minimum guaranteed power cap lines.

   – The soft power capping region is between the Minimum guaranteed power cap and Minimum power cap lines.

   – When a power cap is activated, deactivated, or modified, it takes some time (a polling interval) for the change to be visible through Active Energy Manager.

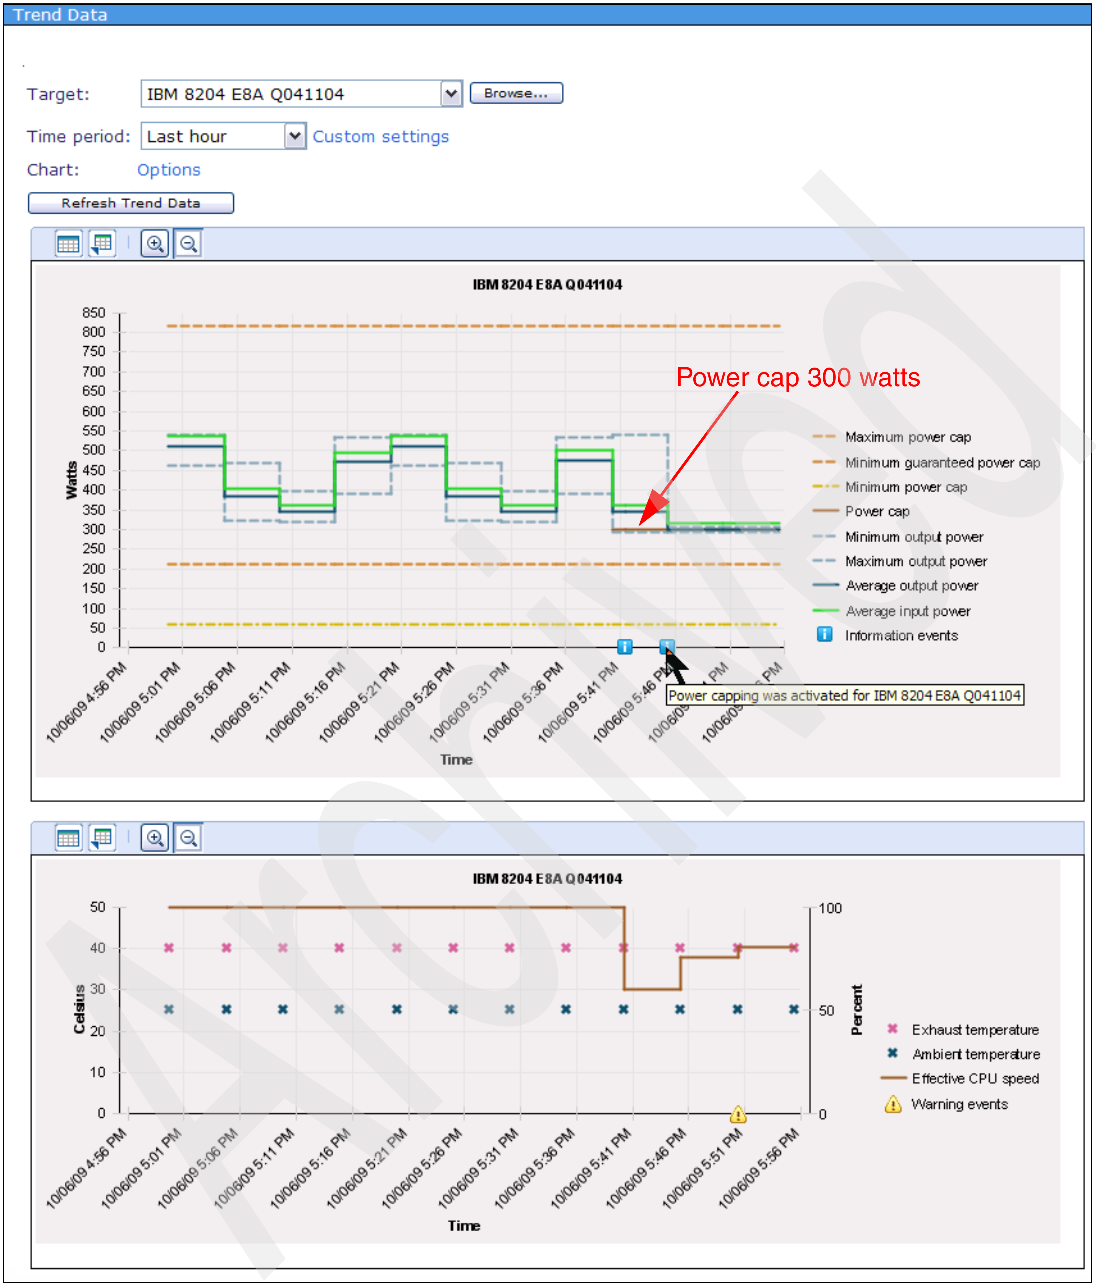


*Figure 7-20 Setting a power cap for a single power managed system - 3*

## 7.3.2 Setting a power cap for a group of systems

When you set a power cap for a group of power managed systems, the power cap you choose is activated on each member of the group. The power cap value you set is for each individual system; it is not an aggregate value that is divided up amongst the members of the group. This is different than when you activate a *group power capping policy*. In this case, the power cap value you set is an aggregate number which is divided up amongst the members of the group, depending on each system's configuration. A power cap value is calculated individually for each member of the group and in fact might be different for each member of the group. To learn how to set a power cap using a group power capping policy, refer to "Creating a group power capping policy" on page 336.

If some resources in a group do not support a specified power cap, the closest available cap is set on that resource.

> **Note:** To activate power capping on a group of power managed systems, all members of the group must support it.

To set a power cap on a group of power managed systems, follow these steps:

1. Close any open Navigate Resources tabs, and then click **Navigate Resources** in the left-hand navigation pane.
2. In the display shown in Figure 7-21, right-click your group of Active Energy Manager resources, and select **Energy** → **Manage Power** → **Power Capping**.

> **Note:** If the context menu does not show the **Power Capping** option, then not all members of the group support it.

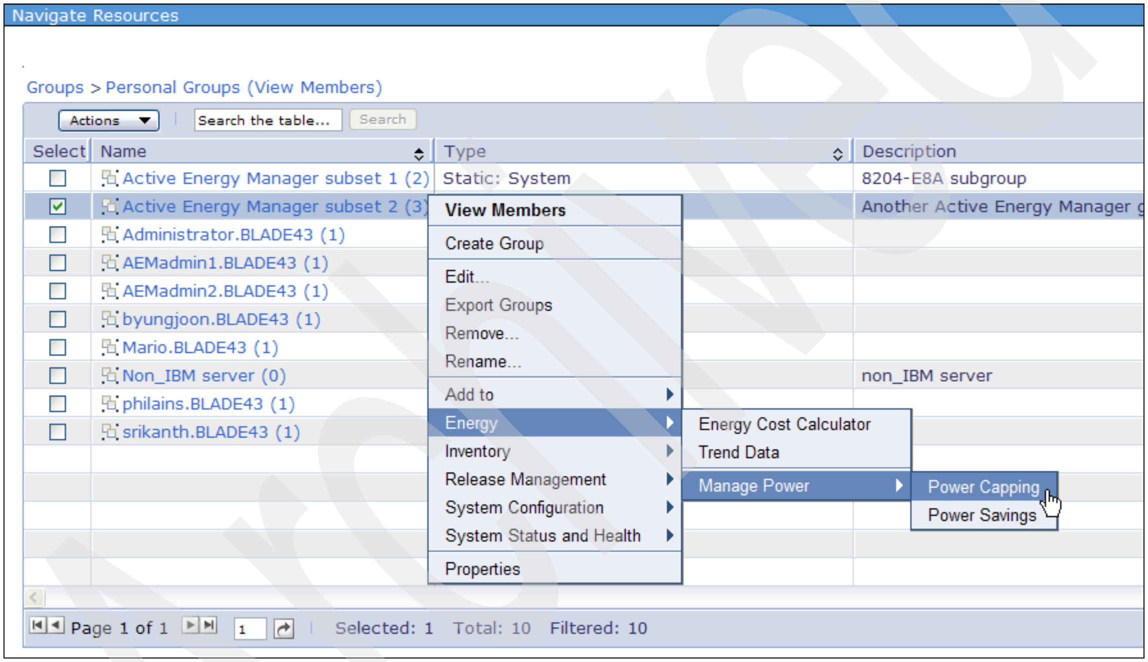


*Figure 7-21 Setting a power cap for a group of power managed systems - 1*

3. In the display shown in Figure 7-22, select **Activate Power Capping**, and then enter a value in the box or move the slider to the desired value.

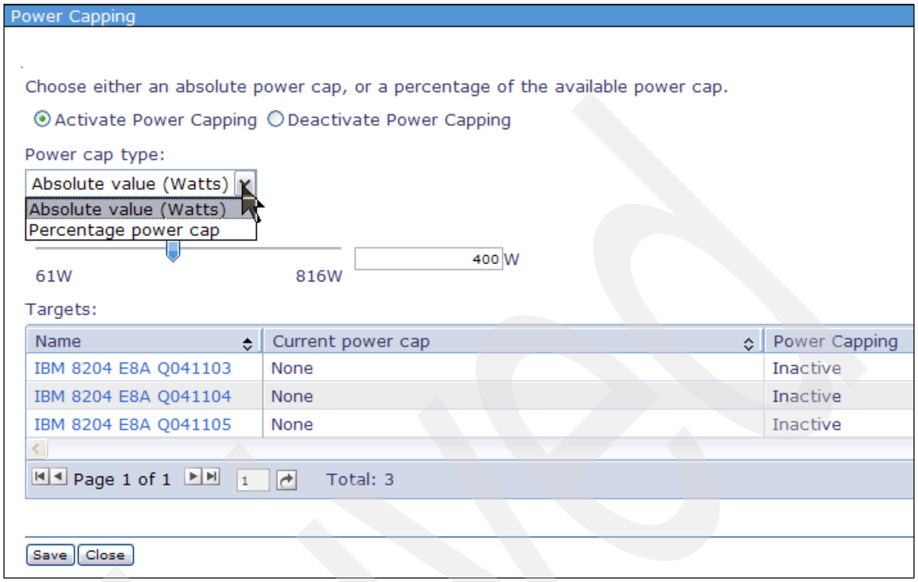


*Figure 7-22 Setting a power cap for a group of power managed systems - 2*

Notice that when setting a power cap for a group of resources, both the absolute and percent power capping options consider the full range of available power caps in all the resources of the group. When applied, if a resource does not support the selected power cap, the closest available cap is applied to that resource.

In Figure 7-22 notice that the soft power capping range is not shown, in contrast to Figure 7-19 on page 309. The hard and soft capping ranges probably differ for each system in the group because these ranges are dependent on the hardware configuration. Therefore, the power cap that you select can be a hard power cap on some members of the group and a soft power cap on others. If you want to set a guaranteed (hard power cap) on each system, it might be better to set power caps individually on each system so that you can be sure that the cap is set in the hard capping region.

We choose a value of 400 W, which means that each system in the group has a power cap setting of 400 W imposed. Click **Save**.

4. After a few minutes you see the status in the Current power cap column change to the setting you chose (Figure 7-23). If you check the trend data for each of the power managed systems, you see that power capping has been activated on each one.

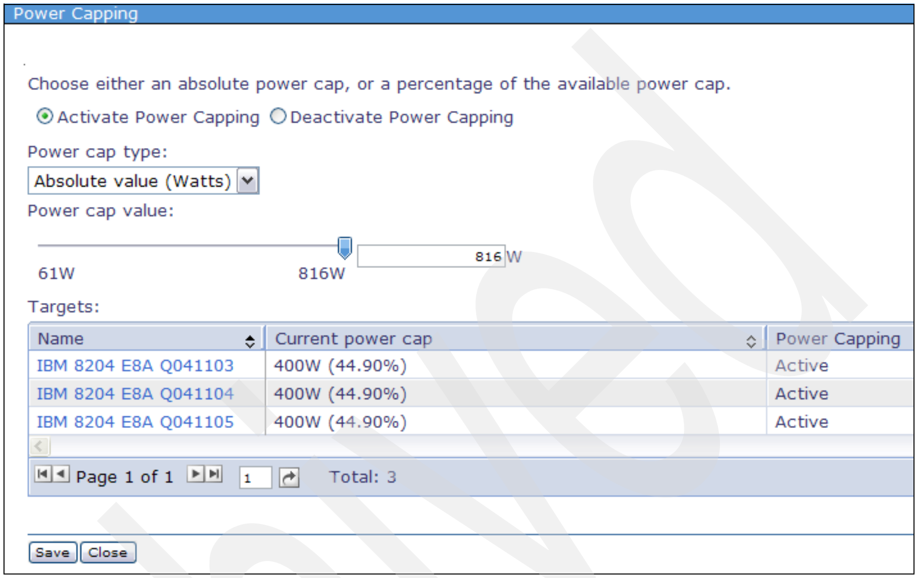


*Figure 7-23 Setting a power cap for a group of power managed systems - 3*

5. Figure 7-24 on page 317 shows the trend data graph for the group after the power cap has been set to 400 watts (around 4:47 pm). Notice that, for a group, only the aggregate input power is shown, where the aggregate power cap for the group is 3 x 400 = 1,200 watts. You can see that the aggregate input power is actually slightly greater than the aggregate power cap of 1,200 watts. The reason for this is that the power cap is based on output power (DC) whereas it is the aggregate input power that appears on the graph. The input power is always slightly greater than the output power because there are power losses in the power supplies that are converting input power (AC) to output power (DC).

   When a power cap is activated, deactivated, or modified, it takes some time (a polling interval) for the change to be visible through Active Energy Manager.

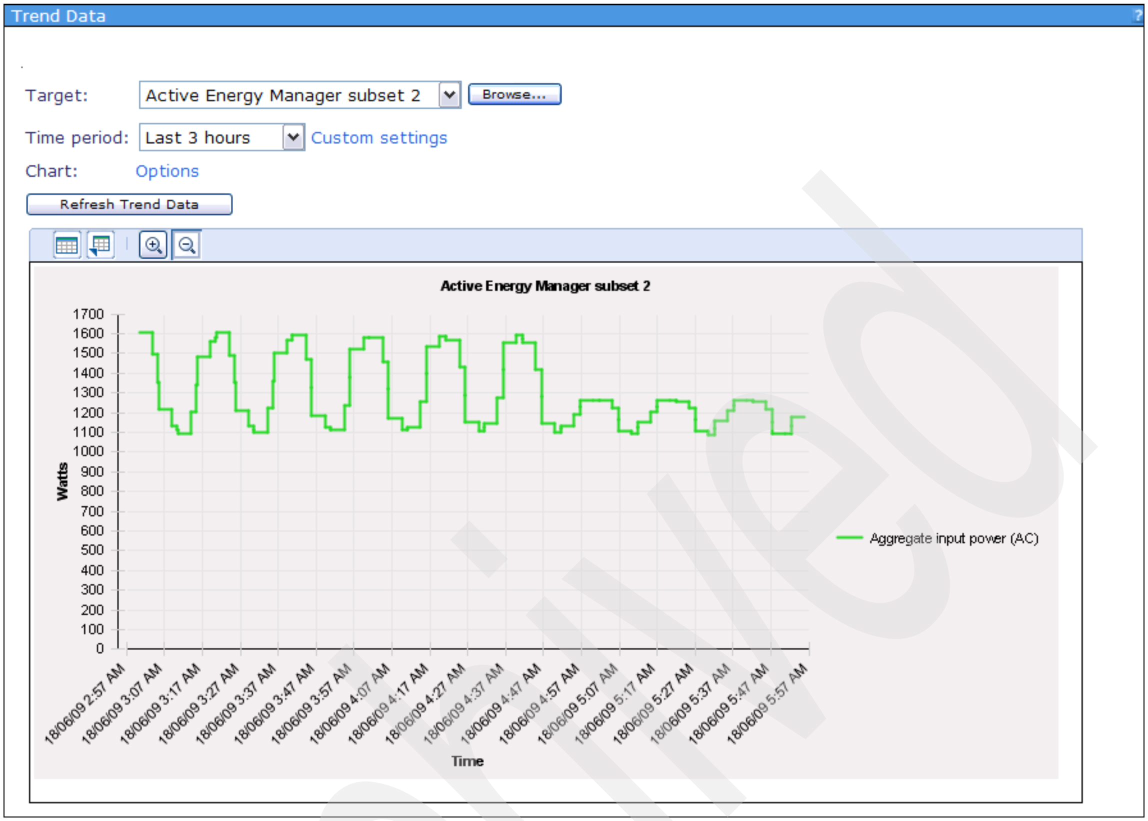


*Figure 7-24 Setting a power cap for a group of power managed systems - 4*

## 7.3.3 Power capping examples

In this section, we provide two examples of how to use power capping.

### Power capping example 1

An installation has twelve JS22 POWER6 blade servers. The UL rating on a JS22 server is approximately 402 W (DC). Based on that value, the data center manager has budgeted approximately 4.8 KW (DC) of power in the data center. Using Active Energy Manager, the system operators notice that the twelve JS22s are only consuming a maximum of 308W (DC) of power while running their normal workload.

The firmware, based on the hardware configuration, allows a minimal hard power cap of 368 W (DC). Based on the power consumption the operators have seen on the Active Energy Manager trending charts, they are confident in setting a power cap at the minimum guaranteed value of 368 W (DC). The operator sets the power cap of each of the twelve systems to 368 W (DC), and the data center manager is now guaranteed that the average power over time for each system will never exceed 368 W (DC). This frees up approximately 408 W (DC) power within the data center, calculated as follows:

```
(402W(DC) – 368W(DC)) * 12 = 408W(DC)
```

The data center manager can now be confident about adding a thirteenth JS22 without exceeding the original power budget.

## Power capping example 2

In the Figure 7-25 on page 319, a POWER6 processor-based system has power capping and power savings disabled. Its workload shows a cyclical behavior, varying between idle and heavy utilization over the course of several hours. Average output power (DC) varies between 770 watts when idle, and 1,001 watts at peak load. The CPU frequency remains constant at 100% for the first three cycles of the workload. The maximum output power (DC) tends to vary around 1,000 watts, and the minimum output power (DC) tends to vary around 750 watts, near the server's idle power consumption.

Shortly before 18:00:00, an administrator sets a power cap of 925 watts DC for this server to reallocate part of the server's power budget to another system. For the next several workload cycles, peak and average power consumption are *clipped* as the system firmware enforces the power cap. This clipping is evident in the lower peak and average power consumption and slight variations in the previously static average effective CPU speed.

At approximately 20:15:00, the power cap was lowered further to a slightly more aggressive 875 watts value. As the workload continues to cycle, peak and average power consumption are again clipped. However, average effective CPU frequency over these periods is impacted more, as the platform firmware adjusts the processor frequency and voltage to enforce the cap. A slight performance loss occurs in the workload during periods of heavy utilization as the system firmware automatically throttles back the frequency and voltage of the system's processors.

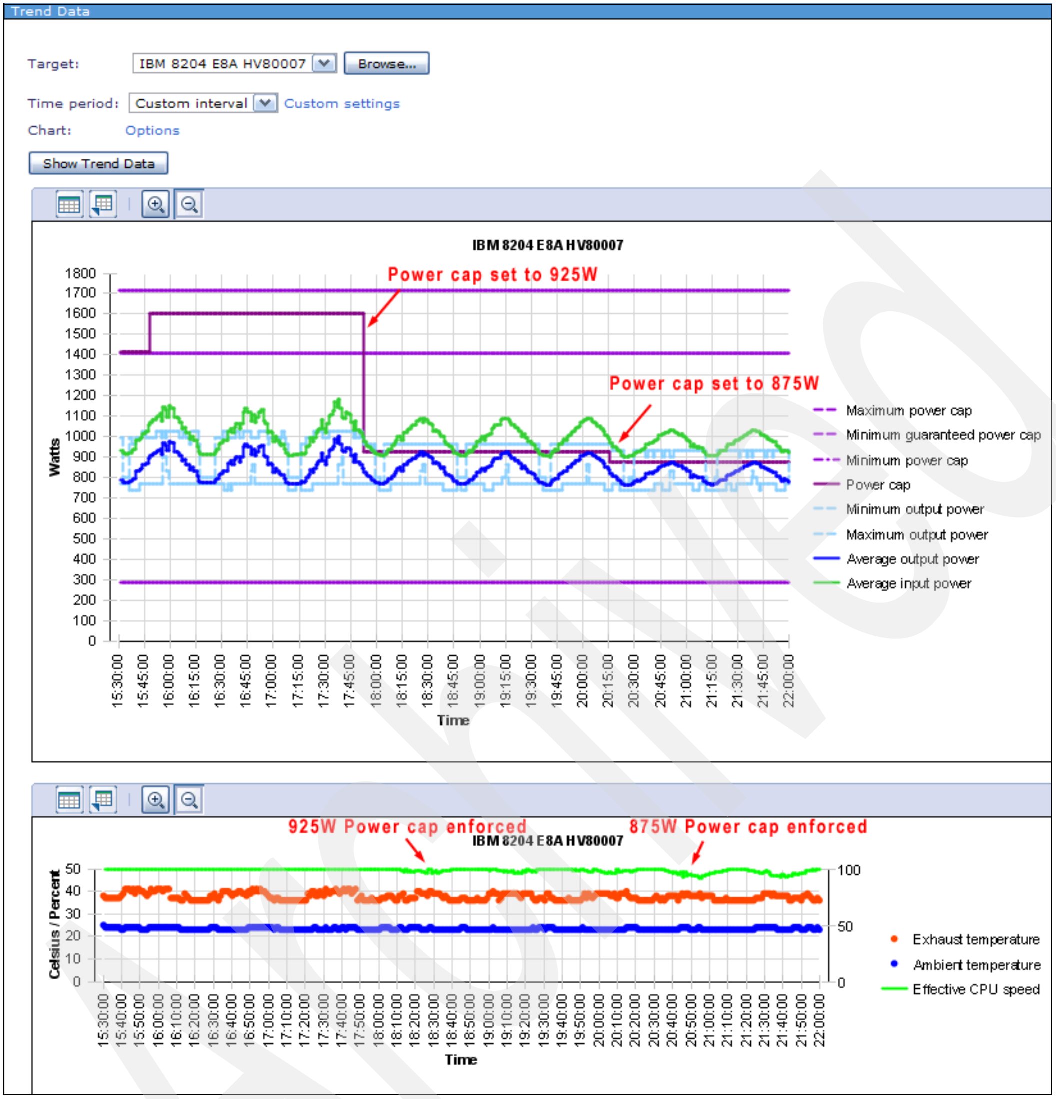


*Figure 7-25 Power capping example 2*

## 7.4 Managing power on a BladeCenter server

As an alternative to Active Energy Manager, you can set a power savings mode or a power cap for a POWER6 blade server through the management module of the associated BladeCenter chassis. The result is the same as though you made the setting through Active Energy Manager.

To set a power savings mode or a power cap for a POWER6 blade server through the management module, follow these steps:

1. Sign on to the BladeCenter management module.
2. On the left hand navigation pane, expand **Monitors**. Then, click **Power Management** and click **Power Domain 1** (or **Power Domain 2**, depending on where the blade is located) as shown in Figure 7-26.

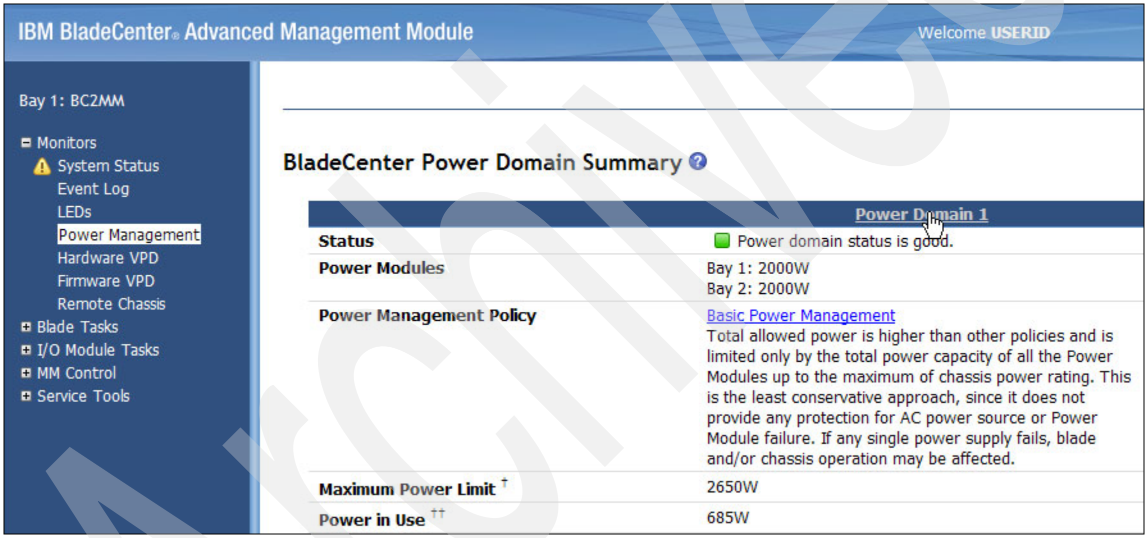


*Figure 7-26 Managing power on a POWER6 blade server - 1*

3. Click the POWER6 blade server for which you want to set power savings or power capping (Figure 7-27).

| Bay (s) | Status | Module | State | Power In Use | Allocated Power Maximum | Allocated Power Minimum | CPU Duty Cycles |
|---|---|---|---|---|---|---|---|
| *Chassis Components* | | | | | | | |
| | | Midplane | On | 5W | 5W | 5W | n/a |
| 1 | | Media Module | On | 5W | 5W | 5W | n/a |
| *Power Module Cooling Devices* | | | | | | | |
| 1 | | Power Module | On | 15W | 15W | 15W | n/a |
| 2 | | Power Module | On | 15W | 15W | 15W | n/a |
| 3 | | Power Module | On | 15W | 15W | 15W | n/a |
| 4 | | Power Module | On | 15W | 15W | 15W | n/a |
| *Management Modules* | | | | | | | |
| 1 | | AMM614637145 | On | 12W | 13W | 13W | n/a |
| 2 | | Advanced Management Module Bay 2 (not present) | | 0W | 8W | 8W | n/a |
| *I/O Modules* | | | | | | | |
| 1 | | Ethernet SM | On | 22W | 23W | 23W | n/a |
| 2 | | Ethernet SM | On | 22W | 23W | 23W | n/a |
| 3 | | Fibre Channel SM | On | 45W | 45W | 45W | n/a |
| 4 | | Fibre Channel SM | On | 45W | 45W | 45W | n/a |
| *Blades* | | | | | | | |
| [1] | | JS22 | Off | 6W | 402W | 368W | n/a †† |
| [2] | | JS12-100EB0A | On | 153W | 246W | 220W | n/a †† |
| [4] | | HS21-KQGZCK2 | On | 135W | 203W | 161W | n/a †† |

*Figure 7-27 Managing power on a POWER6 blade server - 2*

4. If the blade server is a POWER6 blade, you see a display similar to the one shown in Figure 7-28. From this display, you can set power savings and power capping. For x86 blade servers, you can set a power cap on supported blades, but not a power savings mode.

   Notice that in Figure 7-28, **Dynamic Power Saver is** equivalent to dynamic power savings (favor power savings). If you want to set dynamic power savings (favor performance), select **Favor Performance over Power**.

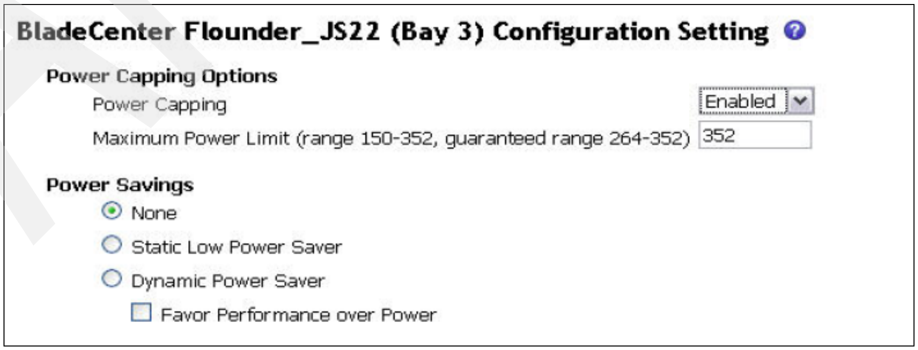


*Figure 7-28 Managing power on a POWER6 blade server - 3*

5. If you now go into Active Energy Manager, the power savings mode or power cap that you just set through the management module is reflected accurately after the next polling interval.

## 7.5 Working with policies

We cover the following topics in this section:

- Introducing policies
- Creating a policy
- Applying a policy
- Deactivating a policy
- Removing a policy
- Editing a policy
- Deleting a policy

### 7.5.1 Introducing policies

You can create a policy to apply power savings and power capping settings to a power managed system or group of systems.

Note the following important points about policies:

- A policy is a stand-alone Active Energy Manager construct. When you create a policy, you do not specify any target systems or groups.
- After you create a policy you can then apply it to target systems and groups. A policy does not take effect until it is applied to a target.
- Policies can be applied to, and removed from, individual power managed systems or groups of systems in the following ways:
  - Interactively
  - By using a schedule
  - By using an automation plan based on an Active Energy Manager event
- Active Energy Manager constantly monitors each power managed system for events or changes that would require a power capping or power savings policy change.

There are two basic types of policies:

- *System policies* can be:
  - System power savings policies
  - System power capping policies

  A system power savings or power capping policy can be applied to individual systems or systems in a group. When a system policy is applied to a group, the policy is applied individually to each system currently in the group. Note that applying a system power savings or power capping policy to a group is simply a shortcut to help you avoid applying the policy to each target system individually. If any system is removed from the group, the policy continues to be associated with that system, although it does continue to be enforced. If a new system is added to the group, the previously applied policy will not be associated with that system.

  To learn how to create a system power savings policy, refer to “Creating a system power savings policy” on page 327.

  To learn how to create a system power capping policy, refer to “Creating a system power capping policy” on page 332.

- *Group policies* apply to power capping only, and allow an administrator to assign a power budget to an entire group of servers.

  A group power capping policy specifies an overall power cap that the systems in the group collectively share and cannot exceed. A group power capping policy can only be applied to a group. It cannot be applied to an individual resource.

  > **Note:** Group power capping policies can only apply hard power caps. You cannot specify a soft power cap.

  It is important to recognize that the group power capping policy is associated with the group itself, not with the individual systems in the group. Therefore, individual systems in the group can have a different power cap value set. If a system is either added to or removed from the group, Active Energy Manager dynamically manages the power budget for each group member as the group changes, or as the power cap is adjusted.

  A system can be assigned to one or more groups dynamically or statically by an administrator. Systems, for example, might belong to groups organized by function, system type, or location, and a group power capping policy can be applied to each group.

  When a group power cap is applied, Active Energy Manager attempts to divide the group cap evenly across the members of the group. If the calculated power cap falls outside of the minimum to maximum power capping range for one or more group members, Active Energy Manager attempts to redistribute the power cap among the remaining group members so that it falls within the minimum and maximum for all systems. If Active Energy Manager cannot calculate an acceptable power cap for each system, the group power cap is not applied.

When a system is added to, or removed from, a group, the power cap is reallocated automatically among the remaining members of the group. If a system is removed from the group, the group policy no longer applies to that system, and the other systems in the group get a greater share of the aggregate power cap value.

Group power capping policies are possible because power consumption can be aggregated across group members. Group power savings policies are not allowed because a power savings policy affects the performance characteristics of individual systems and cannot logically extend across group members.

In contrast to system power savings and system power capping, there is no option to create a group power capping policy to turn off (deactivate) capping. You need to deactivate the group power capping policy and then disable power capping for each system in the group interactively.

To learn how to create a group power capping policy, refer to "Creating a group power capping policy" on page 336.

### Enforcing policies

A key characteristic of Active Energy Manager policies is the way in which they are enforced. Policies are enforced as follows:

Policies are enforced at every poll by Active Energy Manager on the systems and groups to which the policies are applied. For example, if you apply a system power capping *policy* of 700 W to a system, then later *set* the power cap for the system to 800 W, Active Energy Manager changes the power cap back to the value set by the power policy of 700 W.

System power savings, system power capping and group power capping policies do not conflict or interact with each other. Therefore, a system can be subject to the following policies simultaneously:

- One system power savings policy
- One system power capping policy
- Any number of group power capping policies (one policy for each group to which the system belongs)

If a second system power savings or power capping policy is applied to the system, it replaces the original system power savings or power capping policy.

If multiple group power capping policies are applied to a group, each group capping policy is used to compute a per system cap for each system in the group. A single system, by virtue of being in multiple groups, can have several of these power-system caps (one from each group). The lowest of the per-system caps is the one that actually is applied to the system.

If there is a mixture of system power capping policies and group power capping policies, Active Energy Manager applies the policy that specifies the lowest supported power cap for each system.

If a group with a group power capping policy applied has a system added or removed, you need to re-evaluate or re-enforce the group policy.

If a policy is edited by a user, all targets to which the policy is applied will have the new policy re-enforced on them.

If, due to a configuration change on a system, the power cap minimum or power cap maximum changes, then you need to re-enforce the policies if possible. In cases where a re-evaluated and re-enforced policy cannot be enforced successfully, an event is generated and the policy properties page indicates that the policy is not enforced successfully.

## Automating policies using event automation plans

You can tie the activation of a policy to the triggering of an event for an Active Energy Manager monitor using an automation plan. Because applying and removing a power policy are implemented as IBM Systems Director tasks, applying and removing a policy can be specified as actions to take in an event automation plan.

Refer to 6.13, “Working with automation plans” on page 265 for a discussion of how to do this.

## Policies examples

This section provides two examples to help you understand how policies work.

### *Policies example 1*

System A is a member of three groups (1, 2, and 3), and has *all* the policies shown in Table 7-1 on page 326 applied. Groups 1 and 2 have group power capping policies applied, resulting in a per system power cap of 825 watts and 775 watts respectively. An administrator has also applied a system power savings policy to all members of group 3 and a system power capping policy of 900 watts to system A.

As group policies are updated and group members change, Active Energy Manager constantly resolves all of the policies affecting system A and enforces the appropriate power cap or power savings setting.

In this example, Active Energy Manager enables static power savings mode and sets a power cap of 775 watts on system A because the group 2 power cap is the lowest of all applied policies. If the group 2 power cap is removed, Active Energy Manager automatically applies the next lowest power cap of 825 W from group 1 to the target system.

*Table 7-1 Policies example 1*

| Policy type | Policy applied | Policy effect |
|---|---|---|
| System power savings: Group 3 | Yes | Enable static power savings |
| System power cap: System A only | Yes | 900 watts |
| Group power cap: Group 1 | Yes | 825 watts |
| Group power cap: Group 2 | Yes | 775 watts |

### *Policies example 2*

A group contains three systems, each with a minimum power cap of 500 watts and maximum power cap of 1,000 watts as shown in Table 7-2. An administrator sets a group power capping policy of 2,400 watts that are evenly distributed to each system as an 800 watt power cap (2,400 watts / 3 servers = 800 watts per server).

When a fourth system is added to the group, Active Energy Manager automatically reallocates power among the group members, resulting in a 600 watt power cap per system. When a fifth system is added, Active Energy Manager attempts to recalculate the power budget for each system, but cannot meet the group power cap because each server cannot guarantee a 480 watt (2,400 watts / 5 systems) cap. The group power cap is automatically disabled.

When the group power cap is increased to 3,500 watts, Active Energy Manager again attempts to calculate a new power cap per system and can now successfully apply a 700 watt power cap to each system.

*Table 7-2 Policies example 2*

| Group members | Group power cap | Minimum power cap | Maximum power cap | Power cap per system |
|---|---|---|---|---|
| 3 systems | 2,400 watts | 500 watts | 1,000 watts | 800 watts |
| 4 systems | 2,400 watts | 500 watts | 1,000 watts | 600 watts |
| 5 systems | 2,400 watts | 500 watts | 1,000 watts | None |
| 5 systems | 3,500 watts | 500 watts | 1,000 watts | 700 watts |

## 7.5.2 Creating a policy

In this section, we show how to create the following types of policy:

- System power savings policy
- System power capping policy
- Group power capping policy

### Creating a system power savings policy

To create a system power savings policy using the power policy editor wizard, follow these steps:

1. Start the power policy editor wizard by clicking **Work with power policies** in the Manage section of the Active Energy Manager home page, as shown in Figure 7-29.

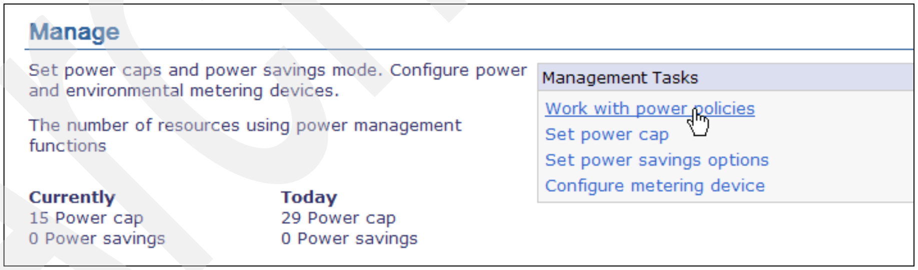


*Figure 7-29 Creating a system power savings policy - 1*

2. The Power Policies display lists all the policies that you created. In the example shown in Figure 7-30, we have not yet created any policies. Click **Create policy** to start the wizard.

   Alternatively, if you want to create a policy that is like an existing policy, click **Create like** to start the power policy editor wizard and to insert automatically the properties of an existing policy that you choose. Then, you need to click only through the options and make changes to the existing policy, rather than creating a new policy from scratch.

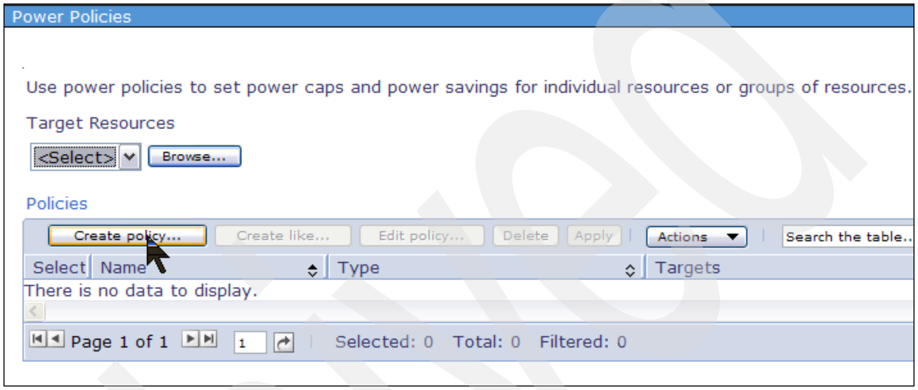


*Figure 7-30 Creating a system power savings policy - 2*

3. In the power policy editor wizard Welcome display, click **Next**.

4. In the display shown in Figure 7-31, enter a name and description for the new policy. Then, click **Next**.

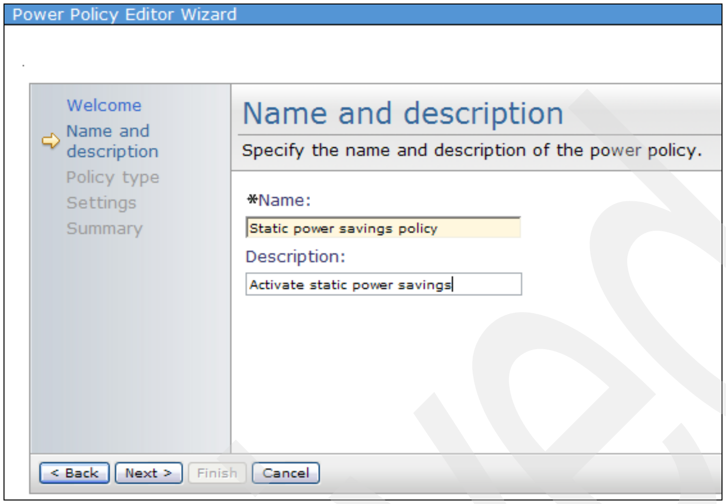


*Figure 7-31 Creating a system power savings policy - 3*

5. In the display shown in Figure 7-32, select the policy type. In this example, we select **System power savings**. Click **Next**.

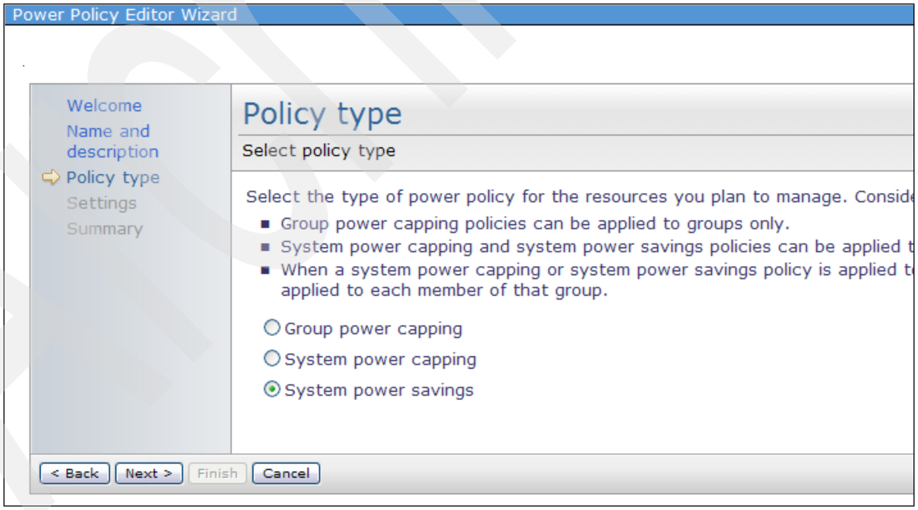


*Figure 7-32 Creating a system power savings policy - 4*

6. Because we selected System power savings in Figure 7-32 on page 329, we need to specify the type of power savings policy. In our example, we select **Static power savings** in Figure 7-33. Click **Next**.

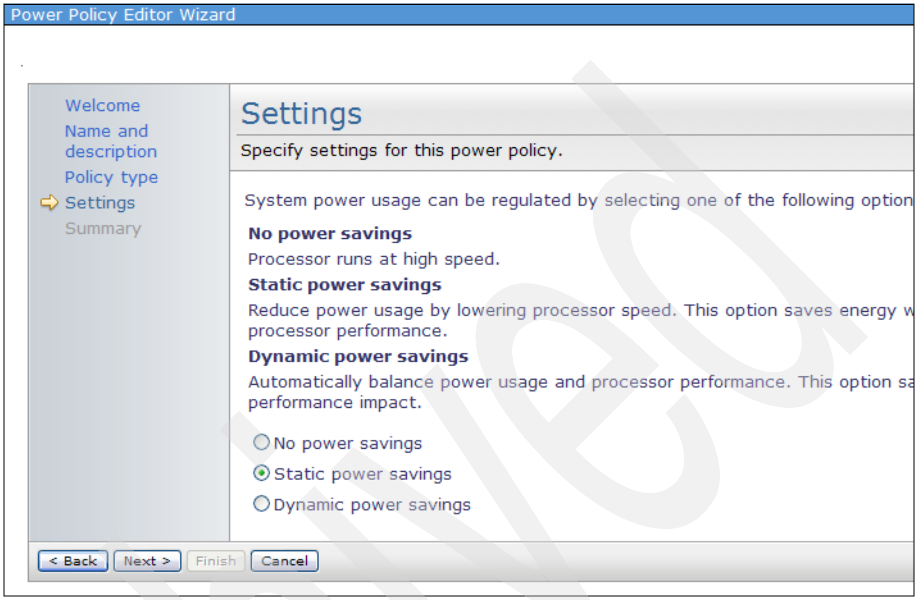


*Figure 7-33 Creating a system power savings policy - 5*

7. The Summary window displays, as shown in Figure 7-34. Click **Finish**.

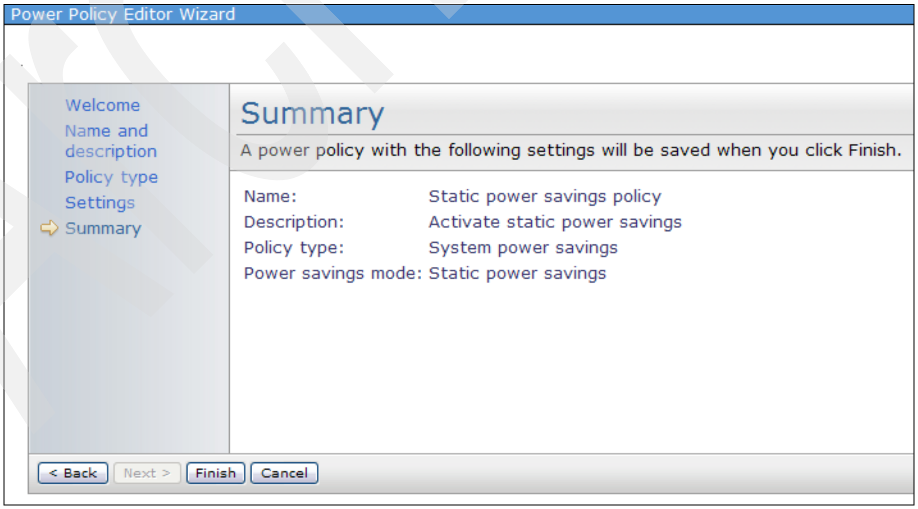


*Figure 7-34 Creating a system power savings policy - 6*

The policy is created and appears in the list of policies shown in Figure 7-35.

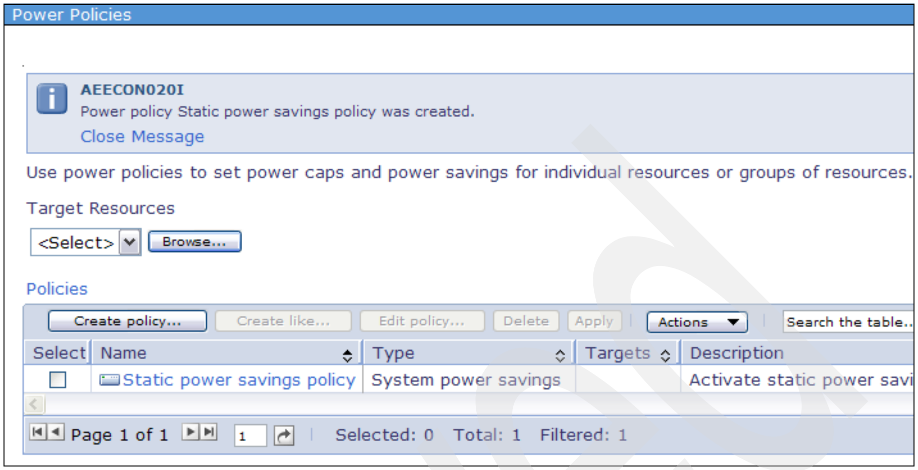


*Figure 7-35 Creating a system power savings policy - 7*

8. Clicking a policy name in the Power Policies view shows the power policy's properties (Figure 7-44 on page 336).

   The Targets tab shows all the systems to which this policy is applied (or groups, if it is a group power capping policy), as shown in Figure 7-36 on page 332.

> **Note:** You use the Targets tab to remove a policy from a target system or group. Refer to 7.5.5, "Removing a policy" on page 353 for more information.

You can now apply this policy to a power managed system or group of systems as described in 7.5.3, "Applying a policy" on page 341.

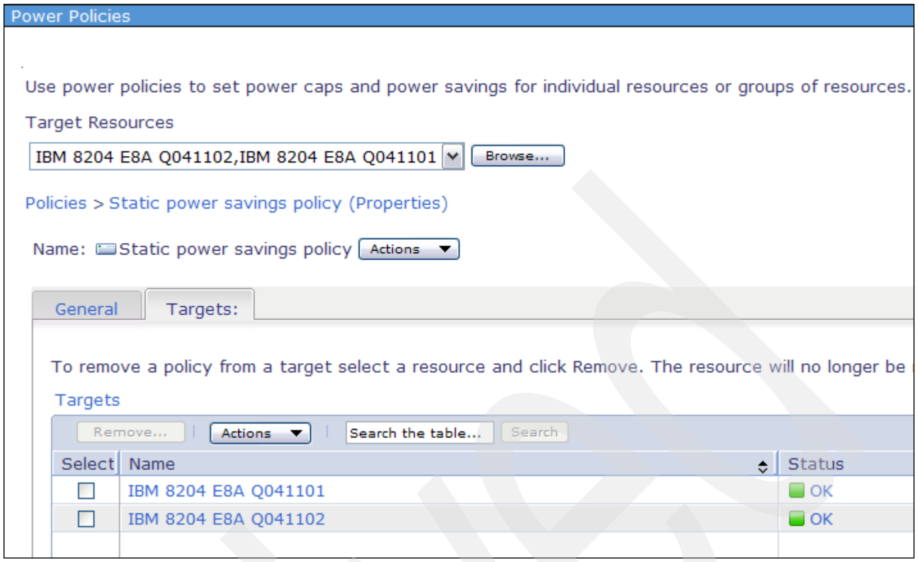


*Figure 7-36 Creating a system power savings policy - 8*

## Creating a system power capping policy

To create a system power capping policy using the power policy editor wizard, follow these steps:

1. Start the power policy editor wizard by clicking **Work with power policies** in the Manage section of the Active Energy Manager home page, as shown in Figure 7-37.

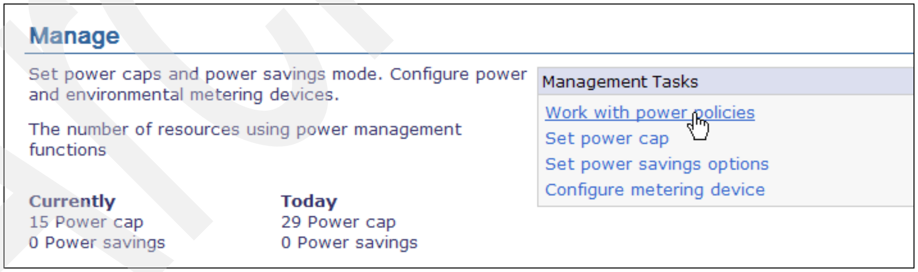


*Figure 7-37 Creating a system power capping policy - 1*

2. The Power Policies display lists all the policies that you have created. The example shown in Figure 7-38 shows the policy created previously in "Creating a system power savings policy" on page 327. Click **Create policy**.

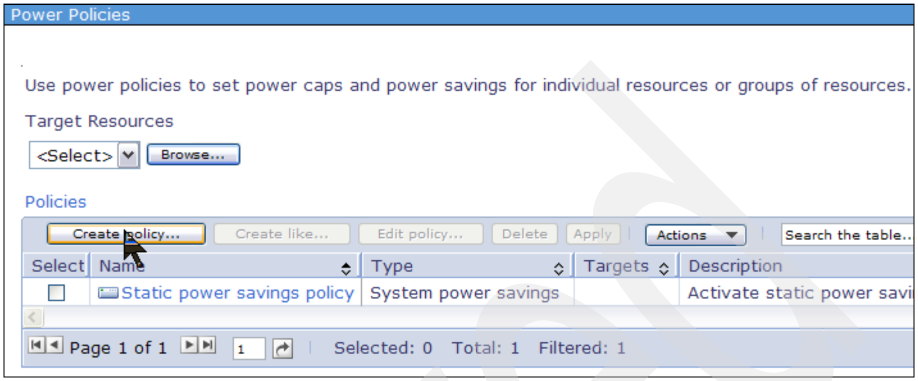


*Figure 7-38 Creating a system power capping policy - 2*

3. In the Welcome display, click **Next**.
4. In the display shown in Figure 7-39, enter a name and description for your new policy. Click **Next**.

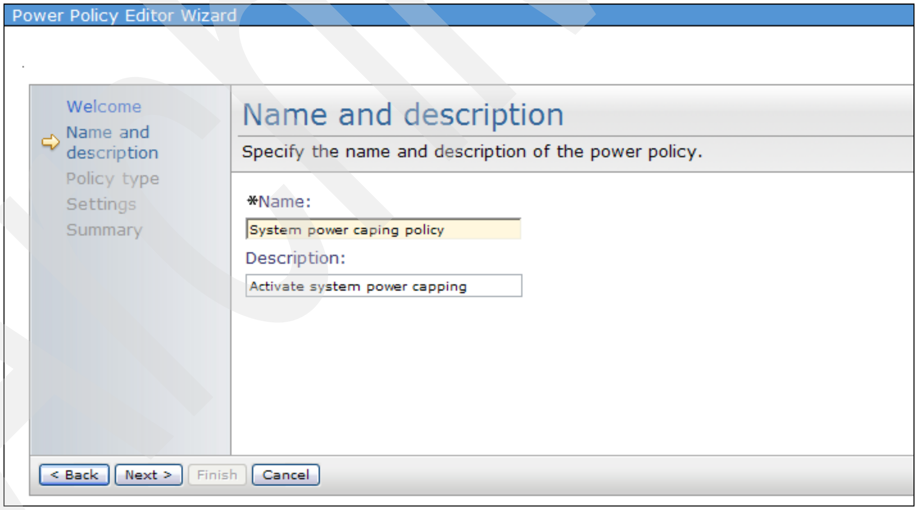


*Figure 7-39 Creating a system power capping policy - 3*

5. In the display shown in Figure 7-40, select the policy type. In this example we select **System power capping**.

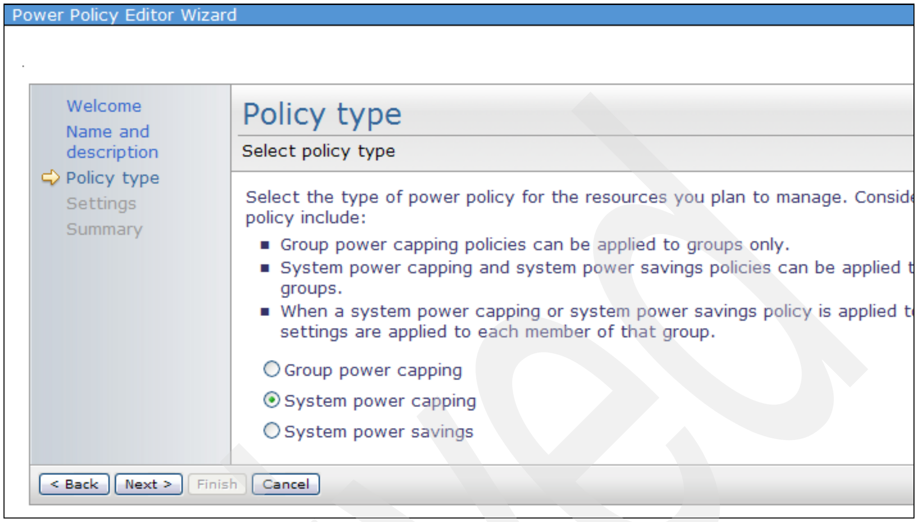


*Figure 7-40 Creating a system power capping policy - 4*

6. In the display shown in Figure 7-41, select **Activate Power Capping**, and select a **Power cap type** and **Power cap value**. If you are unsure about how to select a Power cap type or Power cap value, refer to 7.3, “Setting a power cap” on page 306. Click **Next**.

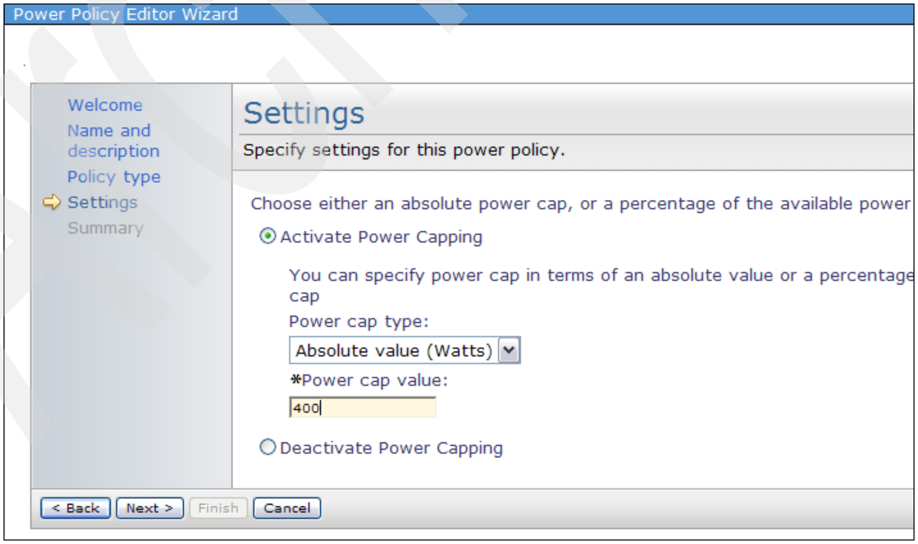


*Figure 7-41 Creating a system power capping policy - 5*

7. In the Summary window shown in Figure 7-42, click **Finish**.

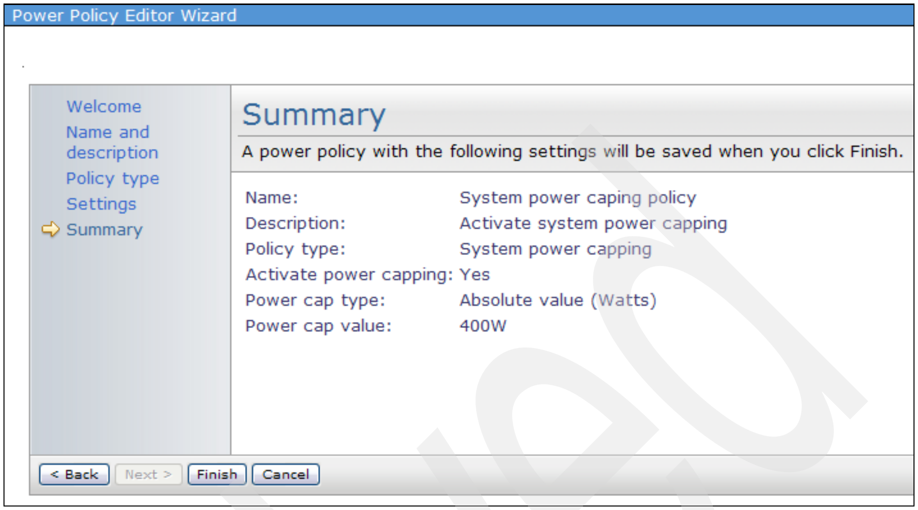


*Figure 7-42 Creating a system power capping policy - 6*

The policy is created and displays in the list of policies shown in Figure 7-43.

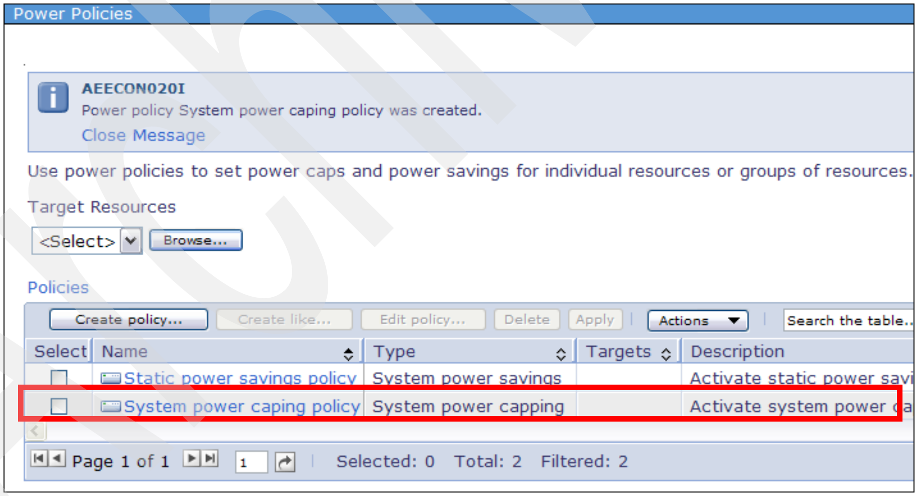


*Figure 7-43 Creating a system power capping policy - 7*

8. Clicking a policy name in the Power Policies view shows the power policy's properties (Figure 7-44 on page 336).

   The Targets tab shows all the systems that this policy is applied to (or groups if it is a group power capping policy).

**Note:** You use the Targets tab to remove a policy from a target system or group. Refer to 7.5.5, “Removing a policy” on page 353 for more information.

You can now apply this policy to a power managed system or group of systems as described in 7.5.3, “Applying a policy” on page 341.

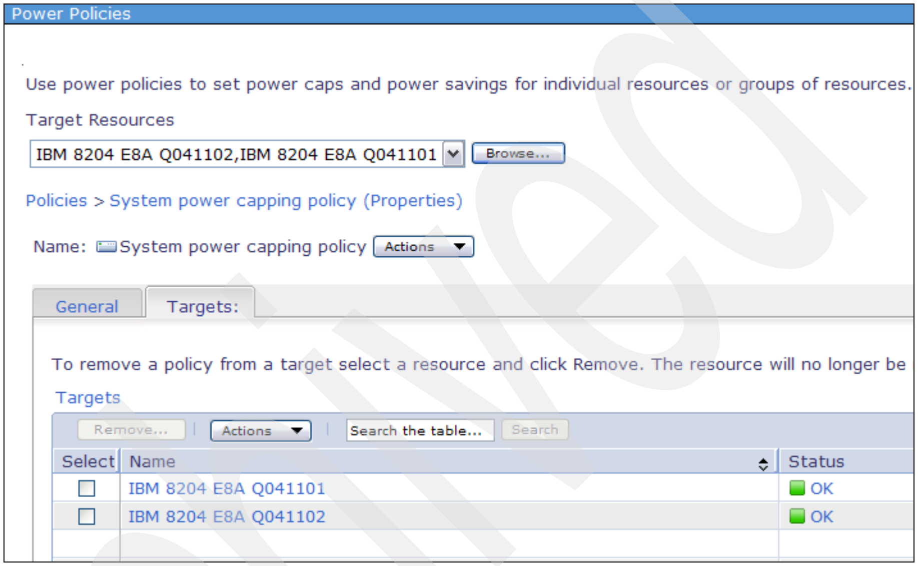


*Figure 7-44 Creating a system power capping policy - 8*

## Creating a group power capping policy

Creating a group power capping policy is different than setting a power cap for a group of power managed systems. When you create a group power capping policy, the power cap value you set is an aggregate number which is divided up amongst the members of the group.

A power cap value is calculated individually for each member of the group and, in fact, can be different for each member of the group. However, all the individual power cap values must add up to the aggregate value. In this way, you can allocate a block of power to be shared amongst a group of systems, each according to its individual needs.

For more information about group power capping policies, refer to 7.5.1, “Introducing policies” on page 322.

To create a group power capping policy using the power policy editor wizard, follow these steps:

1. Start the power policy editor wizard by clicking **Work with power policies** in the Manage section of the Active Energy Manager home page, as shown in Figure 7-45.

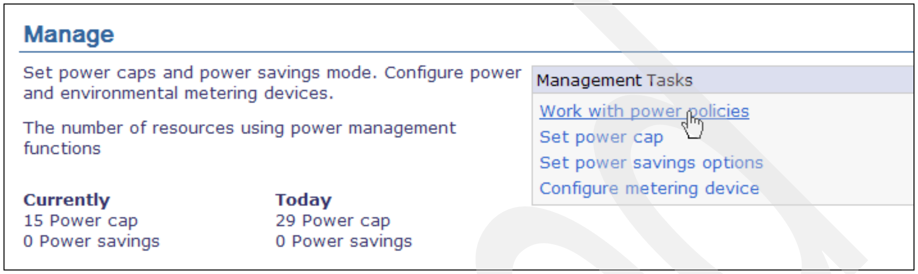


*Figure 7-45 Creating a group power capping policy - 1*

2. The Power Policies display lists all the policies that you have created. The example shown in Figure 7-46 shows the policies created in previous examples. Click **Create policy**.

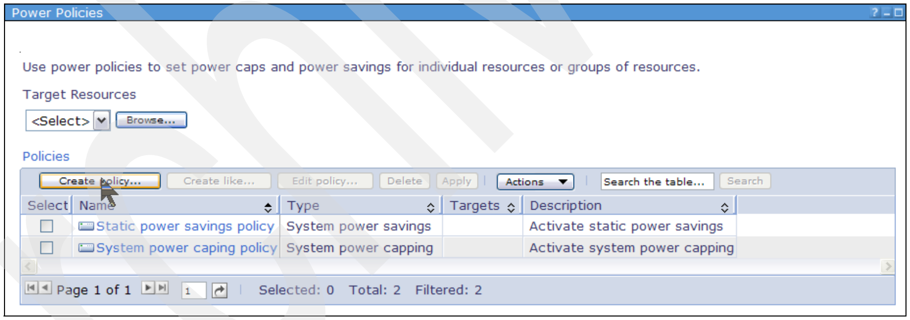


*Figure 7-46 Creating a group power capping policy - 2*

3. In the Welcome display, click **Next**.

4. In the display shown in Figure 7-47, enter a name and description for your new policy. Click **Next**.

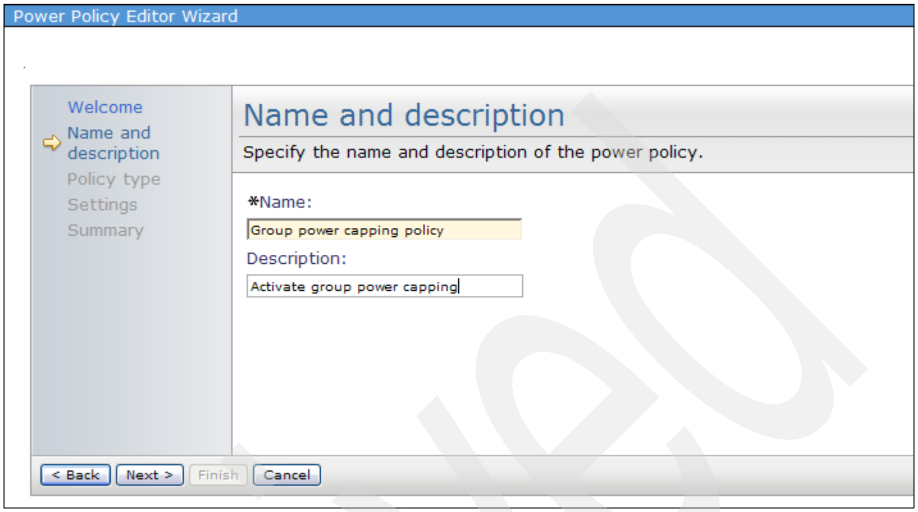


*Figure 7-47 Creating a group power capping policy - 3*

5. In the display shown in Figure 7-48, select the policy type. In this example, we select **Group power capping**.

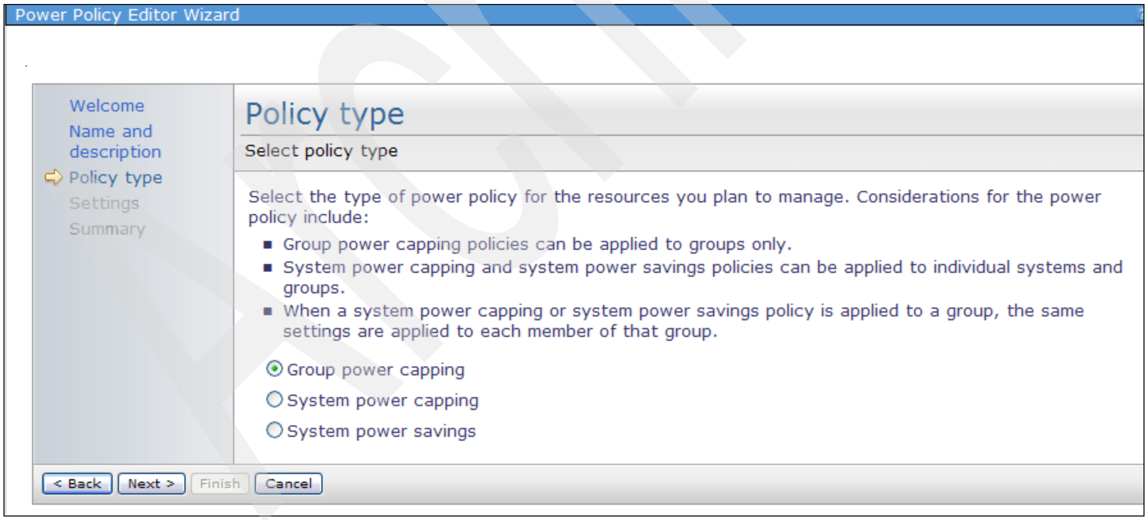


*Figure 7-48 Creating a group power capping policy - 4*

6. In the display shown in Figure 7-49, select a **Power cap type** and **Power cap value**. Click **Next**.

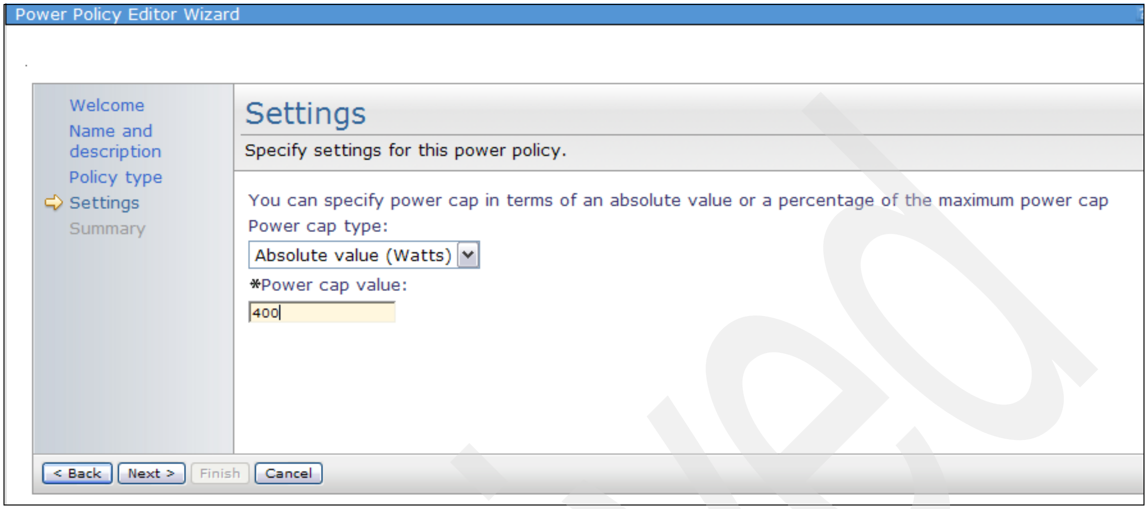


*Figure 7-49 Creating a group power capping policy - 5*

7. In the Summary window shown in Figure 7-50, click **Finish**.

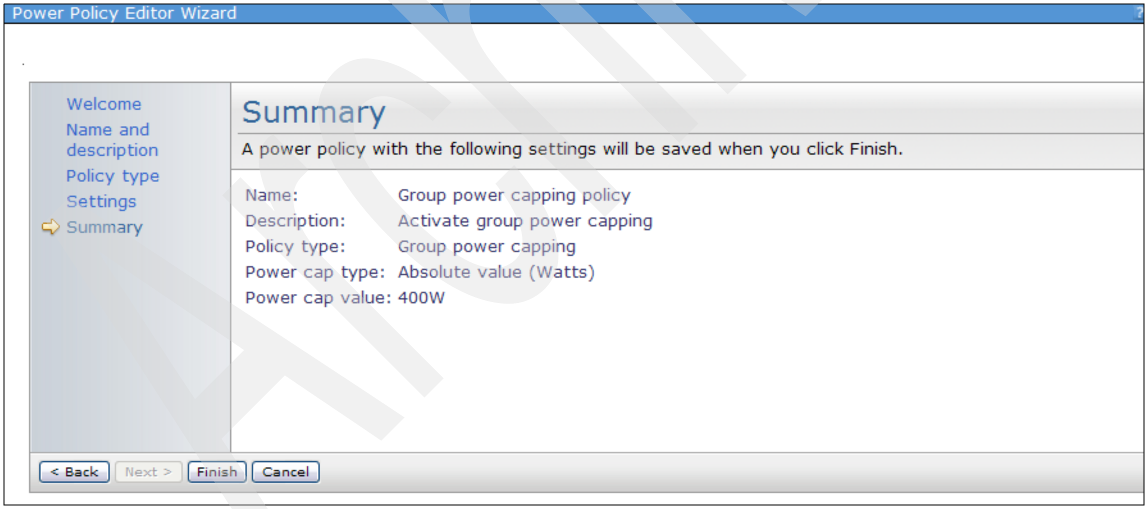


*Figure 7-50 Creating a group power capping policy - 6*

The policy is created and displays in the list of policies shown in Figure 7-51.

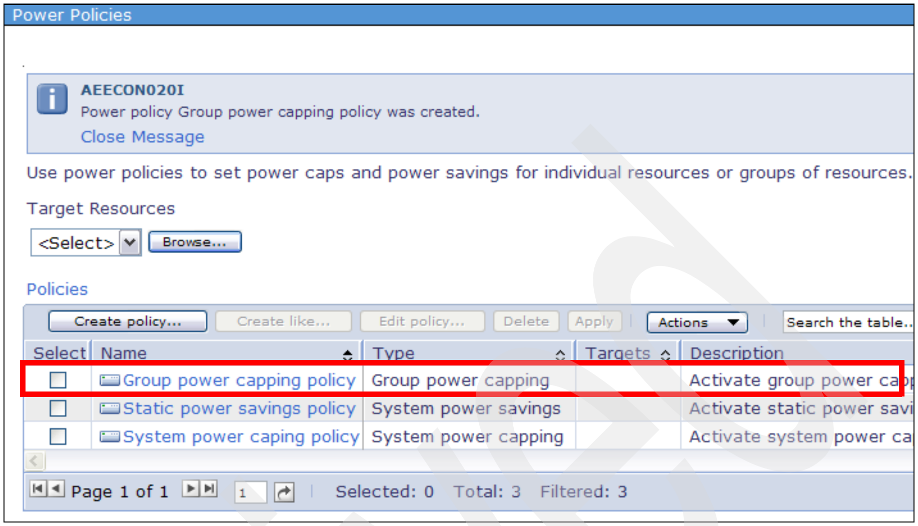


*Figure 7-51 Creating a group power capping policy - 7*

8. Clicking a policy name in the Power Policies view shows the power policy's properties (Figure 7-44 on page 336).

   The Targets tab shows all the systems that this policy is applied to (or groups if it is a group power capping policy), as shown in Figure 7-52 on page 341.

> **Note:** You use the Targets tab to remove a policy from a target system or group. Refer to 7.5.5, "Removing a policy" on page 353 for more information.

You can now apply this policy to a power managed system or group of systems as described in 7.5.3, "Applying a policy" on page 341.

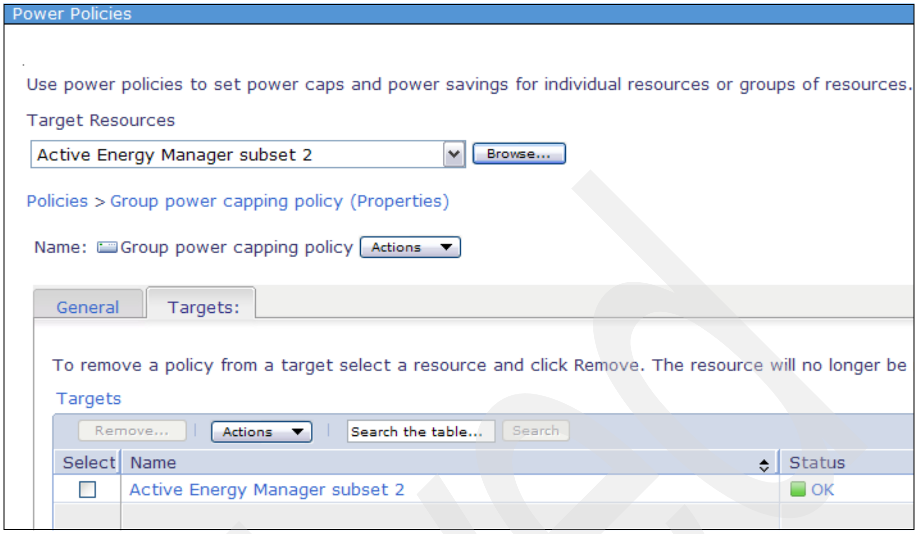


*Figure 7-52 Creating a group power capping policy - 8*

## 7.5.3 Applying a policy

Policies are stand-alone constructs. They are not effective until they are *applied* to a resource or group of resources.

When a power policy is applied to a target, Active Energy Manager makes sure that policy stays in effect until that policy is removed. This involves listening for actions or changes that could require Active Energy Manager to re-evaluate a policy and possibly change the system or group power cap or savings mode. Note that this listening and re-evaluating goes on constantly, but can only happen while the IBM Systems Director Server is running. When IBM Systems Director server is started or restarted, all policies are re-evaluated and enforced on the systems or groups to which they are applied.

Any system power savings or system power capping *policy* overrides and disables previously applied *settings* of the same type. For example, applying a static power savings policy disables a previously applied dynamic power savings (favor performance) setting.

When you apply a system power savings or a system power capping policy to a group of resources, the policy is applied to each member of the group independently. If you apply a group power capping policy to a group of resources, the power cap value you specify is spread across all members of the group. A power cap is calculated for each individual resource, and the cumulative value of all system power caps is equal to the group power cap you specify. Notice that the power cap value assigned to each individual resource can be, and probably will be, different because each value takes into account the hardware installed on the resource.

The rules for applying policies are closely related to the rules used to enforce them. Refer to “Enforcing policies” on page 324 for information about enforcing policies.

When applying policies to a BladeCenter server, note the following special considerations:

- Do not set system power capping or system power savings policies on both a slot and the blade in that slot.
- Do not set group power capping policies on two groups, where one group contains a slot and the other group contains a blade in that slot.
- A policy setting on both a slot and a blade server in that slot affect the blade server, so setting different policies on each can cause problems.

To apply a policy, follow these steps:

1. On the Active Energy Manager home page, click **Work with power policies** in the Manage section (Figure 7-53).

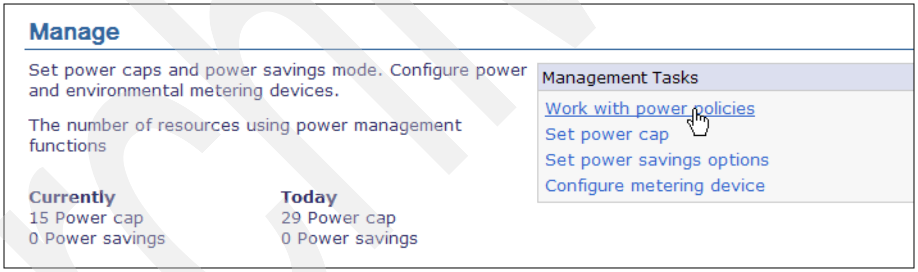


*Figure 7-53 Applying a policy - 1*

2. In the display shown in Figure 7-54, click **Browse** to select the target systems for the apply.

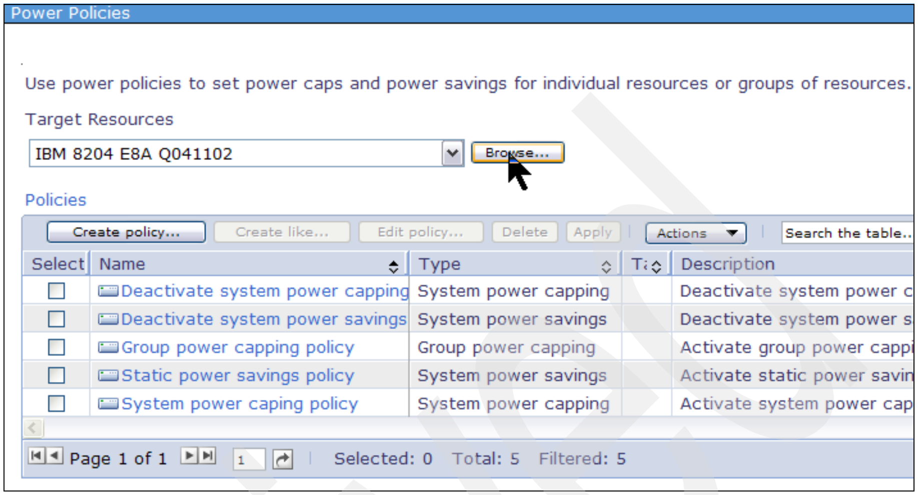


*Figure 7-54 Applying a policy - 2*

3. In the display shown in Figure 7-55 on page 344, perform these tasks:

   a. From the **Show** drop-down menu, select one of the following options:

      - **Active Energy Manager Resources** to select from a list of individual resources
      - **Groups** to select from a list of groups of resources
      - **Recent targets** to select from a list of resources that you have been working with recently

      As shown in Figure 7-55 on page 344, we select from a list of individual Active Energy Manager resources.

   b. Select the systems or groups you to which you want to apply the policy by selecting the appropriate check boxes. When finished, click **Add** to move the systems or groups to the Selected box. Then, click **OK**.

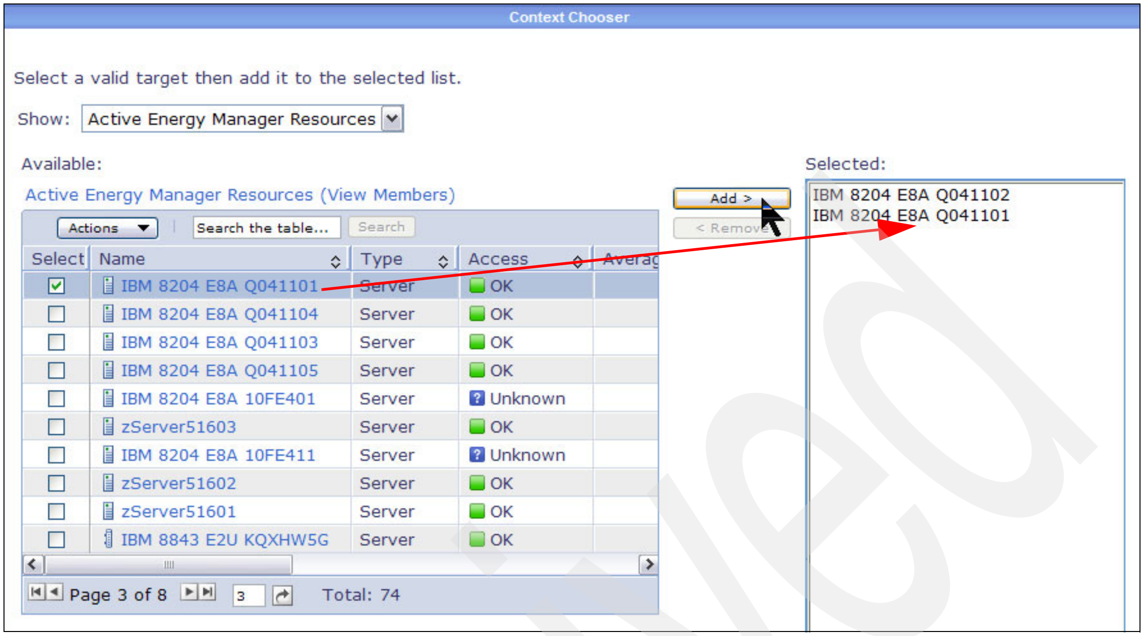


*Figure 7-55 Applying a policy - 3*

4. Select the policy that you want to apply, and then click **Apply** (Figure 7-56).

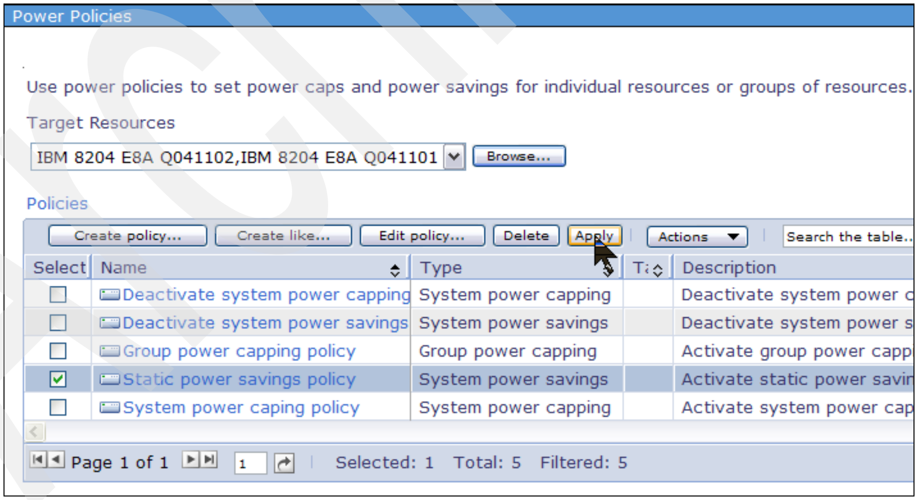


*Figure 7-56 Applying a policy - 4*

5. In the display shown in Figure 7-57, notice that Active Energy Manager inserts a date and time into the job name field. You can choose to run the apply interactively (Run Now), or Schedule the job for another time.

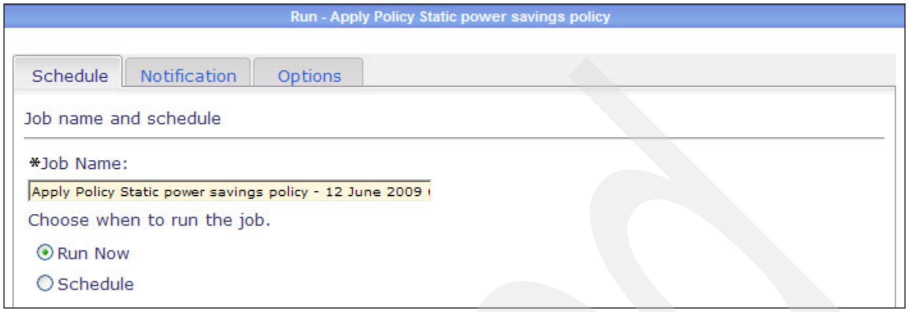


*Figure 7-57 Applying a policy - 5*

6. If you choose to schedule the apply, select **Schedule**. The display shown in Figure 7-58 opens. Select the **Time**, **Date**, and **Frequency** to run the apply, and then click **OK**.

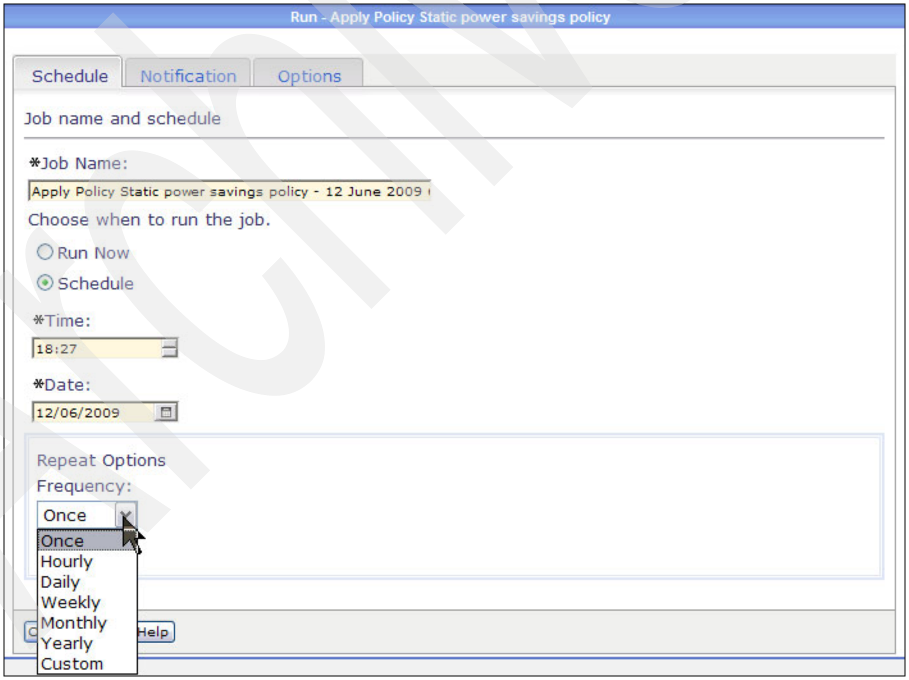


*Figure 7-58 Applying a policy - 6*

7. If you chose the Run Now option, the display shown in Figure 7-59 opens. Note the message confirming that the policy apply job has started. The systems to which the policy was applied are listed in the Targets column.

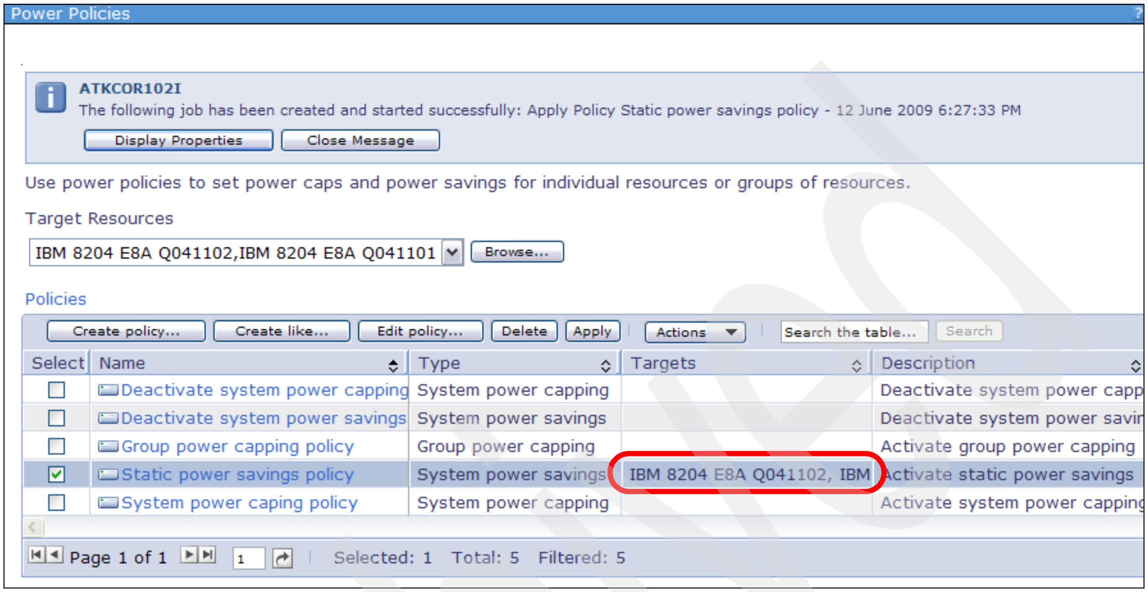


*Figure 7-59 Applying a policy - 7*

8. You can confirm that the policy is applied by checking the task status as follows:

   a. First, click **Task Management** → **Active and Scheduled Jobs** in the IBM Systems Director navigation pane. In the display shown in Figure 7-60, click the apply job that you have just run.

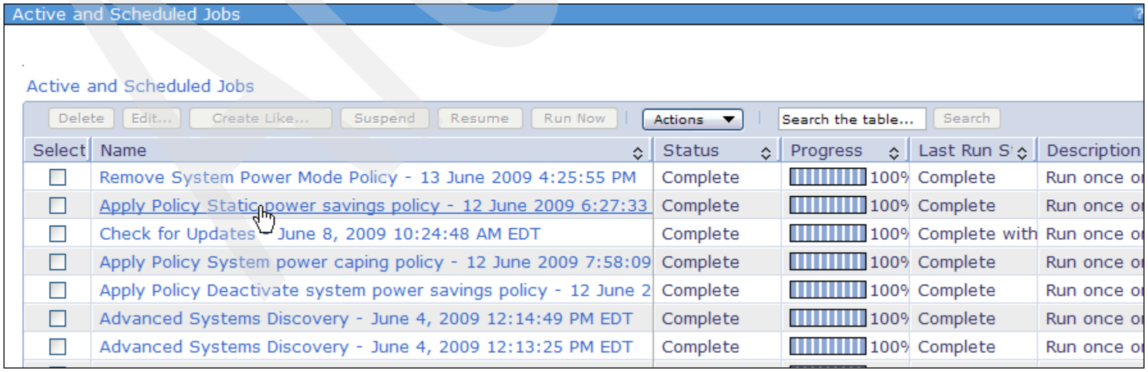


*Figure 7-60 Applying a policy - 8*

b. In the display shown in Figure 7-61, the General tab shows the overall status of the apply or remove task, including scheduling info (for instance if it is a reoccurring scheduled job).

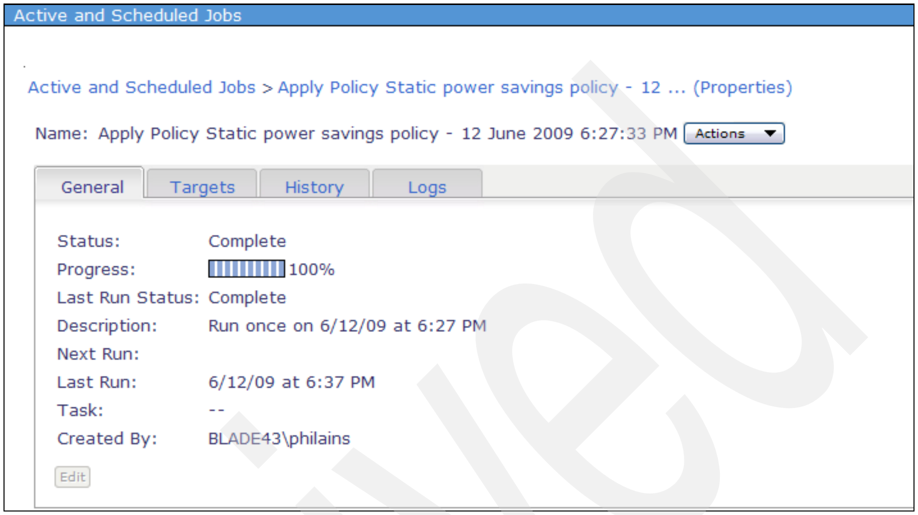


*Figure 7-61 Applying a policy - 9*

c. Click the Targets tab to confirm the list of targets for the apply (Figure 7-62).

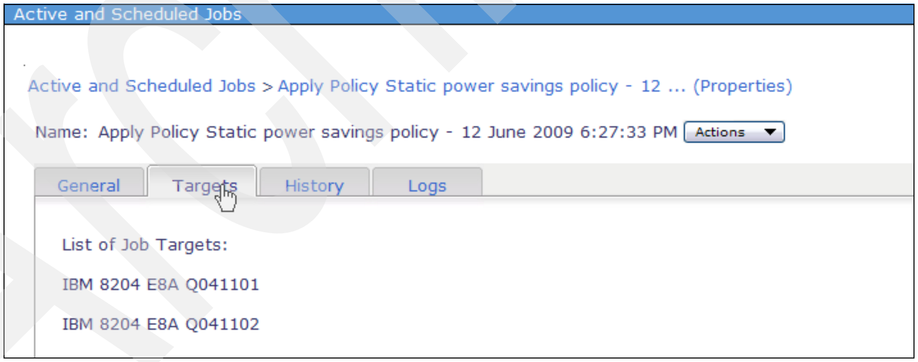


*Figure 7-62 Applying a policy - 10*

d. The History tab (Figure 7-63) has an entry for each time this job was run. Reoccurring scheduled jobs have a row for each time the job was run.

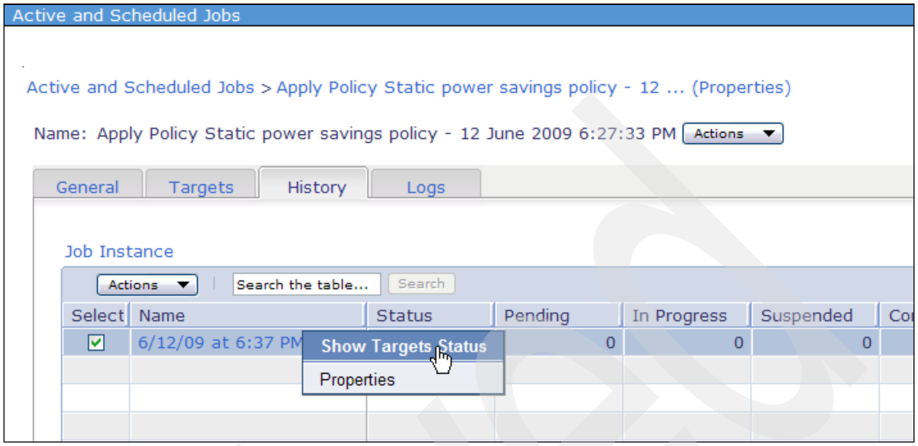


Figure 7-63 Applying a policy - 11

e. Click an entry or right-click an entry in Figure 7-63, and select **Show Targets Status** to see the status of each individual target for that particular run of the job (Figure 7-64).

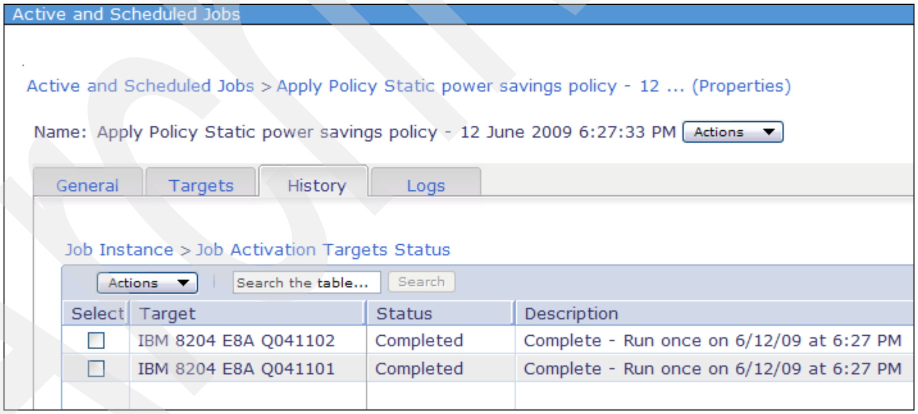


Figure 7-64 Applying a policy - 12

f. The Logs tab is similar to the History tab in that there is an entry for each time this job was run. Again, reoccurring scheduled jobs have a row for each time the job was run.

   Click an entry, or right-click an entry in the table and select **Show Logs** to see the detailed log for that particular run of this apply job (Figure 7-65). You can use this log to help resolve any errors experienced during the apply.

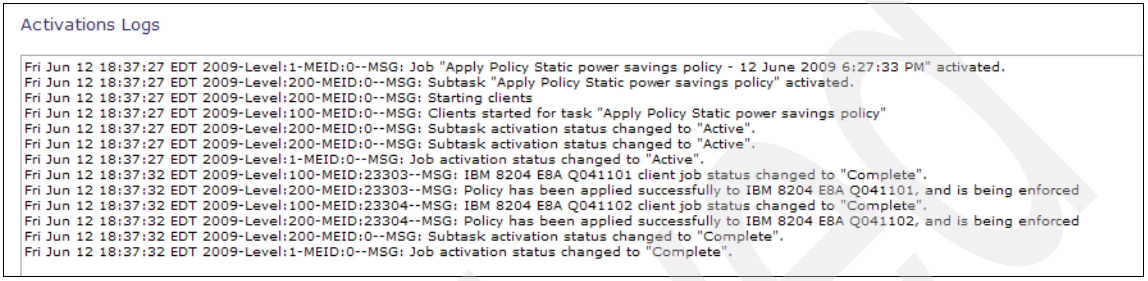
Activations Logs

```
Fri Jun 12 18:37:27 EDT 2009-Level:1-MEID:0--MSG: Job "Apply Policy Static power savings policy - 12 June 2009 6:27:33 PM" activated.
Fri Jun 12 18:37:27 EDT 2009-Level:200-MEID:0--MSG: Subtask "Apply Policy Static power savings policy" activated.
Fri Jun 12 18:37:27 EDT 2009-Level:200-MEID:0--MSG: Starting clients
Fri Jun 12 18:37:27 EDT 2009-Level:100-MEID:0--MSG: Clients started for task "Apply Policy Static power savings policy"
Fri Jun 12 18:37:27 EDT 2009-Level:200-MEID:0--MSG: Subtask activation status changed to "Active".
Fri Jun 12 18:37:27 EDT 2009-Level:200-MEID:0--MSG: Subtask activation status changed to "Active".
Fri Jun 12 18:37:27 EDT 2009-Level:1-MEID:0--MSG: Job activation status changed to "Active".
Fri Jun 12 18:37:32 EDT 2009-Level:100-MEID:23303--MSG: IBM 8204 E8A Q041101 client job status changed to "Complete".
Fri Jun 12 18:37:32 EDT 2009-Level:200-MEID:23303--MSG: Policy has been applied successfully to IBM 8204 E8A Q041101, and is being enforced
Fri Jun 12 18:37:32 EDT 2009-Level:100-MEID:23304--MSG: IBM 8204 E8A Q041102 client job status changed to "Complete".
Fri Jun 12 18:37:32 EDT 2009-Level:200-MEID:23304--MSG: Policy has been applied successfully to IBM 8204 E8A Q041102, and is being enforced
Fri Jun 12 18:37:32 EDT 2009-Level:200-MEID:0--MSG: Subtask activation status changed to "Complete".
Fri Jun 12 18:37:32 EDT 2009-Level:1-MEID:0--MSG: Job activation status changed to "Complete".
```

*Figure 7-65 Applying a policy - 13*

9. You can also confirm that the policy has been successfully applied by checking the trend data for the selected systems. An example is shown in Figure 7-66.

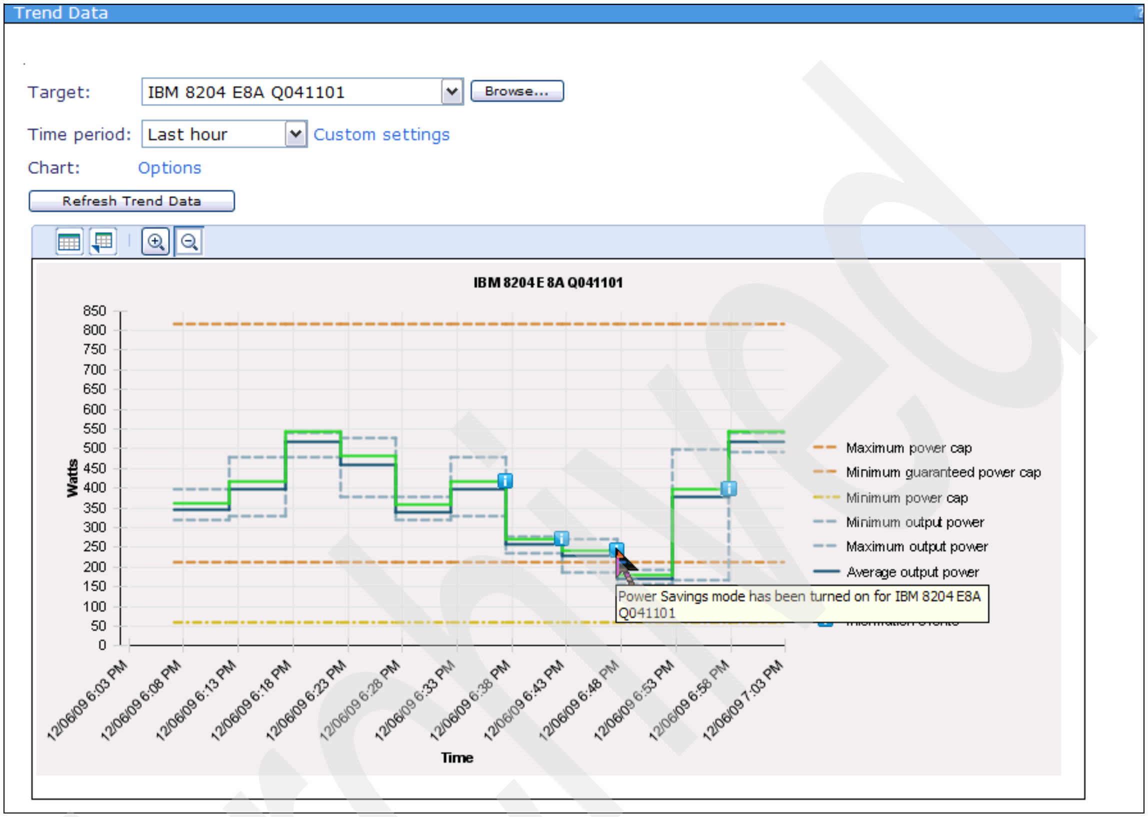


*Figure 7-66 Applying a policy - 14*

## 7.5.4 Deactivating a policy

Completely deactivating a policy is a two-step process:

- *Remove the policy* from individual resources or groups which are subject to the policy.
- *Change the settings* that were set by the policy back to their previous or default values.

In this section, we describe the two methods that you can use to complete these steps.

## Method 1

You can use method 1 to deactivate any type of policy on systems and groups. This method is the most straightforward, but if there are a large number of systems subject to the that policy you want to remove, it might take a long time. In this case, you might want to consider method 2.

The steps involved in using method 1 are:

1. Remove systems or groups from the policy.

   Refer to 7.5.5, "Removing a policy" on page 353.

2. Change the setting on each individual system or group to which the policy has been applied back to the default.

   Refer to 7.2, "Setting a power savings mode" on page 287 or 7.3, "Setting a power cap" on page 306.

## Method 2

This method involves the creation of a policy to deactivate another policy. You might want to use this method if you have a large number of systems subject to a policy you want to deactivate. Otherwise, you must change the policy setting on each individual system or group to which the deactivated policy was applied, after you have removed the policy (method 1).

> **Note:** You cannot use this method to deactivate a group power capping policy because it is not possible to create a policy to deactivate power capping on a group of systems that have been subject to a group power capping policy.

To deactivate a system power savings or system power capping policy by applying another policy, follow these steps:

1. Create a policy to *deactivate* system power savings or system power capping. Specify **No power savings** or **Deactivate Power Capping** on the Settings panel as the case might be.

   Refer to 7.5.2, "Creating a policy" on page 327.

2. Apply the deactivation policy to those systems or groups to which the system power savings or system power capping policy has previously been applied.

   Refer to 7.5.3, "Applying a policy" on page 341.

   > **Note:** The deactivation policy is a valid policy so Active Energy Manager continues to enforce it. You need to keep this in mind when using this method.

3. You can confirm that the initial policy has been deactivated by checking the trend data for the selected systems. An example is shown in Figure 7-67. It shows that the example policy that we previously activated in 7.5.3, “Applying a policy” on page 341 is now deactivated using a deactivation policy.

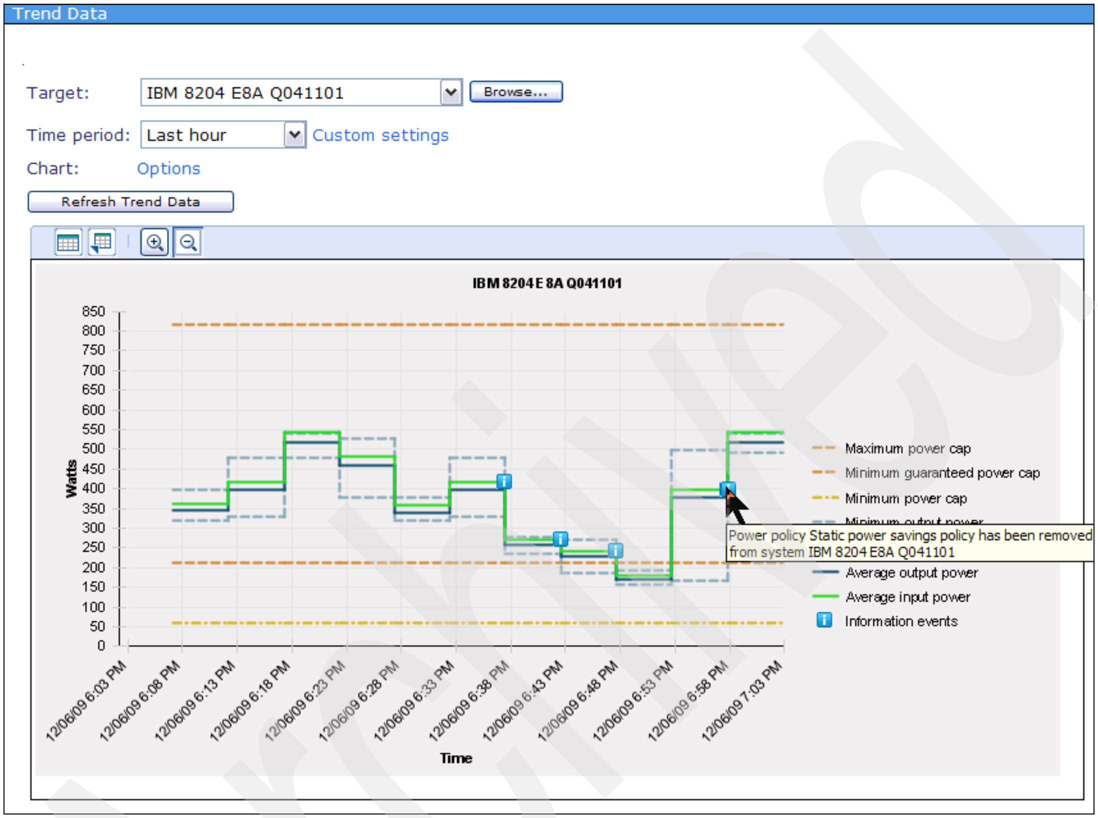


*Figure 7-67   Using a deactivation policy - 1*

4. You can also confirm the deactivation of the initial policy by navigating to the Active Energy Manager (View Resources) display, right-clicking one of the systems that have had the policy deactivated, and selecting **Energy** →

**Manage Power** → **Power Savings**. In the display shown in Figure 7-68, notice that, although power savings is now turned off, the system is still under a policy (the deactivation policy).

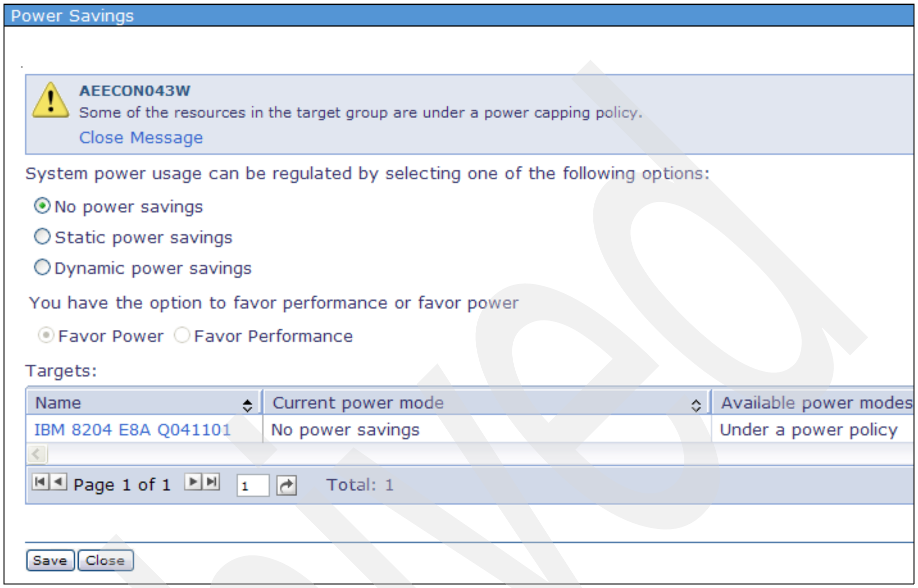


*Figure 7-68 Using a deactivation policy - 2*

## 7.5.5 Removing a policy

You can remove a policy from an individual system or group of systems. When you remove a policy from a system or group, the policy settings are no longer continuously enforced on that system or group. However, the policy is *not* deleted and is available to be applied to other targets in the future.

> **Note:** It is important to understand that the policy settings applied by a removed policy remain on the systems and groups which were subject to the policy, even after the policy has been removed.

Removing a policy simply allows you to change manually the power savings mode and power cap setting on a removed system or group without the policy settings being reimposed. For example, removing a system or group capping policy does not turn power capping off or return the power cap value to previous state. Likewise, removing a power savings policy does not disable the power savings mode in operation.

To change the power savings mode and power cap settings, you can use the Active Energy Manager functions described in 7.2, “Setting a power savings mode” on page 287 and 7.3, “Setting a power cap” on page 306.

To remove a policy, follow these steps:

1. From the Active Energy Manager home page, click **Work with power policies** in the **Manage** section.
2. In the display shown in Figure 7-69, click the policy that you want to remove from the target systems or groups.

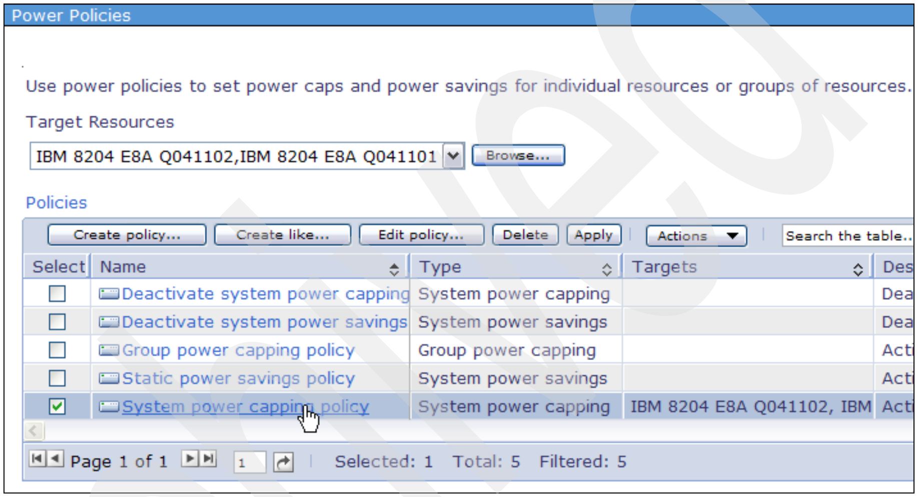


*Figure 7-69 Removing a policy - 1*

3. You see the properties of the policy (Figure 7-70 on page 355). Proceed as follows:

   a. Click the Targets tab.

   b. Click the check boxes next to the targets that you want to remove from the control of this policy.

   c. Click **Actions**.

   d. Select **Remove**.

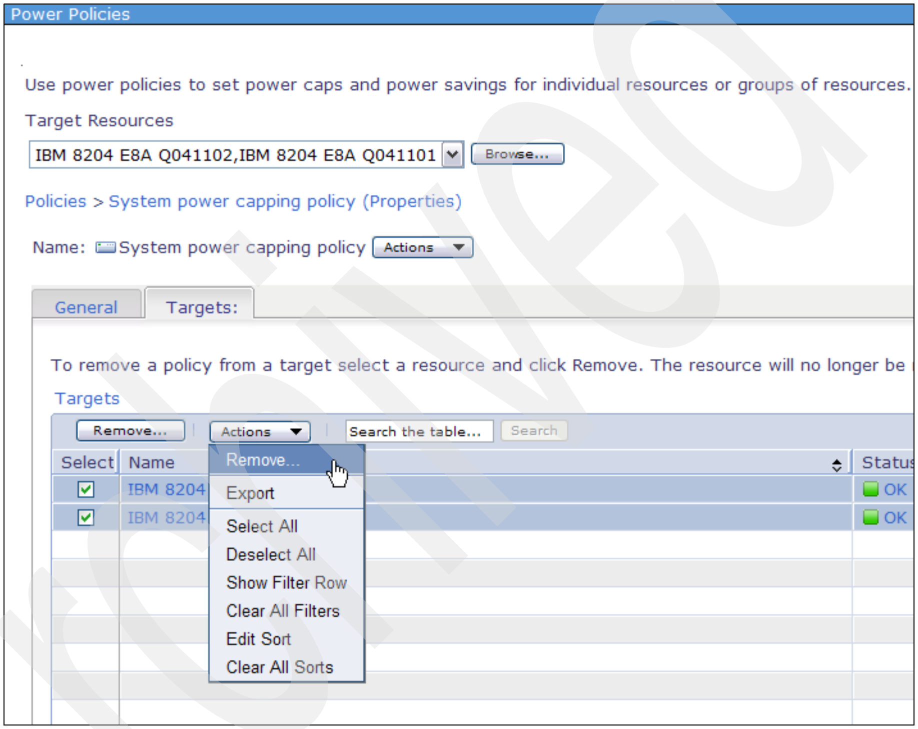


*Figure 7-70 Removing a policy - 2*

4. Click **OK** on the confirmation window.

5. You are prompted to either run the job interactively, or specify a schedule. Make your choice, and then click **OK**.

6. You see a confirmation message, and the target systems or groups are removed from the list on the Targets tab (Figure 7-71).

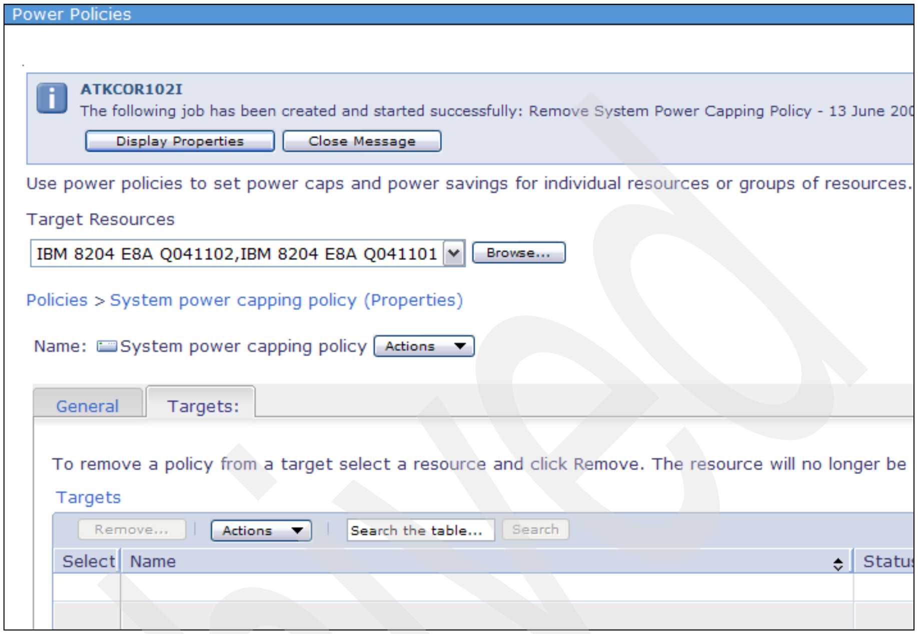


*Figure 7-71 Removing a policy - 3*

7. You can confirm that the policy was removed by checking the task status as follows:

   a. First, click **Task Management** → **Active and Scheduled Jobs** in the IBM Systems Director navigation pane. You see the display shown in Figure 7-72. Click the remove job that you have just run.

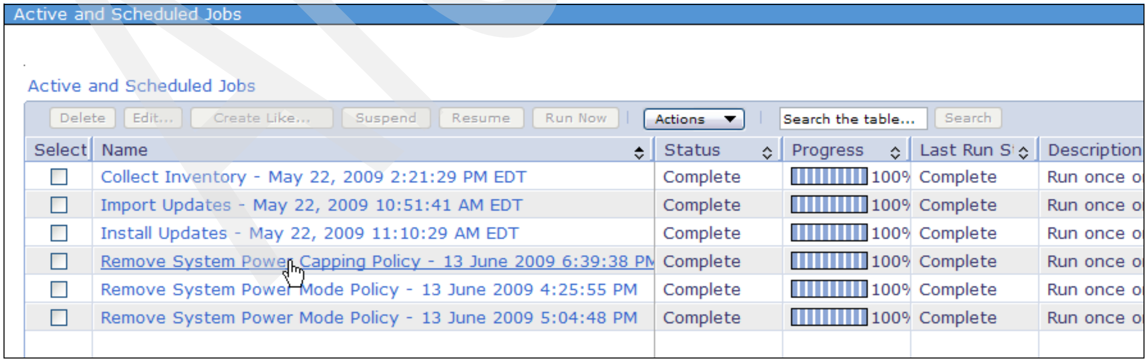


*Figure 7-72 Removing a policy - 4*

b. You see the display shown in Figure 7-73. You can browse the four tabs just as you would for a policy apply.

Two tabs of interest here are the Targets tab and the Logs tab. First click the Targets tab to confirm the list of targets that have been removed.


*Figure 7-73 Removing a policy - 5*

c. Second, click the Logs tab to view the message log for the remove (Figure 7-74). You can use this to help resolve any errors experienced during the remove.

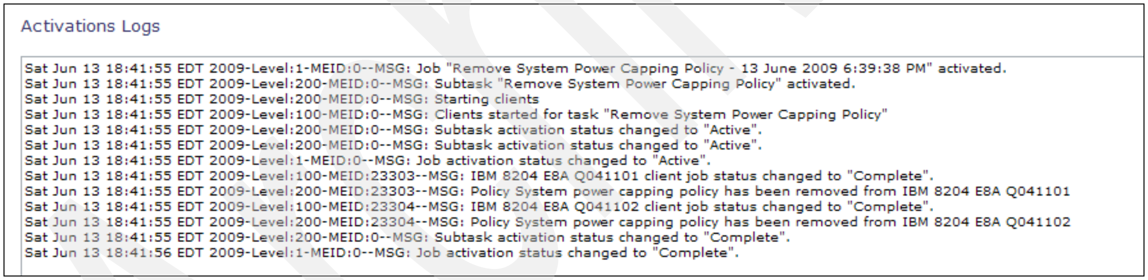


*Figure 7-74 Removing a policy - 6*

8. If you look at the trend data for the systems that have had the policy removed you notice that, although the policy has been removed, the policy setting is still active. An example is shown in Figure 7-75 for one of the systems that we removed the policy from in the preceding steps. The power cap is still set to the policy setting of 400 watts.

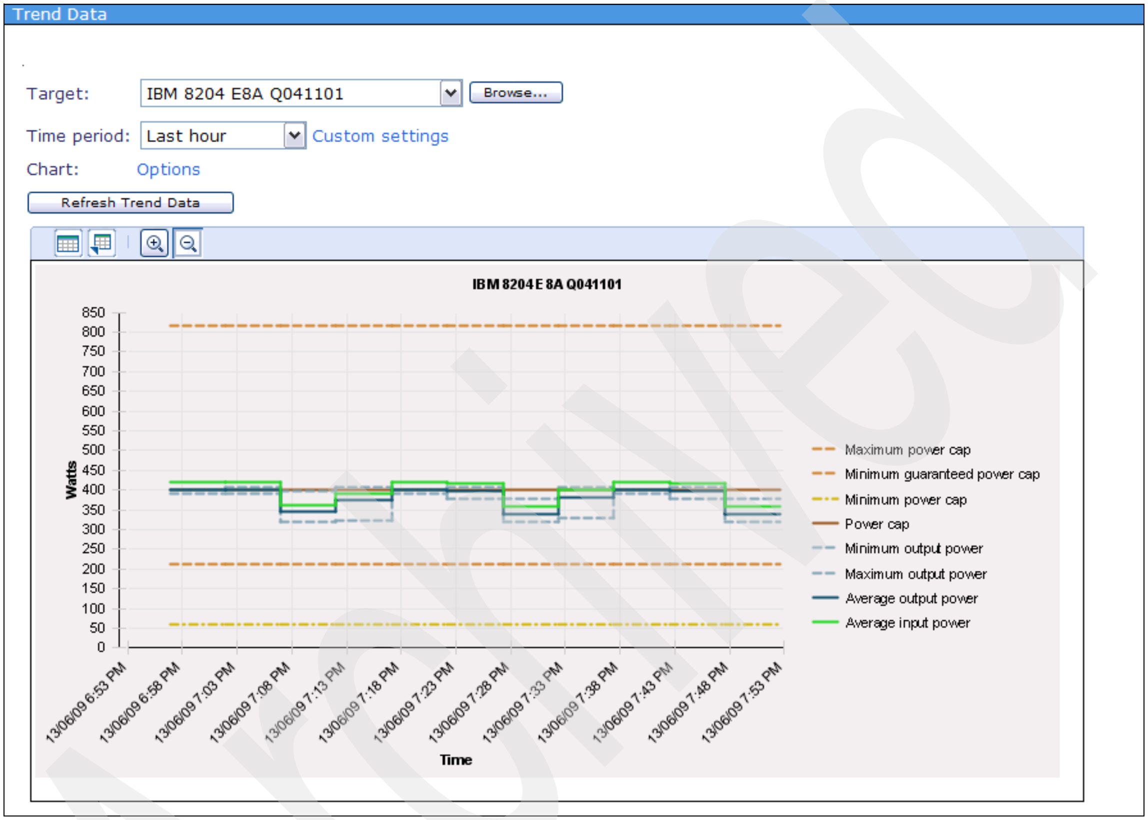


*Figure 7-75 Removing a policy - 7*

9. Now that the policy has been removed we can interactively deactivate the power cap setting as follows:
   a. Navigate to Active Energy Manager Resources (View Members).
   b. Right-click the resource you want to change the setting for.
   c. Select **Energy** → **Manage Power** → **Power Capping** or **Power Savings**.
   d. Deactivate the setting as appropriate.
10. Confirm that the power savings or power capping setting has been disabled by checking the trend data. The trend data for our example system is shown in Figure 7-76 on page 360.

    Note that it also possible to remove a policy by applying another policy. For example, you could effectively remove a power capping policy by applying a power capping policy to deactivate power capping. Refer to 7.5.4, "Deactivating a policy" on page 350 for more details.

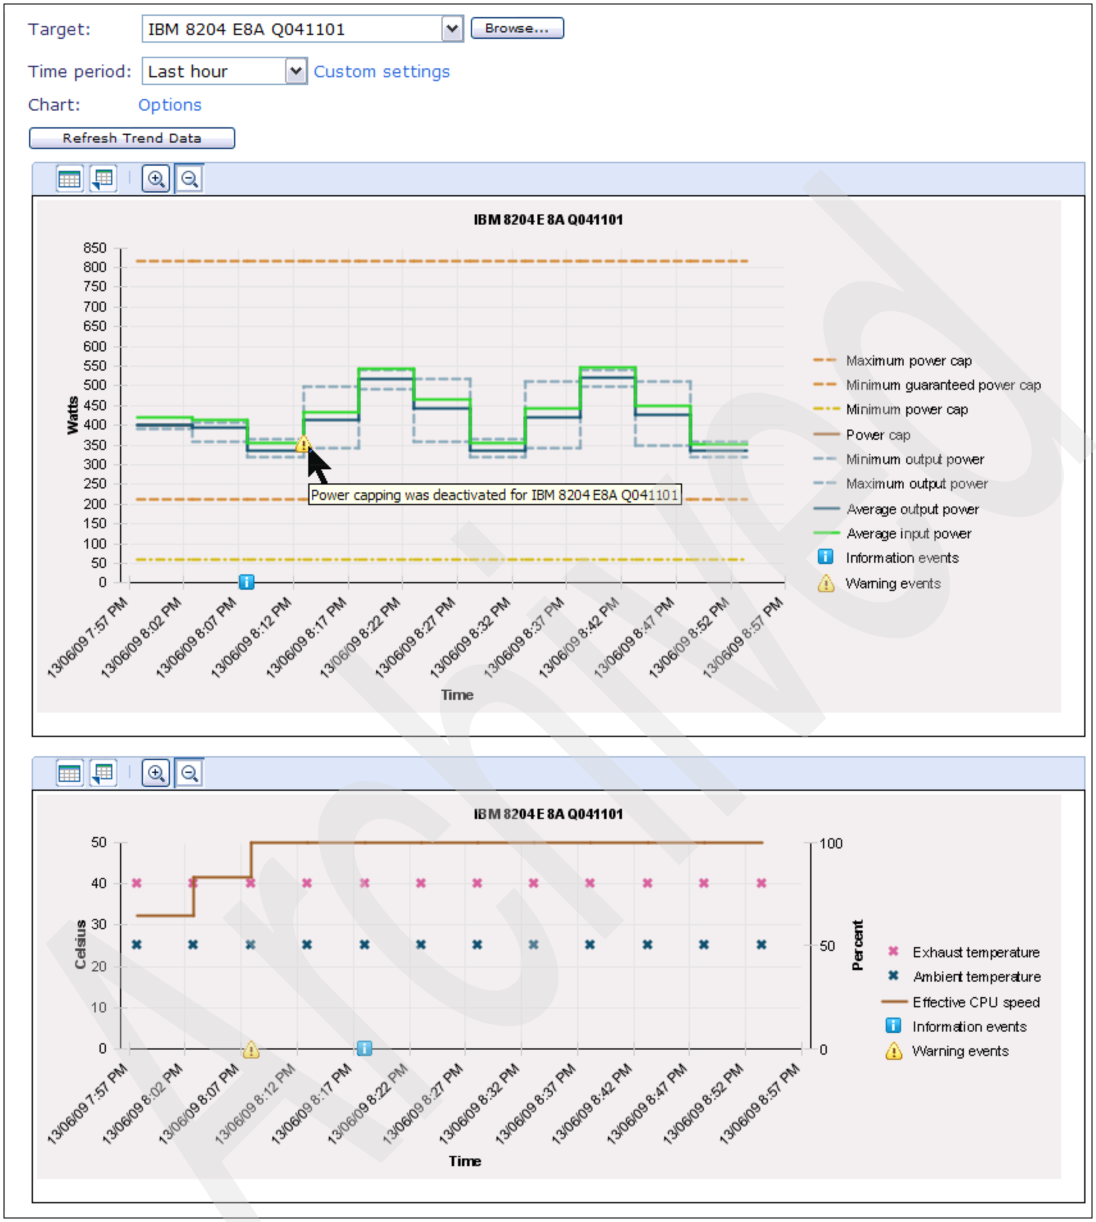


*Figure 7-76 Removing a policy - 8*

## 7.5.6 Editing a policy

You can edit a power policy using the Power Policy Editor wizard.

To edit a policy, follow these steps:

1. From the Active Energy Manager home page, click **Work with power policies** in the Manage section.
2. In the display shown in Figure 7-77, select the policy that you want to edit, and click **Edit policy** to start the Power Policy Editor Wizard.

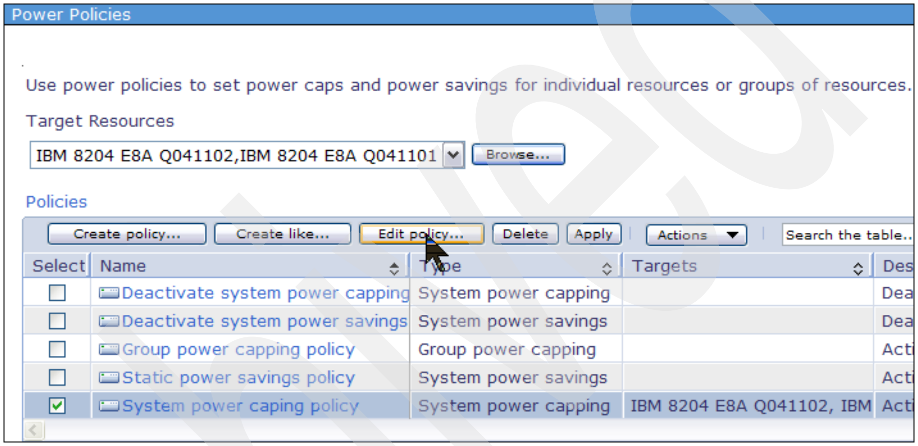


*Figure 7-77 Editing a policy*

3. Use the Power Policy Editor Wizard to make your desired changes. Notice that all policy settings except the policy type can be changed. If you need to change the policy type, you must create a new policy.
4. When you have finished, click **Finish**.

   If the policy you are editing is currently applied to targets, then the new, edited policy is re-enforced on all targets.

## 7.5.7 Deleting a policy

Deleting a policy has the effect of removing the policy from the systems and groups to which it was applied. However, if you simply want to deactivate a policy rather than delete it, a better option would be to *remove* the policy from the systems and groups to which it has been applied. Refer to 7.5.4, “Deactivating a policy” on page 350 and 7.5.5, “Removing a policy” on page 353.

If you really want to delete a policy, follow these steps:

1. From the Active Energy Manager home page, click **Work with power policies** in the Manage section.
2. You see the display shown in Figure 7-78. Select the policy that you want to delete, and then click **Delete**.

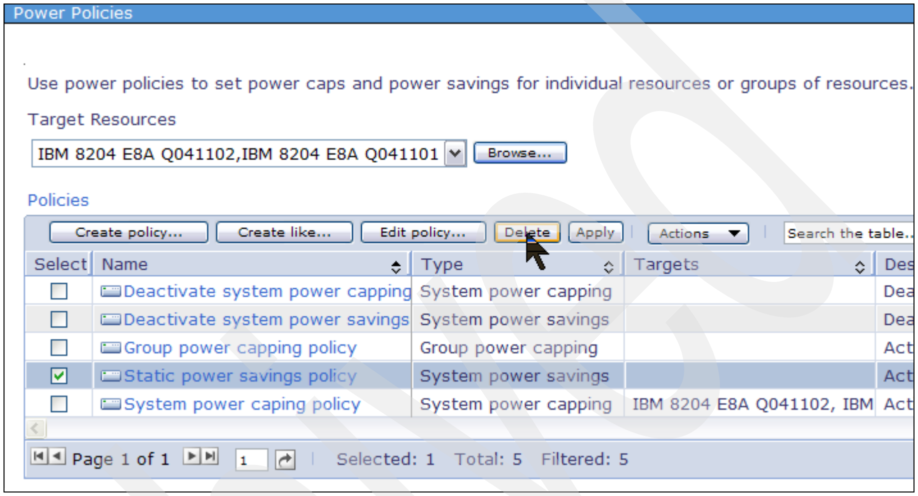


*Figure 7-78 Deleting a policy*

3. Click **OK** on the confirmation window.

   First the policy is removed from any targets to which it is applied, and then the policy is deleted and removed from the Power Policies window.

## 7.6 Using the command-line interface

IBM Systems Director includes a library of commands that you can use to configure or perform many of the systems-management operations that can be accomplished from IBM Systems Director Web console.

For a complete list of IBM Systems Director commands, refer to:

`http://publib.boulder.ibm.com/infocenter/director/v6r1x/index.jsp?topic=/director.cli_6.1/fqm0_r_cli_a_to_z.html`

Active Energy Manager commands are invoked through the standard IBM Systems Director systems management command-line interface (**`smcli`**) mechanism.

For detailed information about Active Energy Manager commands, refer to:

http://publib.boulder.ibm.com/infocenter/director/v6r1x/index.jsp?topic=/aem_410/frb0_r_ref_commands_preferences.html

All Active Energy Manager commands support two general types of output: comma-separated values (CSV) file and verbose. CSV file output is used for importing data directly into another database application or to be viewed with a spreadsheet application. The CSV output column headers are not translated whereas verbose output is translated. Error messages are sent to the standard error output device and all commands return a status code.

## 7.6.1 Active Energy Manager CLI commands description

Table 7-3 shows the general information commands. These commands are used to display and change general information about Active Energy Manager resources.

*Table 7-3 General information commands*

| Command | Description |
|---|---|
| lspowerobjects | List the Active Energy Manager resources |
| lspowerinfo | List the power information for the specified Active Energy Manager resources |
| chpowerinfo | Update power information for the specified Active Energy Manager resources |

Table 7-4 shows global preferences commands. These commands are used to display and change global Active Energy Manager properties.

*Table 7-4 Global preferences commands*

| Command | Description |
|---|---|
| lsproperties | Display the global Active Energy Manager properties |
| chproperties | Change the global Active Energy Manager properties |

Table 7-5 shows the metering commands. These commands are used to manage metering for, and query metered data from Active Energy Manager resources.

*Table 7-5 Metering Commands*

| Command | Description |
|---|---|
| lscollect | List the status of data collection for the specified objects |
| startcollect | Start data collection for the specified objects |
| stopcollect | Stop data collection for the specified objects |
| lspowerlast | Display the summary average watts for Active Energy Manager resources |
| lstrenddata | List the metering information for the specified Active Energy Manager resources |

Table 7-6 shows power management commands. These commands are used to manage the power settings for Active Energy Manager resources.

*Table 7-6 Power Management Commands*

| Command | Description |
|---|---|
| getpcap | Display the current power cap values for the specified Active Energy Manager resources |
| setpcap | Set the current Pcap values for the specified Active Energy Manager resources |
| getpsaver | Display the current power savings settings for the specified Active Energy Manager resources |
| setpsaver | Set power savings for the selected Active Energy Manager resources |
| lspolicy | List policy attributes |
| chpolicy | Create, modify and delete policies |
| setpolicy | Activate a policy for the specified Active Energy Manager resources |

### 7.6.2 Active Energy Manager CLI command examples

> **Note:** System names might not be unique. Certain commands act on all systems with the specified name. To target a particular system that has a name that is not unique, identify the system by specifying its unique hexadecimal ID. (Use the **`smcli lspowerobject`** command to list them all.)

The following list includes Active Energy Manager CLI command examples:

- Display the global properties:

  `smcli lsproperties`

- List objects for a particular system type named *Powerunit*:

  `smcli lspowerobjects -t "PowerUnit"`

- List the collection status for all Active Energy Manager resources:

  `smcli lscollect -a`

- List the collection status for a system with the name Machine:

  `smcli lscollect -n Machine`

- Get the input power for the specified resource at the specified interval:

  `smcli lstrenddata -n 0x171 -r "9/30/08 10:10:49 AM,9/30/08 10:20:49 AM"`

- Get all input power for the specified resource until the specified time:

  `smcli lstrenddata -n 0x171 -r "9/30/08 10:20:49 AM"`

- Display the average watts for a system with the name *Machine*:

  `smcli lspowerlast -n Machine`

- Get the input power for the specified resource from the last hour:

  `smcli lstrenddata -n 0x171`

- Get the Pcap for a system with the name *Machine*:

  `smcli getpcap -n Machine`

- Get the power savings for a system with the name *Machine*:

  `smcli getpsaver -n Machine`

- Get the information for a policy with the name *mypolicy*:

  `smcli lspolicy -p mypolicy`

- Change the global property named *meteringInterval*:

  `smcli chproperties -p meteringInterval=5`

- Set the static power savings mode for a system with the name *Machine*:

```
smcli setpsaver -n Machine -p static
```

- Set dynamic power savings (favor performance) mode off for a system with the name *Machine*:

```
smcli setpsaver -n Machine -p dynamic,favorperformance=off
```

- Set the power information for the resource specified:

```
smcli chpowerinfo -p altitude=600
```

- Set the Pcap for a system with the name *Machine*:

```
smcli setpcap -n Machine -p 165
```

- Start collecting data for a system with the name *Machine*:

```
smcli startcollect -n Machine
```

- Stop collecting data for a system with the name *Machine*:

```
smcli stopcollect -n Machine
```

- Create a policy named *myPolicy*:

```
smcli chpolicy -p myPolicy,policytype=capping,pcap=200
```

- Activate a policy named *myPolicy*:

```
smcli setpolicy -p myPolicy -n 0x0ad1
```

- Update a policy named *myPolicy*:

```
smcli chpolicy -p myPolicy,pcap=350
```

- Remove a policy named *myPolicy*:

```
smcli chpolicy -r -p myPolicy
```

It is possible to combine these commands with shell scripts from Linux or AIX. For example to display all systems with Pcap activated, and their values, use this command:

```
for j in $(/DirRoot/bin/smcli lscollect -a |cut -d, -f2);do smcli
getpcap -n $j|grep -i true;done
```

In Windows, you can use a batch file or script. For example, to list the collection status every 5 seconds, create a `.vbs` file and include the following lines:

```
set wshshell = wscript.CreateObject("WScript.Shell")
while true
WshShell.Run "c:\DirRoot\bin\smcli lscollect -a"
wscript.sleep 5000
wend
```

Run the script by typing `cscript file.vbs` from the command line.

8

# Using a PDU with Active Energy Manager

You can use an intelligent power distribution unit (IBM brands its intelligent PDUs as *PDU+*) to monitor power data of systems that cannot connect to Active Energy Manager through a service processor (such as a BMC, IMM, or HMC). Using a PDU+ with Active Energy Manager allows you to broadly monitor systems in the data center. After the PDU+ is discovered by Active Energy Manager, you can monitor power data from each PDU outlet group and the entire PDU.

> **Note:** This chapter primarily describes the PDU+, the intelligent power distribution unit from IBM. Active Energy Manager also supports iPDUs from other vendors, such as Raritan.

In this chapter, we discuss the following topics:

# 8.1 Understanding the concept of using a PDU+

In this section, we explore the concept of using an IBM PDU+ with Active Energy Manager.

## 8.1.1 An overview of the PDU+

A PDU+ performs the same function as a standard PDU in that it controls power distribution and provides circuit protection to rack mounted devices. In addition, a PDU+ has extra intelligence that enables it to act as an Active Energy Manager endpoint.

A PDU+ can report power and thermal data to Active Energy Manager for existing systems that do not have these capabilities built in. In other words, a PDU+ provides some Active Energy Manager functionality for devices that otherwise cannot be managed by Active Energy Manager.

The PDU+ provides the following Active Energy Manager functions:

- Monitor power data
- Monitor thermal data

You can connect the following devices to a PDU+ and then use the PDU+ to report information to Active Energy Manager:

- Existing systems
- Storage
- I/O drawers
- Non-IBM systems
- Older IBM systems

PDU+ devices provide data to the IBM Systems Director server using built-in Simple Network Management Protocol (SNMP) support. SNMP is an open-industry standard for monitoring and managing network attached devices.

The way a PDU+ manages power depends on whether the connected devices support Active Energy Manager functionality:

- Devices that do not support Active Energy Manager functionality

  The PDU+ reports power usage at the outlet group level, not by outlet. Therefore, we recommend that you connect only one device per outlet group. Otherwise, you cannot distinguish the power consumption of one device that is connected to a outlet group compared with another device that is connected to the same outlet group. In this case, you need to consider which devices to connect to each power outlet of the PDU+ to obtain meaningful data.

- Devices that support Active Energy Manager functionality

  Power can be monitored at the outlet level. A child node for each outlet, rather than for each outlet group, is present under the PDU+ node in the Active Energy Manager console display.

Figure 8-1 shows an example of a rack with four managed servers. The power for each server is supplied by a corresponding PDU+ that is installed in the sides of the rack. Each of the power supplies in each of the servers is connected to a separate PDU+ outlet group. This configuration is recommended but is not mandatory.

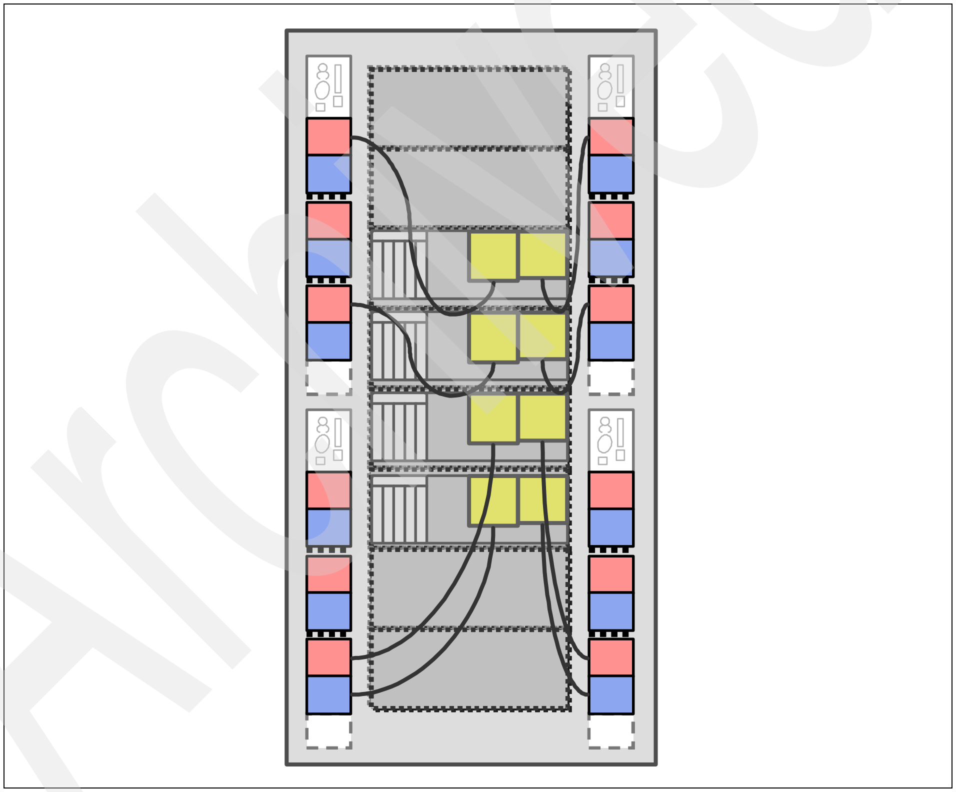

*Figure 8-1 A rack with four system drawers and four PDU+ supplying power*

## 8.1.2 Supported PDU+ models

Support for the PDU+ is new with Active Energy Manager 4.1.1. At the time of writing, the following PDU+ models are supported:

- IBM DPI C13 PDU+ (IBM part numbers 39M2816, 39M2818)
- IBM DPI C13 3-phase PDU+ (IBM part numbers 39M2817, 39M2819, 43V6045, 43V5994, 44V3897)
- IBM Ultra Density Enterprise PDU C19 PDU+ (part number: 71762MX | model: 43V5967)
- IBM Ultra Density Enterprise PDU C19 3 phase 60A PDU+ (part number: 71763MU | model: 43V5968)

You can find support information about these PDU+ models at:

http://www.ibm.com/systems/support/

For detailed information about these PDU+ models, refer to the installation and maintenance guide, which is available at:

ftp://ftp.software.ibm.com/systems/support/system_x_pdf/43v6030.pdf

### IBM DPI C13 PDU+ and IBM DPI C13 3-phase PDU+

The IBM DPI C13 PDU+ and IBM DPI C13 3-phase PDU+ models have 12 power outlets for connecting devices. Figure 8-2 shows the components and controls on the front of the DPI C13 PDU+ and the DPI C13 3-phase PDU+.

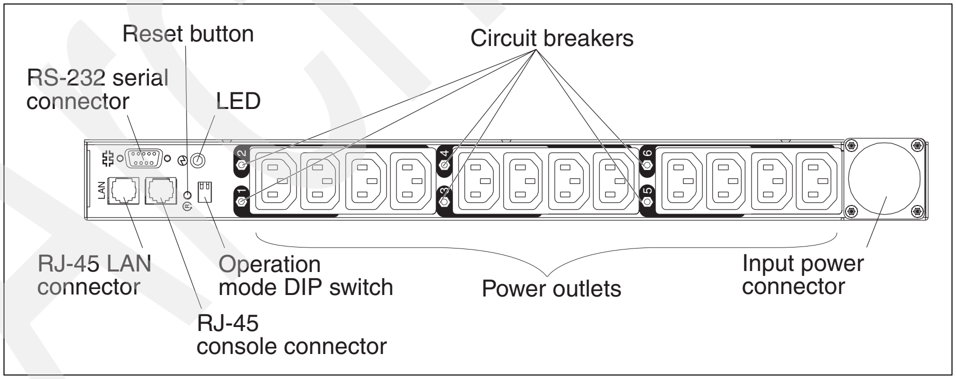


*Figure 8-2 Front view of DPI C13 PDU+ and the DPI C13 3-phase PDU+*

### IBM DPI C19 PDU+ and IBM DPI C19 3-phase PDU+

The IBM DPI C19 PDU+ and IBM DPI C19 3-phase PDU+ models have nine power outlets for connecting devices such as workstations, servers, and printers. Figure 8-3 shows the components and controls on the front of the DPI C19 PDU+ and the DPI C19 3-phase PDU+.

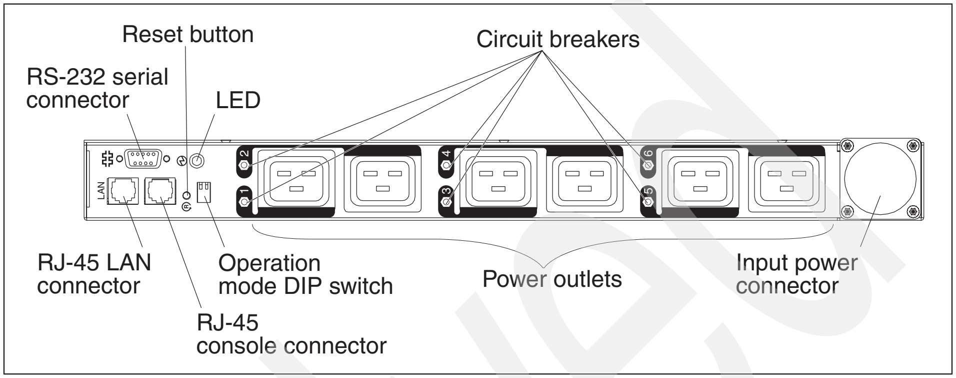


*Figure 8-3 Front view of DPI C19 PDU+ and the DPI C19 3-phase PDU+*

Figure 8-4 shows the power outlets on the rear of the DPI C19 PDU+ and the DPI C19 3-Phase PDU+.

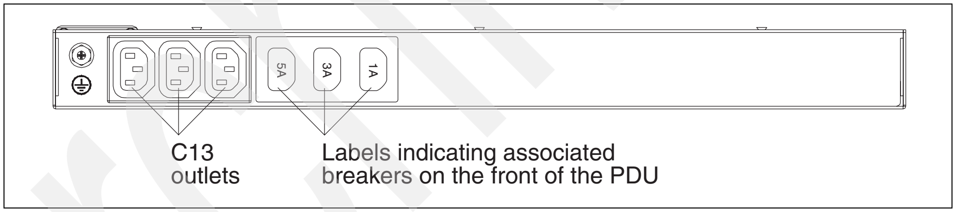


*Figure 8-4 Rear view of DPI C19 PDU+ and the DPI C19 3-phase PDU+*

For a description of the terms used in these figures, refer to the installation and maintenance guide, which is available at:

ftp://ftp.software.ibm.com/systems/support/system_x_pdf/43v6030.pdf

## 8.1.3 PDU+ outlet groups

PDU+ power outlets are combined together into *outlet groups* (also known as *load groups*). Each outlet group is protected by one of the PDU+ circuit breakers. If the power that is drawn by the outlets associated with a outlet group exceeds the rating of the circuit breaker, the breaker trips, and power is cut to the outlet group.

Table 8-1 shows the power outlets associated with each outlet group of the DPI C13 PDU+ and DPI C13 3-phase PDU+, as shown in Figure 8-2 on page 370.

*Table 8-1 Outlet groups of the DPI C13 PDU+ and DPI C13 3-phase PDU+*

| Circuit breaker (outlet group) number | Associated front power outlets |
|---|---|
| 1 | 1 and 2 |
| 2 | 3 and 4 |
| 3 | 5 and 6 |
| 4 | 7 and 8 |
| 5 | 9 and 10 |
| 6 | 11 and 12 |

**Note:** C13 PDU+ monitors power consumption from each outlet group, not from each outlet. If you want to recognize power consumption by each systems, connect system power cords into different outlet group separately.

Figure 8-5 shows how the outlet groups are displayed in the Active Energy Manager console.

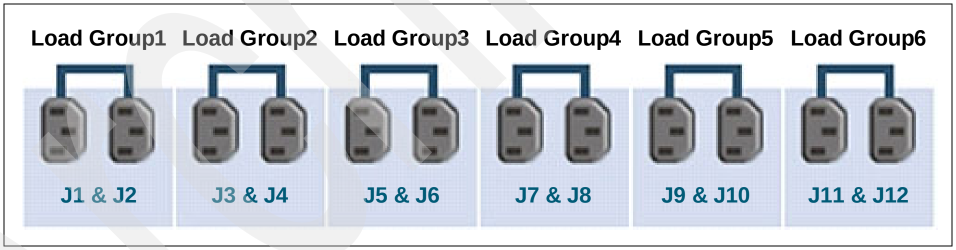


*Figure 8-5 Outlet group graphical view*

Table 8-2 shows the power outlets that are associated with each outlet group of the DPI C19 PDU+ and DPI C19 3-phase PDU+, as shown in Figure 8-3 on page 371 and Figure 8-4 on page 371.

*Table 8-2 Outlet groups of the DPI C19 PDU+ and DPI C19 3-phase PDU+*

| Circuit breaker (outlet group) number | Associated front power outlet | Associated rear power outlet |
|---|---|---|
| 1 | 1 | 1A |
| 2 | 2 | Not applicable |
| 3 | 3 | 3A |
| 4 | 4 | Not applicable |
| 5 | 5 | 5A |
| 6 | 6 | Not applicable |

## 8.2 Setting up the PDU+

In this section we describe the tasks necessary to set up the PDU+ for first-time use.

## 8.2.1 Connecting the PDU+ to a LAN

By connecting the PDU+ to a local area network (LAN), you can monitor the PDU power outlets and digital outputs over a network through the PDU+ Web interface. You connect a router or switch to the RJ-45 LAN connector on the PDU+ using an Ethernet cable, as shown in Figure 8-6. You can then monitor the PDU+ from a computer that is connected to the same TCP/IP network.

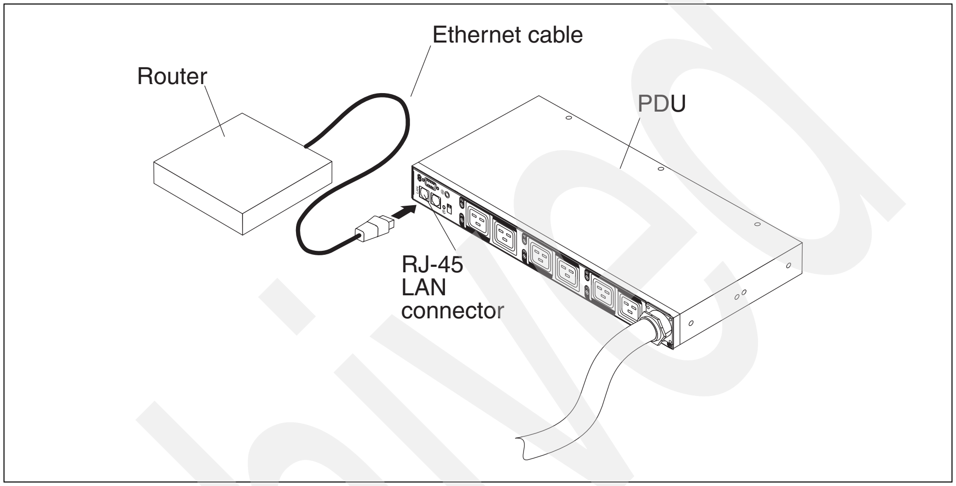


*Figure 8-6 Connecting the PDU+ to a LAN*

## 8.2.2 Configuring the IP address

To communicate with the PDU+, you must first determine the device's default IP address. After you determine the default address, you can log in using the PDU+ Web interface and change the IP address to one suitable for your network.

To configure the PDU+ IP address follow these steps:

1. First, connect the console. Use the DB9-to-RJ45 cable that comes with the PDU+ to connect the serial (COM) connector on a workstation or notebook computer to the RJ45 console connector on the PDU, as shown in Figure 8-7.

> **Note:** A PDU+ default IP address is not published. An initial IP address is configured when you first log in to console with the DB9-to-RJ45 cable, but there is no guarantee that this IP address is the default IP address for any other PDU+. Therefore, we recommend that you connect the console to change the IP address. Also, you must use the cable that ships with the PDU+. It is a special cable. Standard serial cables will not work.

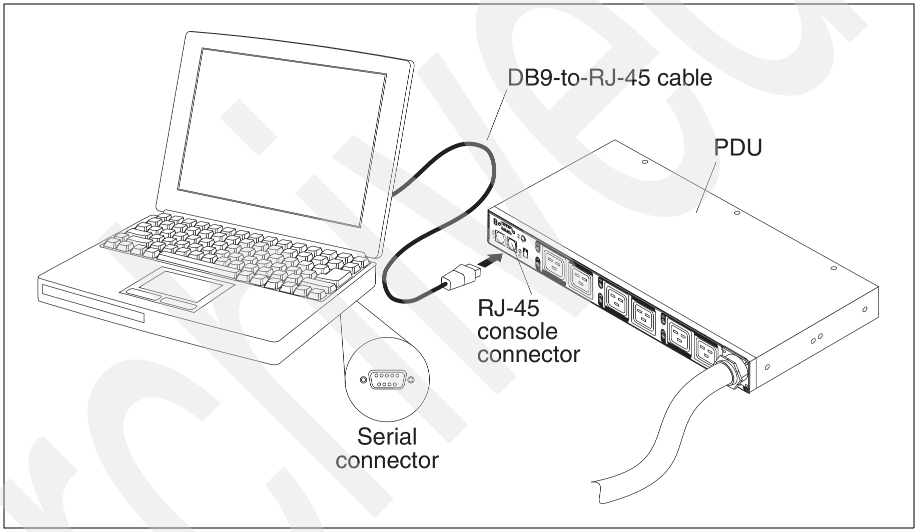


*Figure 8-7 Connecting the console using the DB9-to-RJ45 cable*

2. After you connect the PDU+ using the DB9-to-RJ45 cable, run HyperTerminal to access the PDU+ console. Use the following values to set the HyperTerminal options:
   - Baud rate = 9600 bps
   - Data bit = 1
   - Parity = None
   - Stop bits = 1
   - Flow control = None

3. Click **OK** to start the session, then press any key to display the password prompt dialog box, shown in Figure 8-8. Enter the password. The default is `passw0rd`.

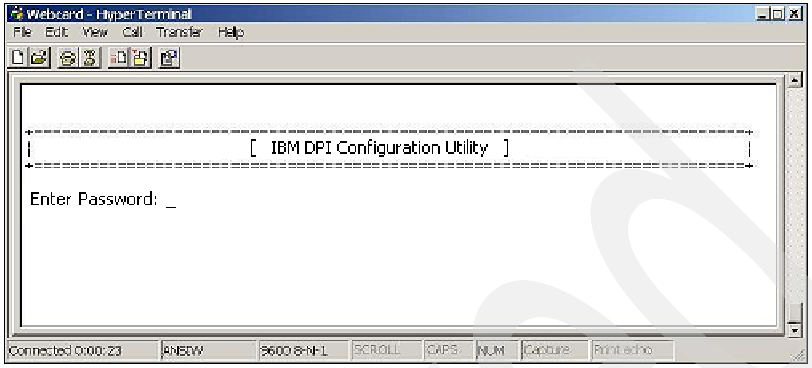


*Figure 8-8 Enter a password*

4. In the IBM DPI Configuration Utility panel, select option **1. IBM DPI Settings**, as shown in Figure 8-9, to set the IP address and other information.

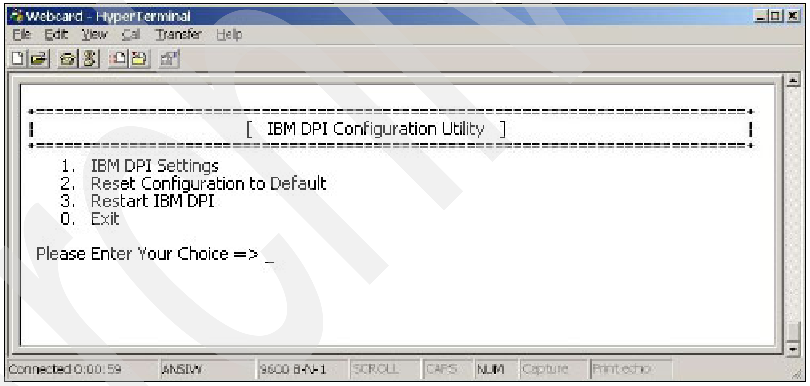


*Figure 8-9 IBM DPI PDU+ console top menu*

5. Select option **1. Set the IP Address, Gateway Address and MIB System Group** to configure the IP as shown in Figure 8-10. You can also set the other network information.

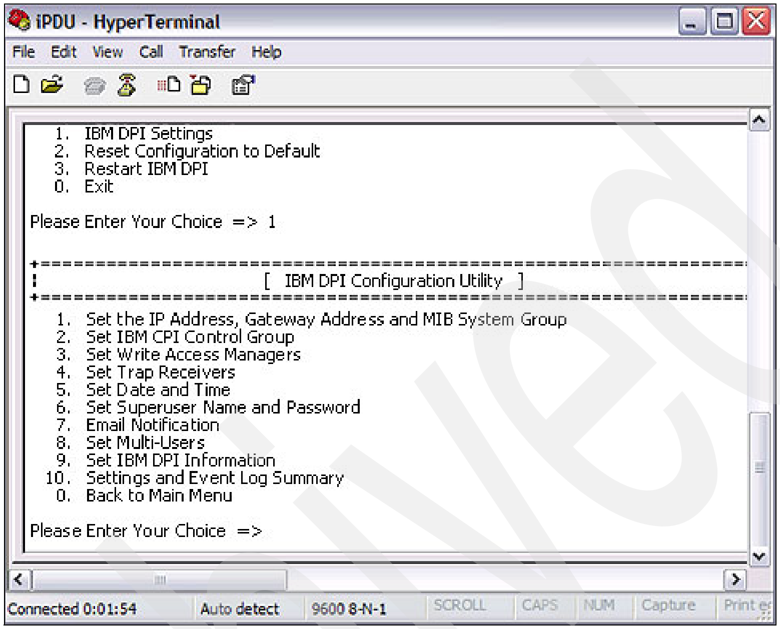


*Figure 8-10 IBM DPI Settings menu*

6. When completed, you can then access the PDU+ from the network, as shown in Figure 8-11.

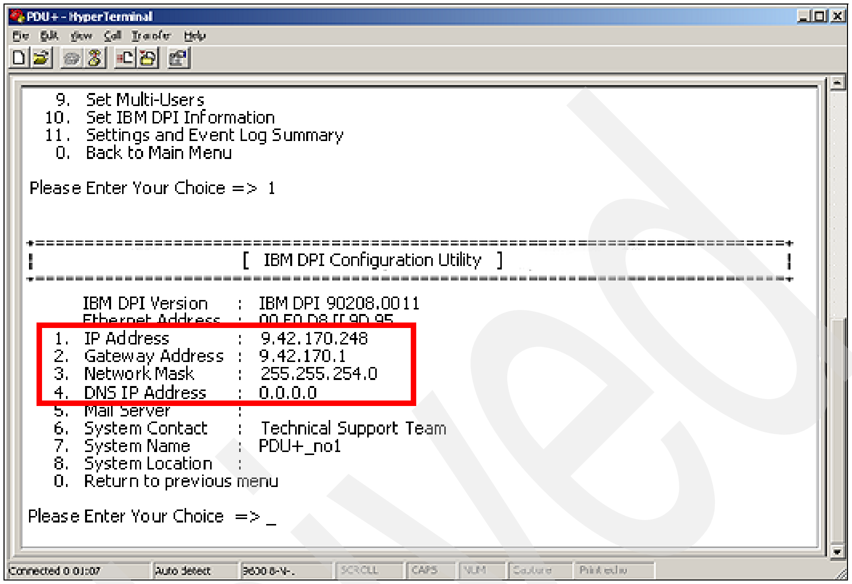


*Figure 8-11 IP Address Settings menu*

7. When the IP configuration is complete, you can use the PDU+ Web interface to change network settings as well as other information. Open a browser, and then use `http://ip_address` to start the Web interface to the PDU+ as shown in Figure 8-12.

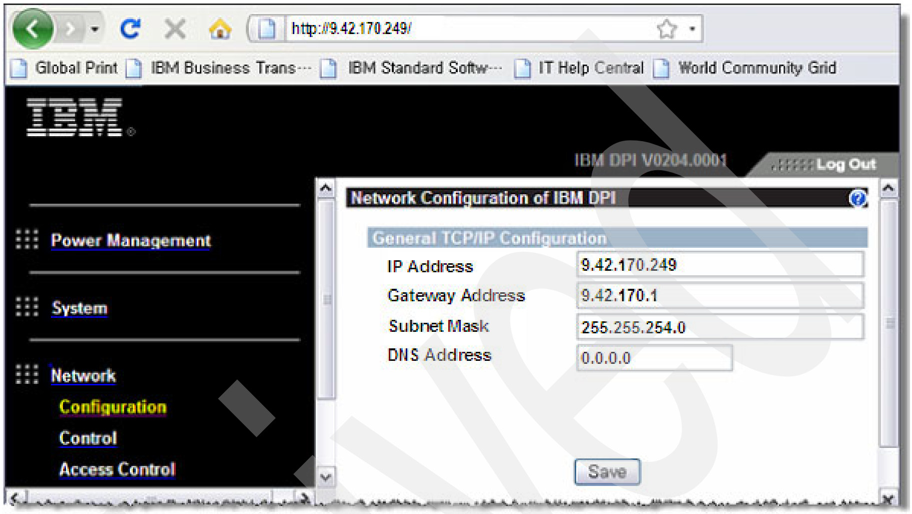


*Figure 8-12 Network Configuration panel on the Web interface*

## 8.2.3 Changing the default user ID and password

By default, authentication is not required to allow Active Energy Manager to access the PDU+, since AEM first attempts access using the default credentials:

- User name = `USERID`
- Password = `passw0rd`

If you prefer, you can change the user name and password to prevent unauthorized users from accessing the PDU+ using the following default settings:

To change the user name and password, follow these steps:

1. Open a browser to the PDU+ Web interface.
2. In the left panel area, click **System** → **Configuration** as shown in Figure 8-13.

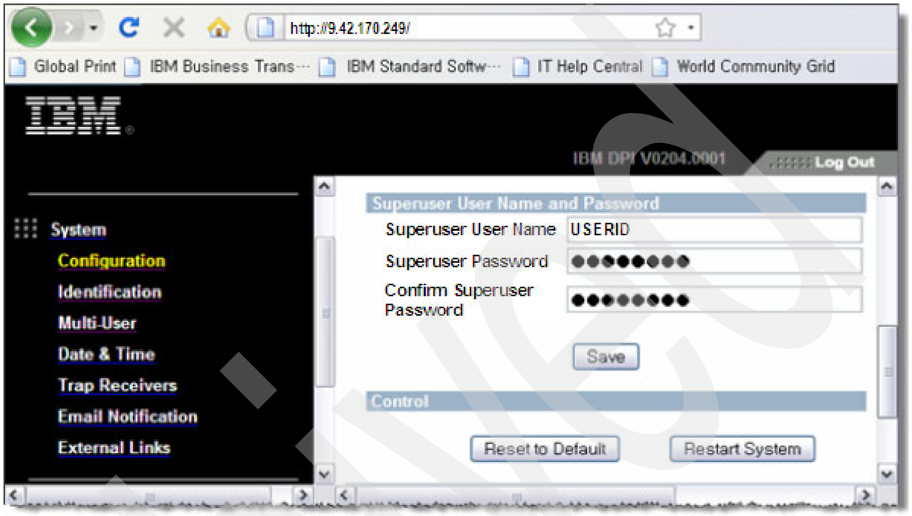


*Figure 8-13 User name and password panel on the Web interface*

3. Enter the user name and password that you want to use.
4. Click **Save** to apply it.

## 8.2.4 Connecting an environmental monitoring probe to the PDU

The environmental monitoring probe has a built-in temperature and humidity sensor and enables you to monitor remotely the temperature and humidity of the environment in which the probe is operating.

The probe is standard with the PDU+, except for the DPI C13 3-phase model, where it is available separately. Connect the environmental monitoring probe to the RJ-45 console connector on the PDU+ as shown in Figure 8-14.

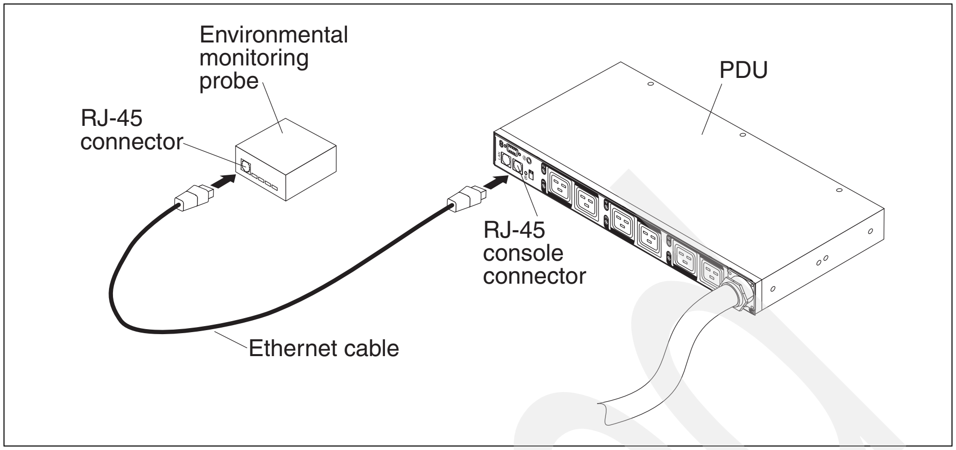


*Figure 8-14 Connecting the environmental monitoring probe to the PDU+*

# 8.3 Configuring the PDU+

In this section, we explain how to configure PDU+ to connect to Active Energy Manager 4.1.1 through a TCP/IP network.

After the PDU+ is set up, power and thermal data from the PDU+ is delivered to Active Energy Manager by SNMP communication. In addition, Active Energy Manager can monitor power consumption from the entire PDU+ as well as from each outlet group.

To configure the PDU+ in Active Energy Manager, you need to first discover the PDU+ so that it can communicate with Active Energy Manager. Then, you can assign managed systems to any of the outlet groups. We describe these steps in the sections that follow.

## 8.3.1 Discovering the PDU+

Before you can configure the PDU+, you must discover it. Active Energy Manager provides many different ways to discover the PDU+. We describe the simplest method here. Follow these steps:

1. In IBM Systems Director 6.1, start the System Discovery task (Figure 8-15 on page 382) and follow these steps:
   a. Select **Single system (IP address)** if you need to discover only one PDU+ by IP address.
   b. Enter the IP address that is assigned to PDU+.

c. Choose **Power Unit** as a resource type.

d. Click **Discover**.

The PDU+ is updated in the Discovered Systems table.

> **Note:** If the PDU+ that you discover already exists in Active Energy Manager, the PDU+ is not updated in Discovered Systems table.

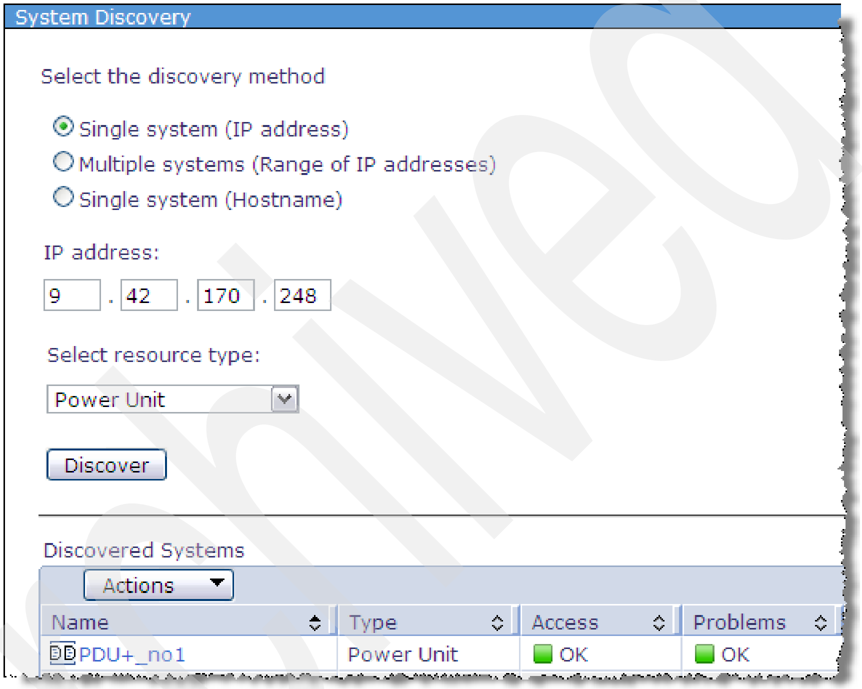


*Figure 8-15 Discovering the PDU+ using the System Discovery task*

> **Note:** PDU+ devices that are configured with SNMP v3 or an SNMP v1/2c with community name other than `public` can be discovered only using the Advanced System Discovery option. For more information about discovering PDUs using Advanced System Discovery, refer to “Discovering PDU devices using advanced system discovery” on page 156.

2. After discovering the PDU+, it appears in the Active Energy Manager Resources group as a device of type Power Unit.

   To find the device, in the Active Energy Manager Resources window, enter `power unit` and then click **Search**, as shown in Figure 8-16.

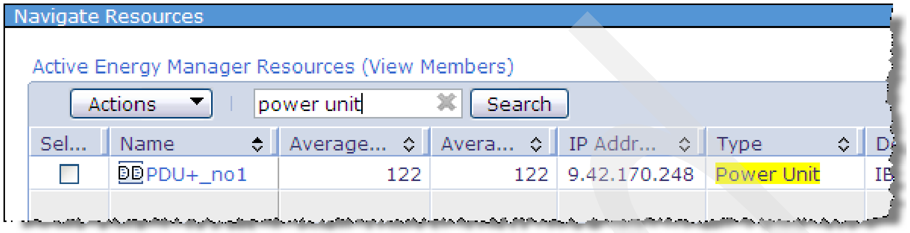


*Figure 8-16 Discovered PDU+ is shown in Active Energy Manager Resource group*

3. Click the name of PDU+ (in this example, **PDU+_no1**). Then, go to the Active Energy tab to verify whether PDU+ firmware fully supports Active Energy Manager.

   If the Active Energy Manager support level is listed as `Update needed` (see Figure 8-17), it means that an update to the PDU+ firmware is required before Active Energy Manager can properly communicate with the device.

   PDU+ firmware should be Version 0208 or later to work with Active Energy Manager. The update is available from:

   http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5073205

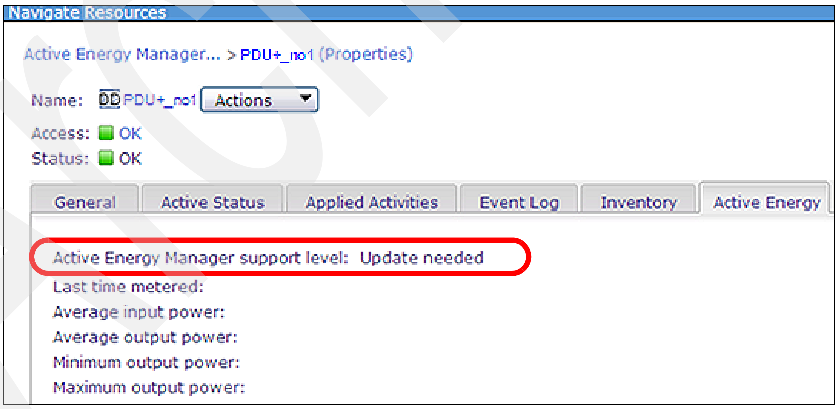


*Figure 8-17 Active Energy tab indicating the PDU+ requires a firmware update*

4. After updating the firmware, Active Energy Manager support level should be listed as `Full`. Also various power information is displayed, as shown in Figure 8-18.

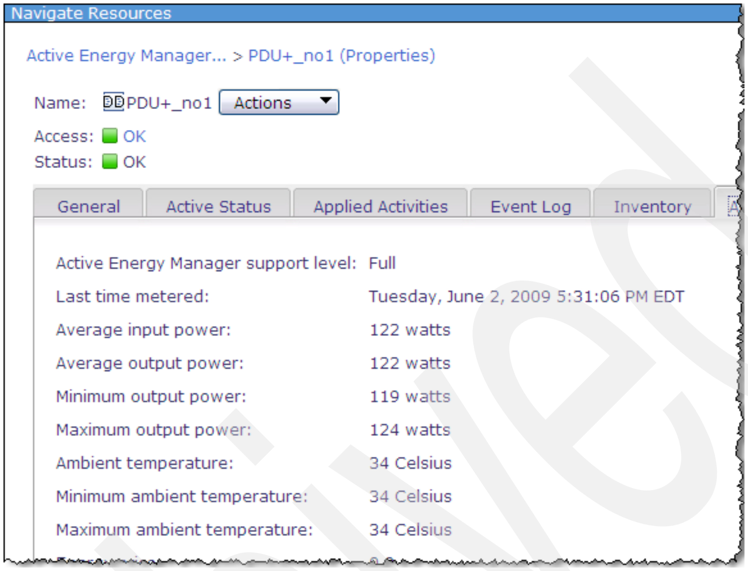


*Figure 8-18 PDU+ Active Energy tab after updating firmware*

## 8.3.2 Assigning systems to an outlet group

In this section, we describe how to configure the outlet groups. To explain the process, we base our instructions on a configuration where a server named `non_IBM_server` is connected to PDU+ Outlet Group 1 as shown in Figure 8-19.

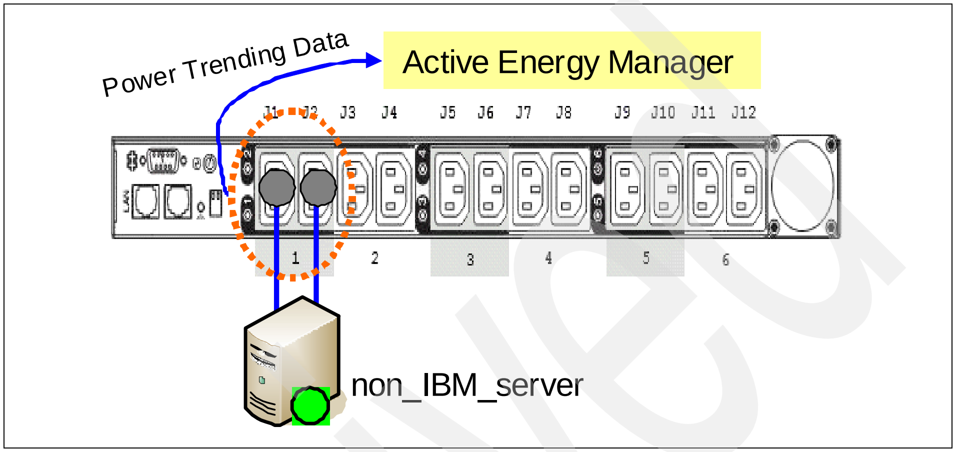


*Figure 8-19 Non_IBM_server connected to Outlet Group 1*

With the DPI C13 PDU+ and DPI C13 3-phase PDU+, power outlets are divided into six outlet groups. To monitor power data from each outlet group, edit the outlet group and assign a managed system.

To assign a managed system to each outlet group in this example, follow these steps:

1. Right-click the name of PDU+ (in this example PDU+_no2), and then select **Energy** → **Configure Metering Device** as shown in Figure 8-20.

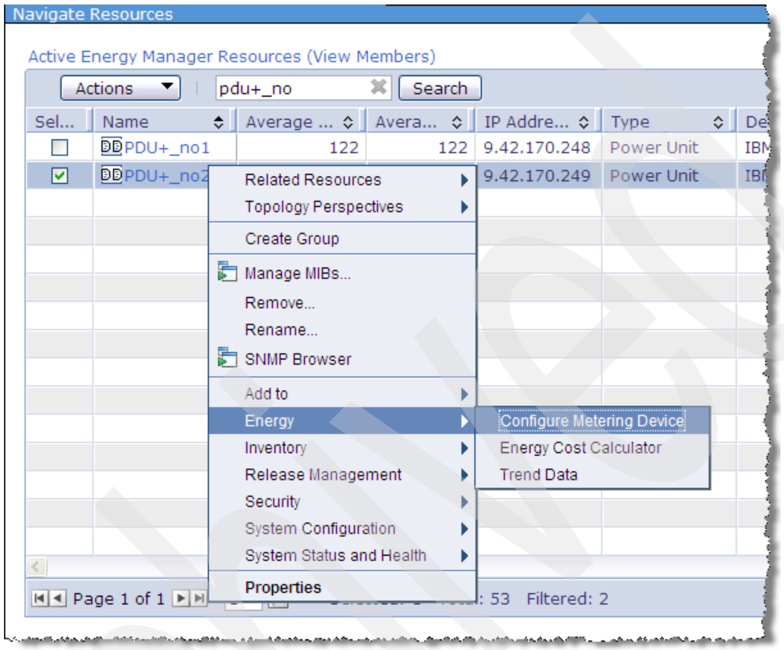


*Figure 8-20 Selecting Configure Metering Device to edit each outlet group*

2. Select **Outlet Group1**, as shown at **1** in Figure 8-21, and then click **Edit** to assign a managed system.

> **Note:** As shown at **2** in the figure, click **Advanced configuration** if you want to launch the PDU+ Web console interface directly to change PDU+ parameters.

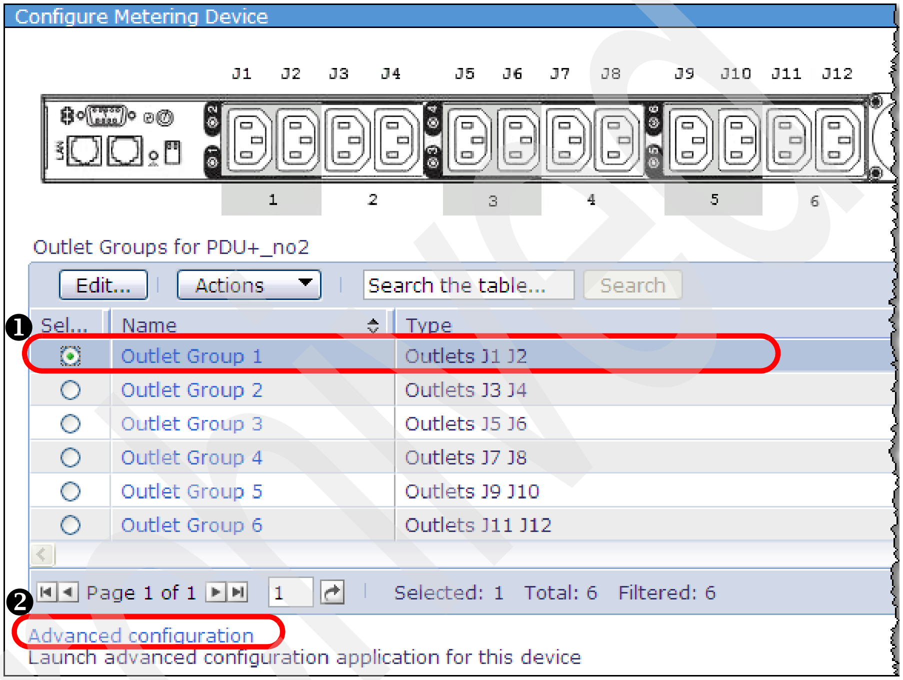


*Figure 8-21 Configure Metering Device task for outlet groups*

3. In the Outlet Group 1 Edit task, select **Metered Resources**, and then click **Add resource** to create a managed system as a metered resource, as shown in Figure 8-22. If the metered resource already exists, click **Browse** to find it.

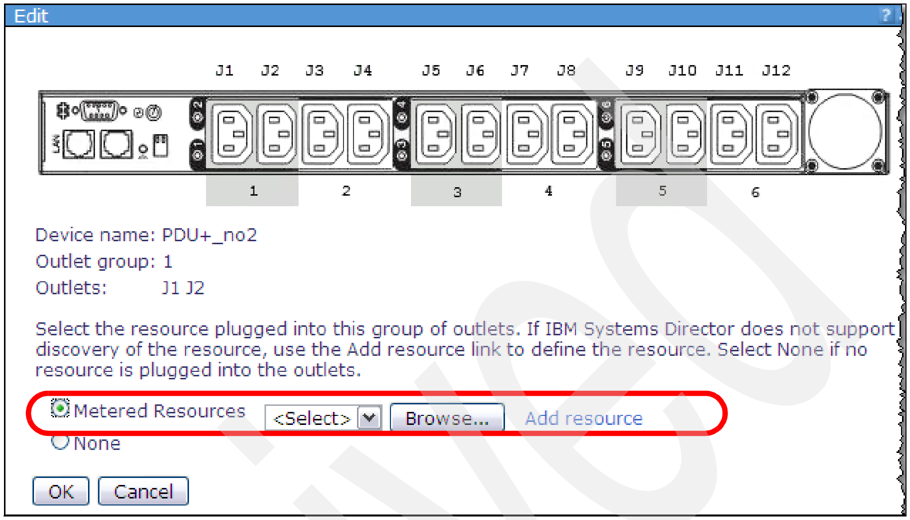


*Figure 8-22 Assign a managed device to Outlet Group 1*

4. Enter the name of system to which you connected Outlet Group 1, as shown in Figure 8-23, and a description (which is optional). The name is shown as a logical module in the Active Energy Manager resources group.

**Tip:** For IT equipment with multiple line cords, associate the same resource name each of the PDU outlets or outlet groups into which the line cords are plugged to show the total power for that resource.

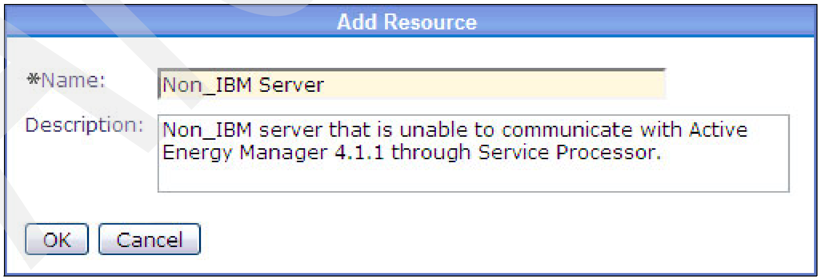


*Figure 8-23 Add Resource window*

5. Finally, the system (in this example `Non_IBM Server`) that is connected to Outlet Group 1 is created as a Logical Module, as shown in Figure 8-24. Thus, Active Energy Manager monitors power data at the level of the individual system just as it monitors power consumption by the service processor.

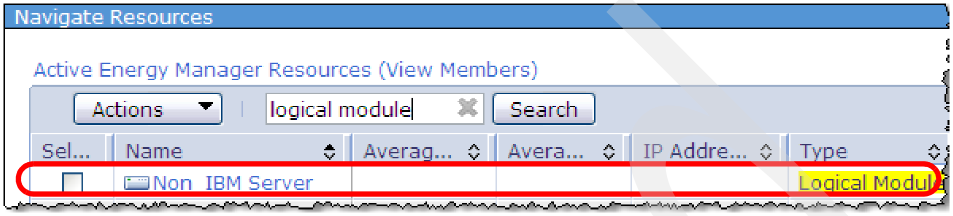


*Figure 8-24 Logical Module from PDU+ Outlet Group is created*

# 8.4 Monitoring the PDU+

In this section, we explain how to collect power data using the IBM PDU+ as well as a supported Raritan PDUs with a scenario that uses different levels of monitoring. We also cover how to monitor the power trend that is provided by Power Domain redundancy environment.

A general PDU cannot monitor power; however, the IBM PDU+ enables you to monitor remotely the power of systems that are connected to a PDU+ Outlet Group using the TCP/IP network.

## 8.4.1 Monitoring power consumption

By using a PDU+, you can associate a system that cannot communicate with Active Energy Manager with a PDU+ Outlet Group to monitor power through the service processor and to provide power data to Active Energy Manager.

A *logical module* is an individual system under the Outlet Group and is the minimum level of monitoring power. By integrating grouping features to Active Energy Manager, you can monitor from the individual system level to the data center level using raw power data that is provided from the logical module.

In this scenario, we monitor power trends from the system that is associated with the Outlet Group and then extend the monitoring range to the data center level as shown in Figure 8-25. The monitoring area can be specified to fit your needs.

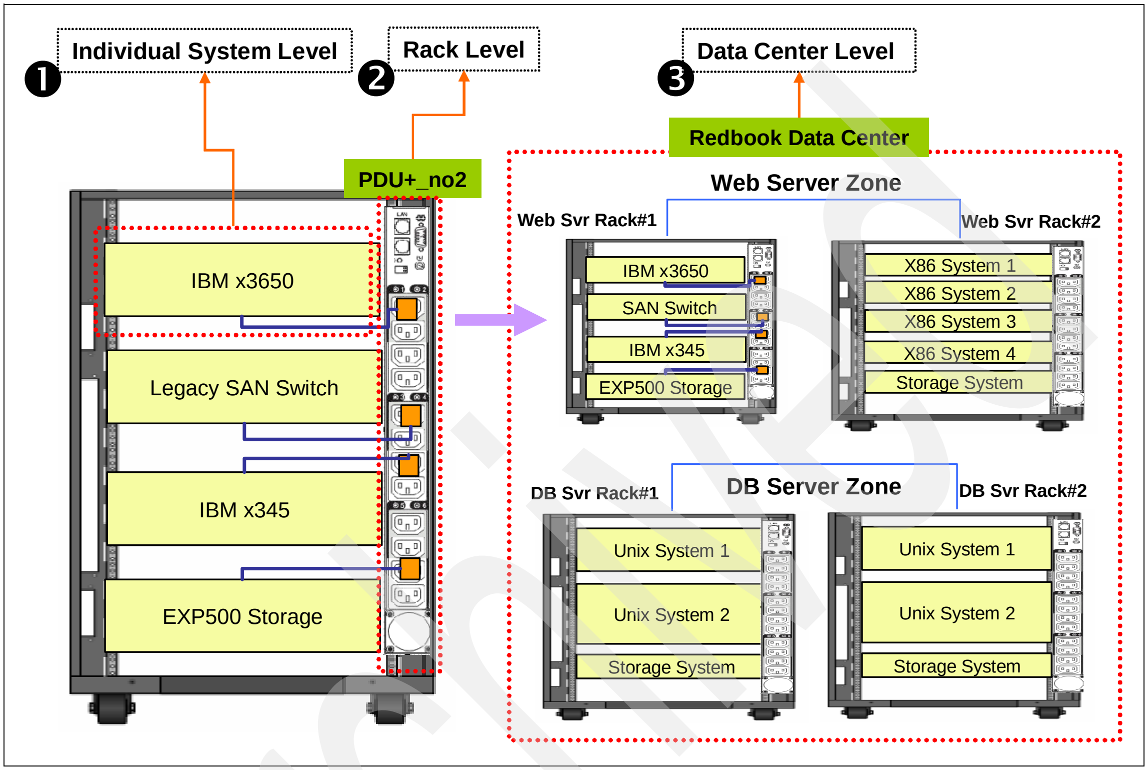


*Figure 8-25 Scenario diagram: Monitoring power from individual system to data center*

In this scenario, we cover the following topics:

1. Outlet Group monitoring (individual system level)
2. PDU+ (rack level) monitoring
3. Data center level monitoring

## Outlet Group monitoring (individual system level)

Monitoring power of an individual system is key to determining the power consumption of the data center. In this example, four different types of systems are associated with each Outlet Group.

As shown in Figure 8-26, the Metered Resources field shows the IBM x3650, which is an individual system that is monitored and which is considered a Logical Module.

**Note:** To configure an Outlet Group, refer to 8.3.2, “Assigning systems to an outlet group” on page 385.

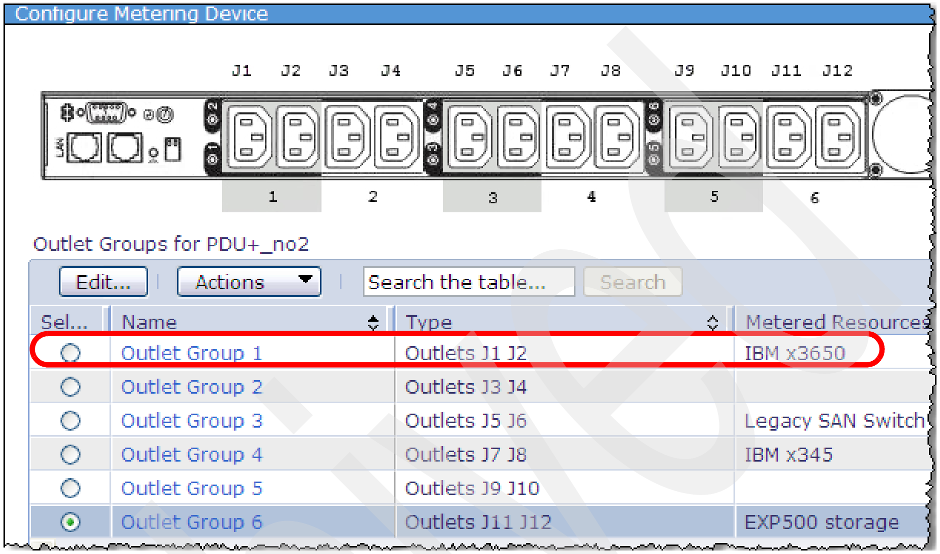


*Figure 8-26 The systems connected to the PDU+ in this example*

To monitor an individual system that is connected to each Outlet Group, follow these steps:

1. In the Active Energy Manager Resources Group on the Active Energy Manager home page, enter `logical module`, and then click **Search** to find individual systems that are connected to outlet groups (Figure 8-27).

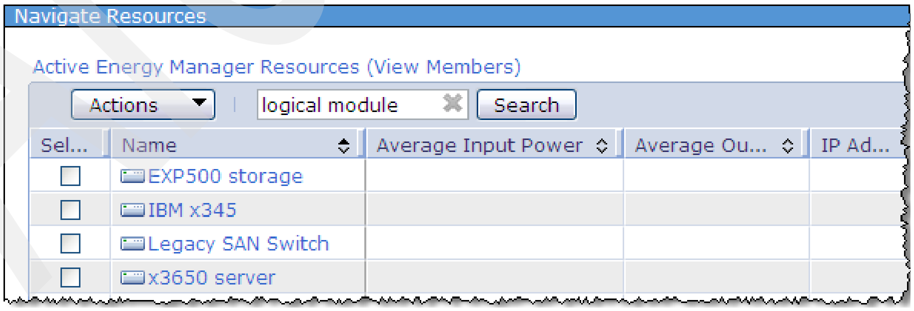


*Figure 8-27 Individual systems as a logical module connected to outlet groups*

**Note:** Because a logical module in Active Energy Manager uses the Average Input Power (externally metered) field to display power-trend data, its power usage information does not display as a default field setting, as shown in Figure 8-27.

2. To change the field setting, click **Actions**, and then click **Columns** as shown in Figure 8-28.

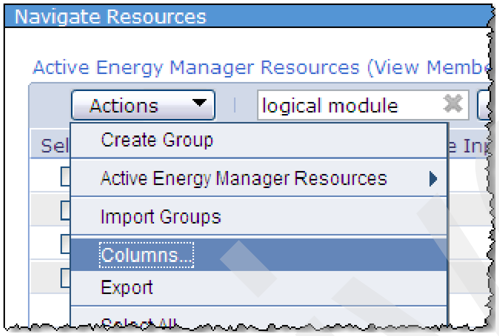


*Figure 8-28 Changing the field setting in the Navigate Resources table*

3. Select **Average Input Power (externally metered)** from the Available Columns box, and then click **Add** to move to it to the Selected Columns box (as shown in Figure 8-29). Click **Up** until it is located at the very top of the list so that this column appears left-most in the Figure 8-30 on page 393 just after the Name field.

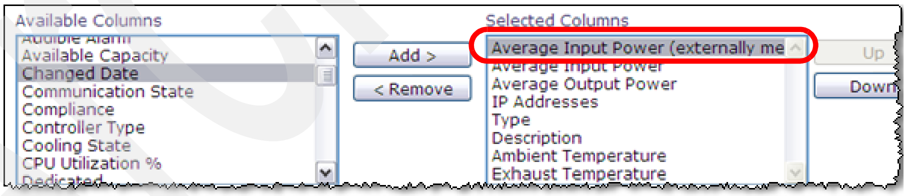


*Figure 8-29 Selecting the field*

The power-trend data now displays under the “Average Input Power (externally metered)” column, as shown in Figure 8-30.

> **Note:** Refer to 2.2.2, “Terminology” on page 45 for information about Average Input Power (externally metered).

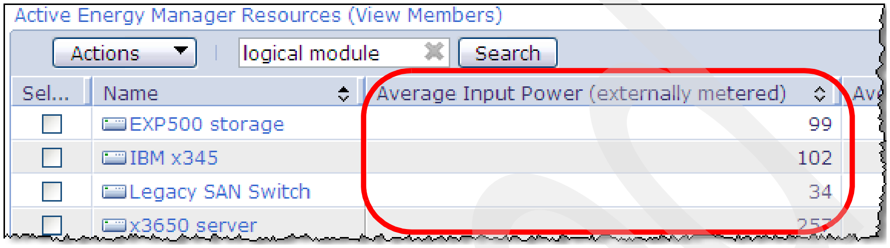


*Figure 8-30 Power-trend data displays*

4. To view the power-trend data of the individual system (in this example the IBM x3650), right-click **IBM x3650**, and then select **Energy** → **Trend Data**, as shown in Figure 8-31.

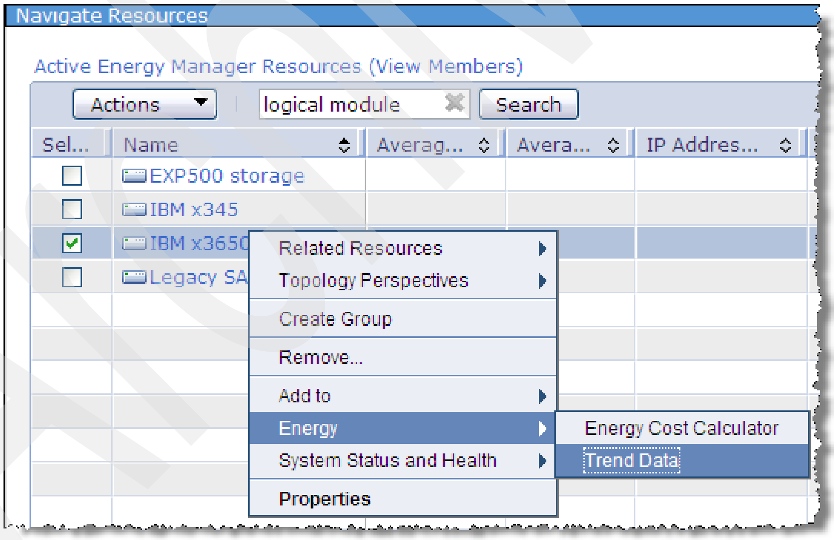


*Figure 8-31 Going to view power-trend of the individual system*

The power-trend data of the IBM x4650 (the individual system) displays on the chart, as shown in Figure 8-32.

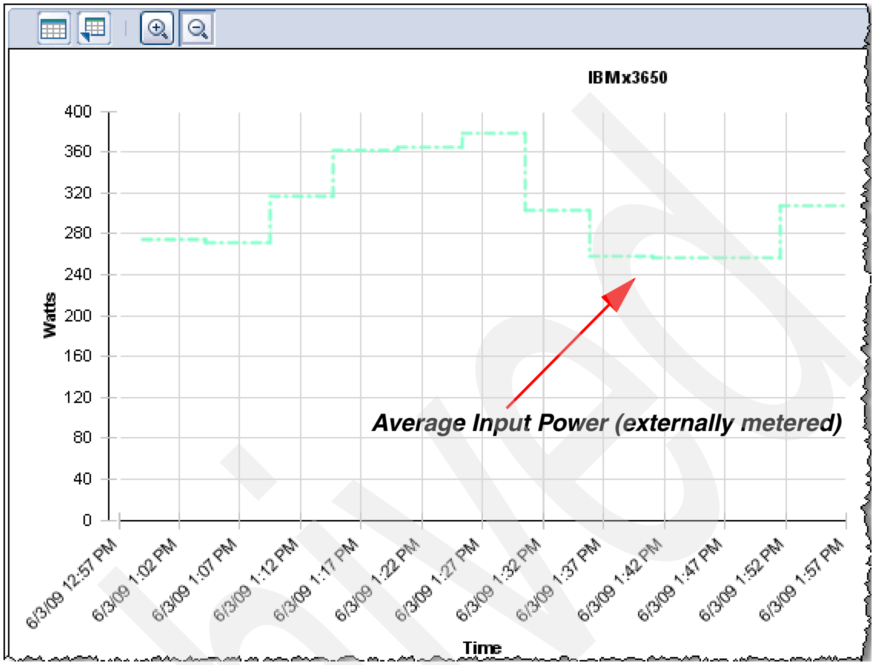


*Figure 8-32 The individual system (IBM x3650) level power-trend monitoring view*

## PDU+ (rack level) monitoring

At the next level above monitoring individual systems, you can monitor power trends of a rack, including all the systems that are connected to the PDU+, as shown at 2 in Figure 8-25 on page 390.

To monitor a PDU+ that connects many different types of systems, follow these steps:

1. In the Active Energy Manager Resources Group, enter `power unit`, and click **Search** to find the PDU+.
2. Right-click the PDU that you want to manage (in this example **PDU+_no2**) and click **Energy** → **Trend Data**, as shown in Figure 8-33.

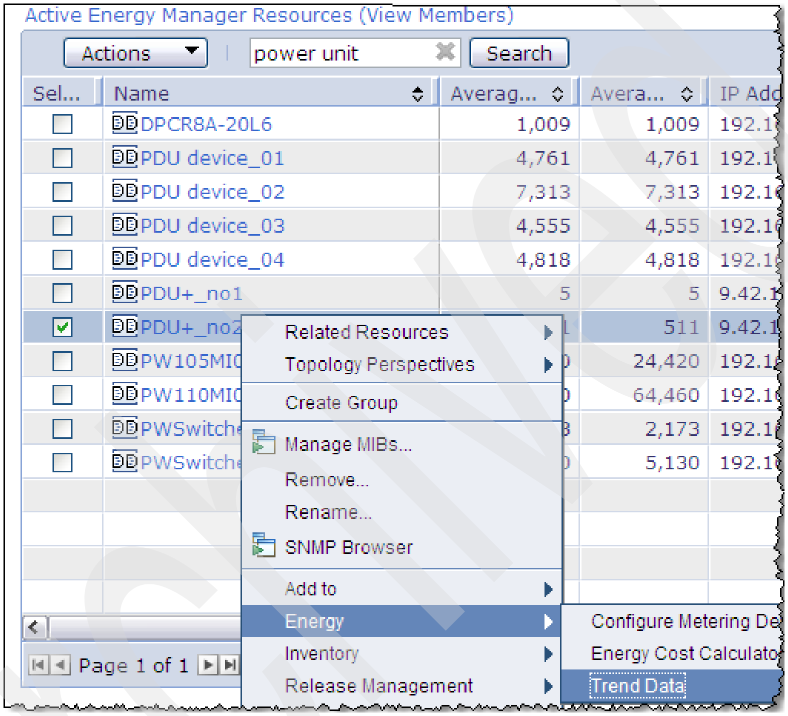


*Figure 8-33 Searching for the PDU+*

3. The power-trend information about PDU+no2 with rack-level monitoring displays, as shown in Figure 8-34 on page 396.

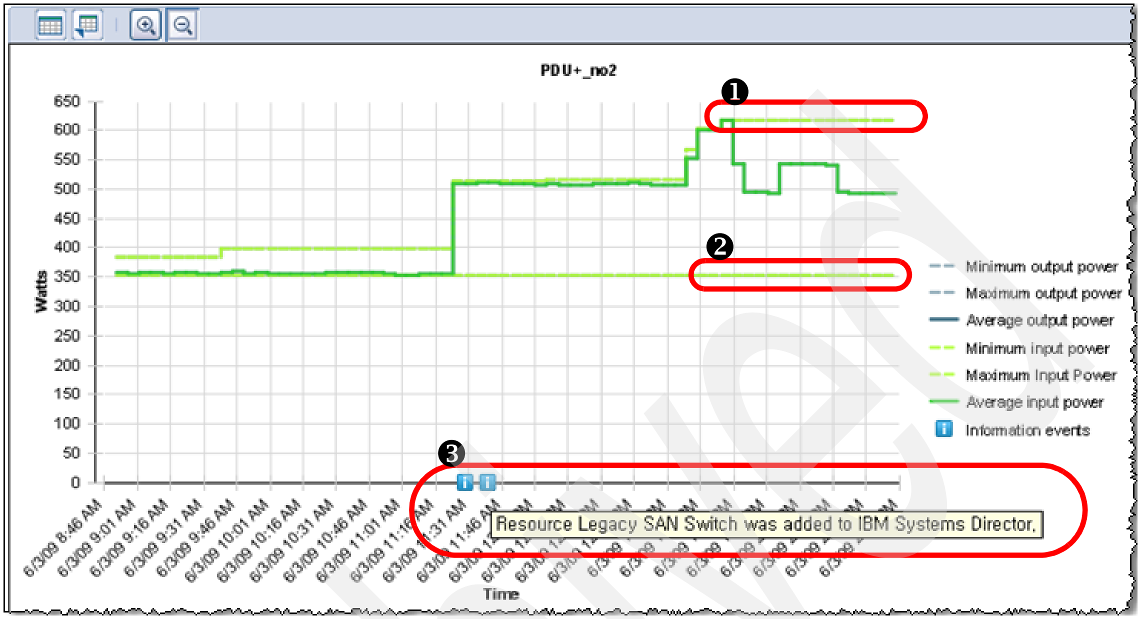

*Figure 8-34 Last 6 hours power trend data of the PDU+ (rack level) monitoring*

In the figure at 1, when the average input power reaches the maximum during the monitoring period, it is recorded. Alternatively, as shown at 2, when the average input power reaches the minimum during the monitoring period, it is recorded. If any event information occurs (in this example the Logical Module is added to the PDU+ Outlet Group), for example as shown at 3, it also displays on the power-trend chart.

## Data center level monitoring

By using IBM Systems Director grouping feature, you can monitor the power trends of entire systems in the data center as a *group*. (For more information, see **3** in Figure 8-25 on page 390.)

For data center level monitoring using the grouping feature, follow these steps:

1. In the Groups Navigate Resources area, click **Create Group** to generate the group (in this example *Redbook DataCenter*) that includes all the system in the data center.

> **Note:** Refer to 4.6.3, “Customizing the Active Energy Manager resource group” on page 131 for more information.

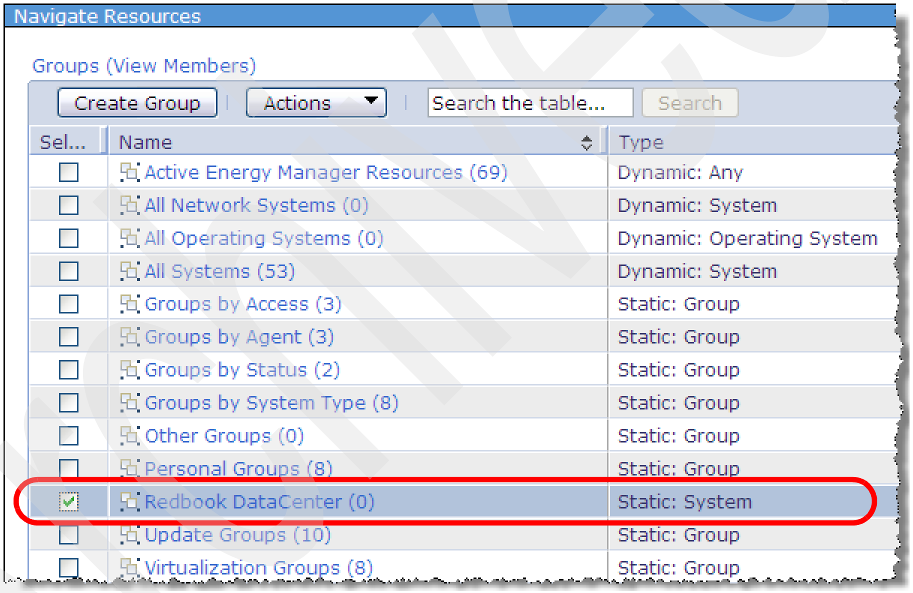


*Figure 8-35 Creating Group (Redbook DataCenter) to monitor entire systems*

In this scenario, the Redbook DataCenter group includes two different power-trend monitoring zones:

– DB Server Zone
– Web Server Zone

As shown in Figure 8-36, each zone includes the following racks (PDU+):

– DB Server Zone, shown at 1, which is *DB Svr Rack#1* and *DB Svr Rack#2*
– Web Server Zone, shown at 2, which is *Web Svr Rack#1* and *Web Svr Rack#2*

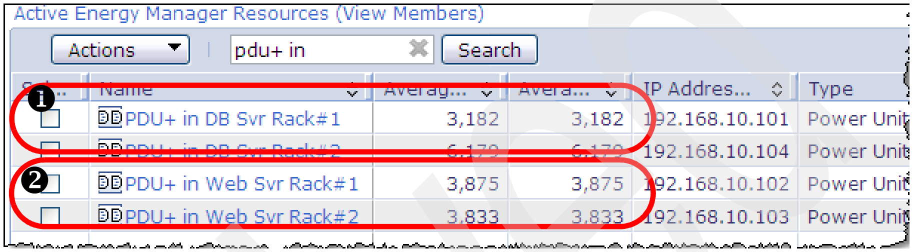


*Figure 8-36 Redbook DataCenter sub-group by monitoring zone*

2. To create the DB Server Zone group with PDU+ in DB Svr Rack#1 and PDU+ in DB Svr Rack#2, select **DB Svr Rack#1** and **DB Svr Rack#2**, and then click **Actions** → **Create Group**, as shown in Figure 8-37.

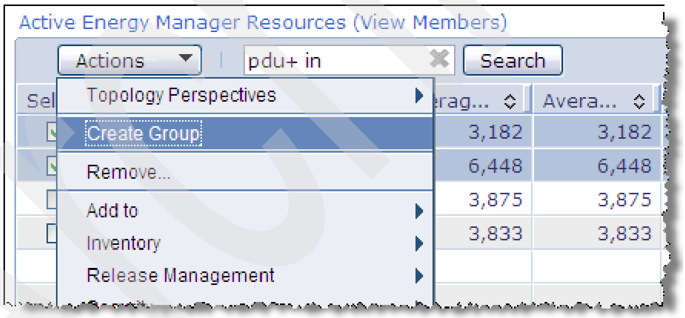


*Figure 8-37 Creating the DB server zone*

3. In the Group Editor Wizard, enter `DB Server Zone`, and click **Next**.

4. To locate the DB Server Zone group in the Redbook DataCenter group, click **Browse** and select **Redbook DataCenter**, as shown in Figure 8-38. Then, click **Next**.

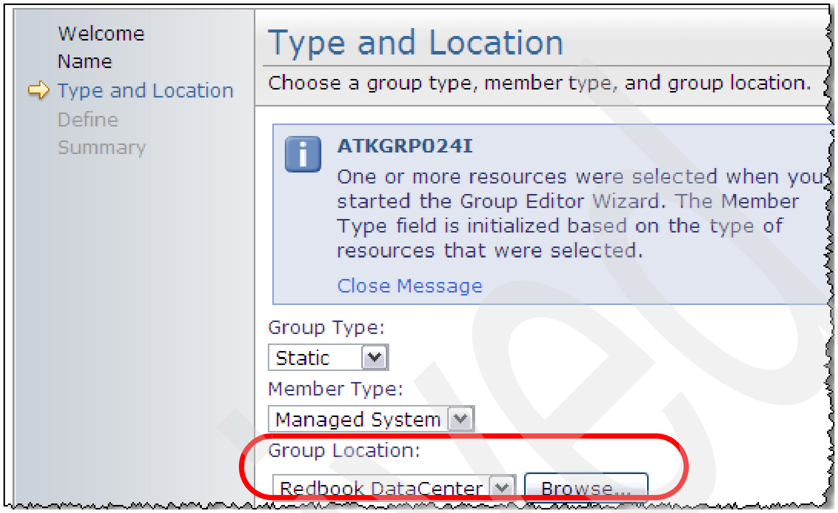


*Figure 8-38 Type and Location of creating group*

5. In the Define task panel, click **Next**.

6. Verify the information in the Summary panel, shown in Figure 8-39, and then click **Finish** to create the group.

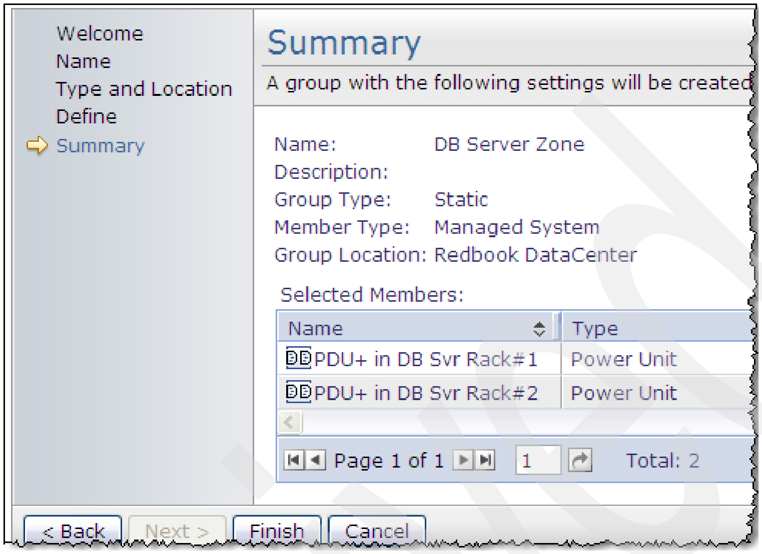


*Figure 8-39 Summary panel for the DB Server Zone group*

7. You can create the Web Server Zone using the same steps that you used to create the DB Server Zone (following steps 2 on page 398 through 6).

   The summary of the Redbook DataCenter group hierarchy is as follows (Figure 8-40):

   – DB Server Zone: DB Svr Rack#1, DB Svr Rack#2
   – Web Server Zone: Web Svr Rack#1, Web Svr Rack#2

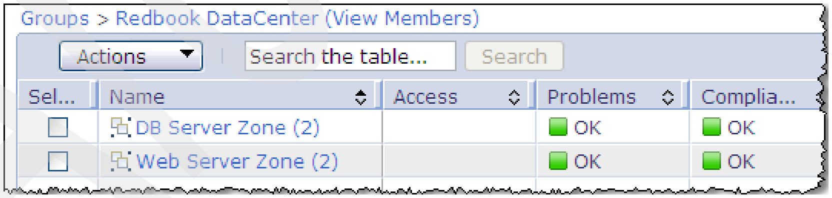


*Figure 8-40 Redbook DataCenter group hierarchy*

8. To view the power-trend data for the DB Server Zone, right-click **DB Server Zone (2)** and then select **Energy** → **Trend Data** as shown in Figure 8-41.

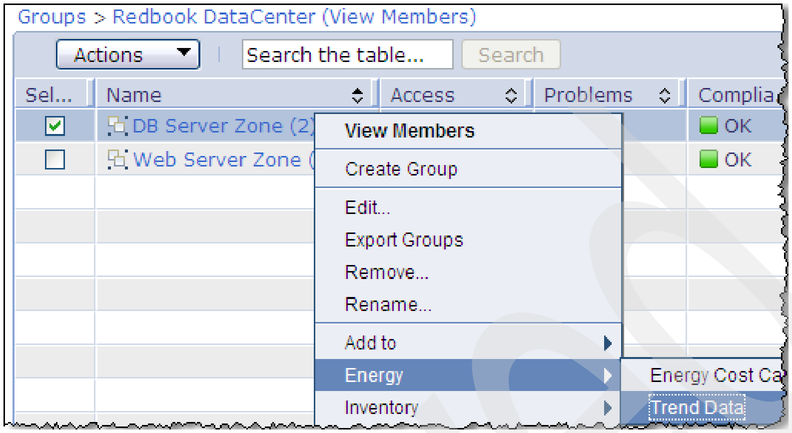


*Figure 8-41 Viewing the power-trend data of the DB Server Zone systems*

Figure 8-42 shows the power-trend data chart for the DB Server Zone.

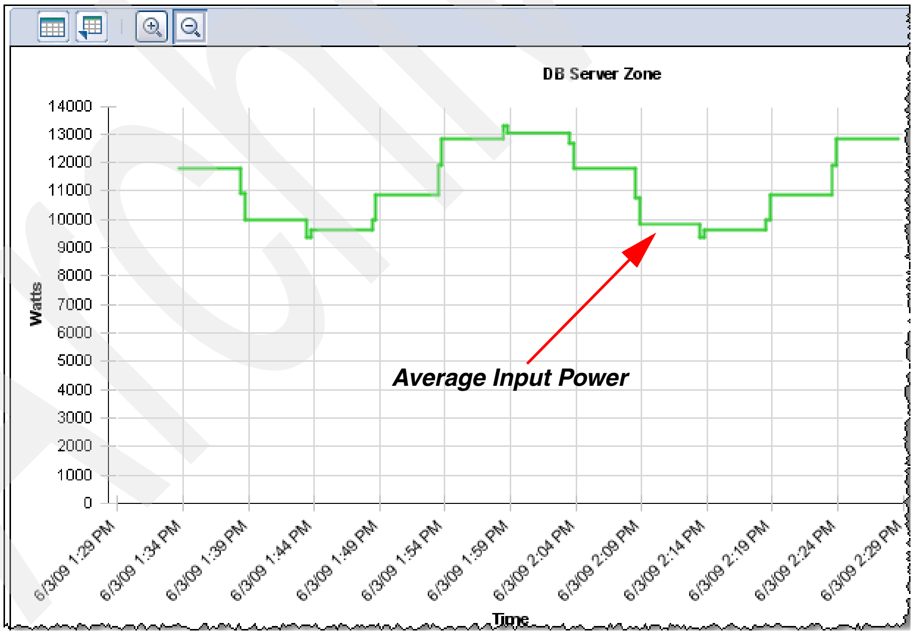


*Figure 8-42 The power-trend data chart for the DB Server Zone*

Figure 8-43 shows the power-trend data chart for the Web Server Zone.

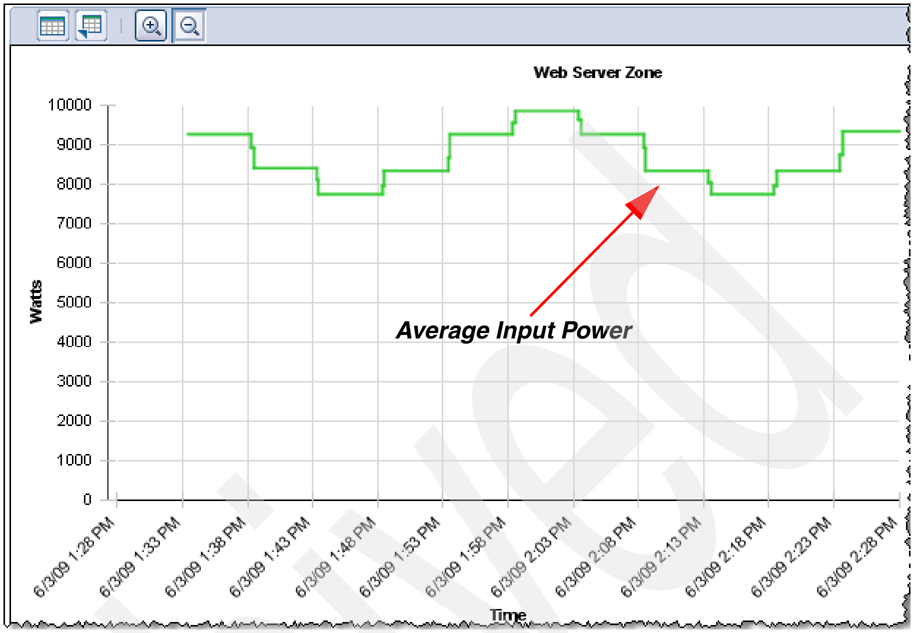


*Figure 8-43 The power-trend data chart for the Web Server Zone*

9. To view the power-trend data for the Redbook DataCenter group, which is data center level power monitoring, right-click **Redbook DataCenter (2)** in the Group Navigation, area and then select **Energy** → **Trend Data**, as shown in Figure 8-44.

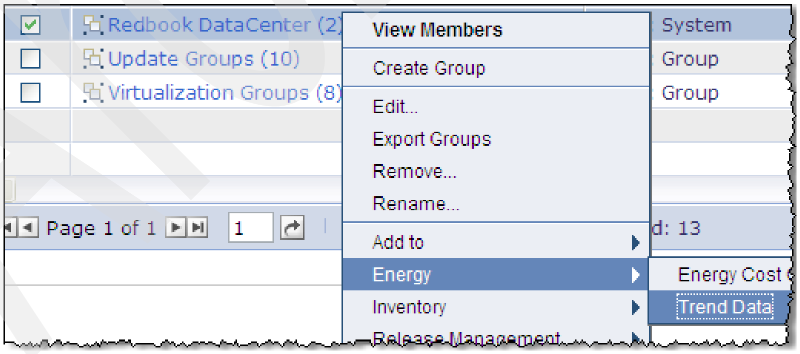


*Figure 8-44 Viewing the power-trend of all systems in the Redbook DataCenter group*

Figure 8-45 shows the power-trend data chart of all the systems in the Redbook DataCenter group as data center power monitoring level.

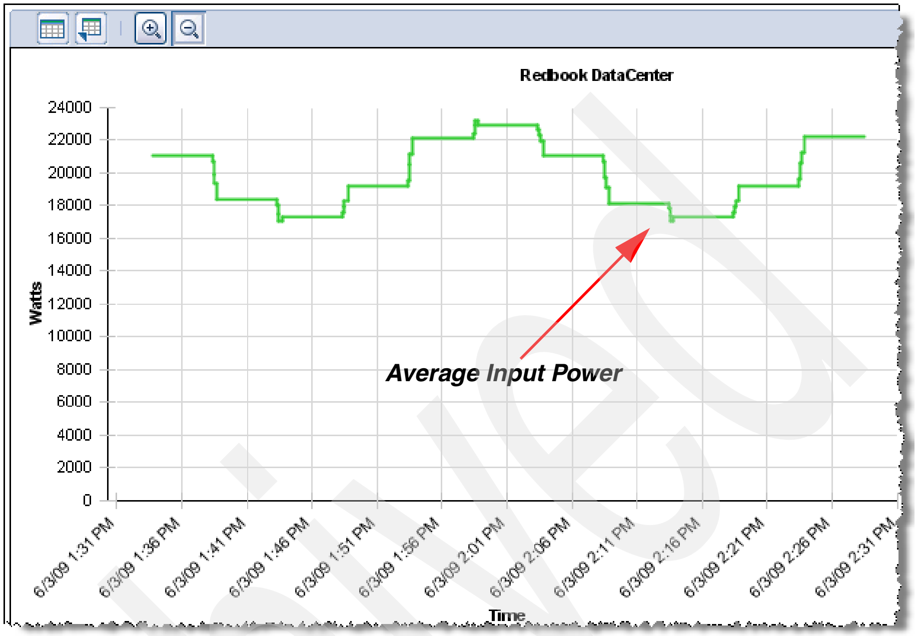


*Figure 8-45 The power-trend data chart for the Redbook DataCenter group*

## 8.4.2 Monitoring temperature and humidity using a PDU+ probe

By connecting an environmental monitoring probe to the PDU+, you can remotely monitor the temperature and humidity of the environment in which the probe operates.

> **Note:** For more detailed information about the environmental monitoring probe, refer to 8.2.4, “Connecting an environmental monitoring probe to the PDU” on page 380.

To monitor temperature and humidity using the environmental monitoring probe, follow these steps:

1. From the Active Energy Manager home page in IBM Systems Director, scroll down to Manage and in the Management Tasks box, click **Configure metering devices**.
2. Select the PDU with the probe installed, then select the **Probe Humidity** entry as shown in Figure 8-46, and click **Edit**.

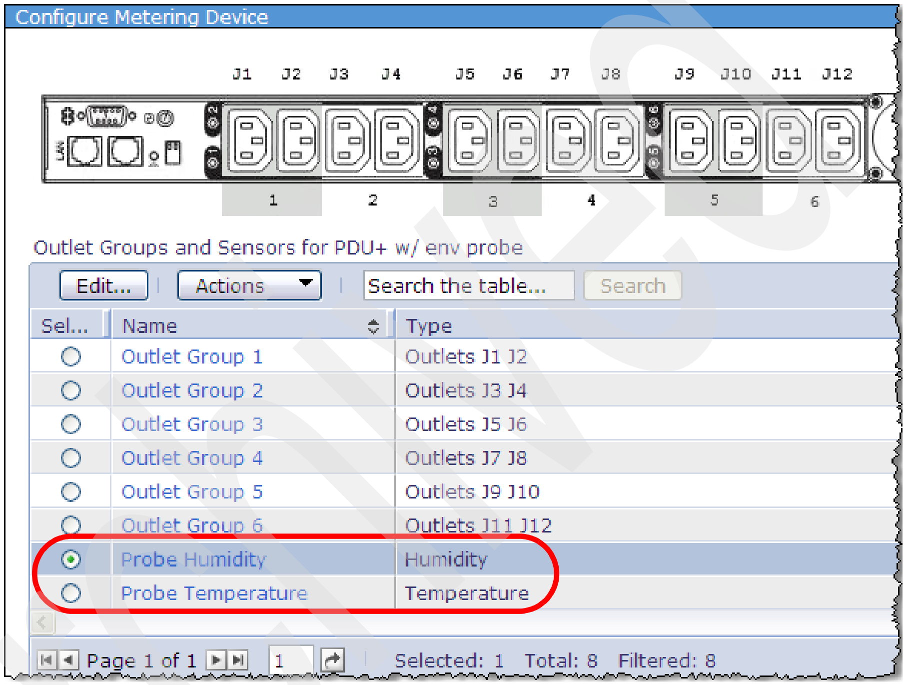


*Figure 8-46 Configure Metering Device panel*

3. Select **Metered Resources** and click **Add resource** to create the logical module (in this example *PDU+ w/ env probe Humidity*), then click **OK**, as shown in Figure 8-47.

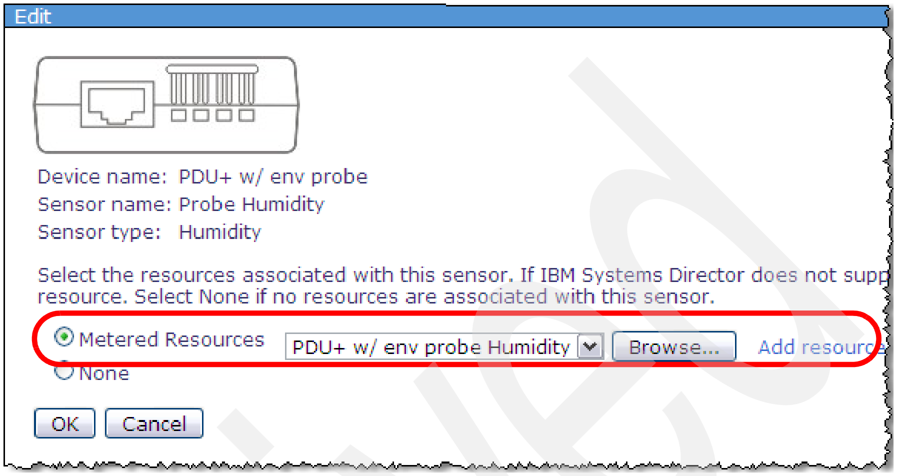


*Figure 8-47 Configuring probe humidity sensor*

4. To view the humidity-trend data for this PDU+, right-click **PDU+ w/ env probe Humidity** and select to **Energy** → **Trend Data** as shown in Figure 8-48.

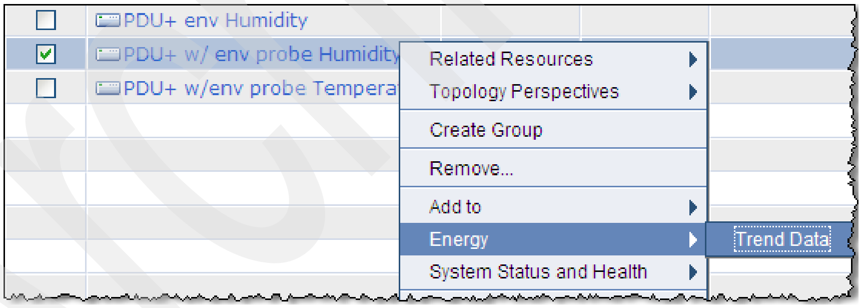


*Figure 8-48 Viewing the humidity-trend data of PDU+ w/ env probe Humidity*

Figure 8-49 shows that the humidity of the environment in which the PDU with this humidity probe is operating. In the trend-data chart, the percent range is used for humidity.

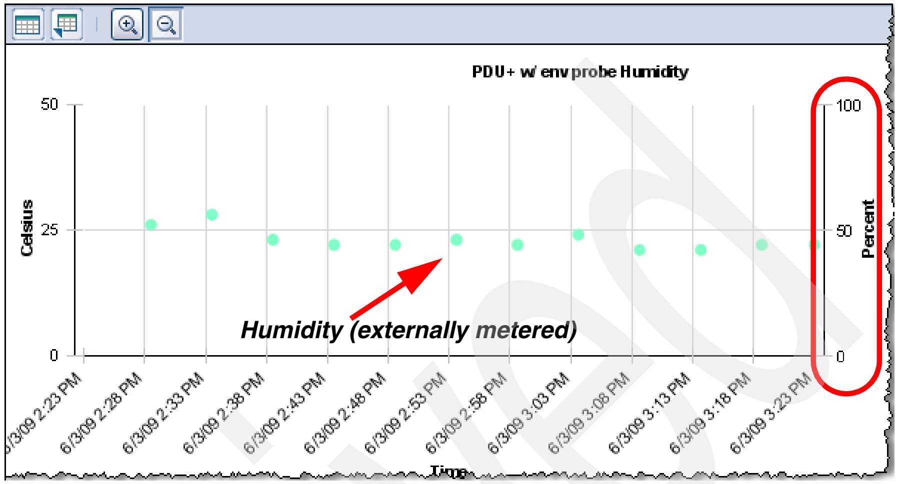


*Figure 8-49 The humidity-trend data for the PDU+ with the humidity probe*

Figure 8-50 shows that the temperature of environment in which the PDU with this temperature probe is operating. In the trend-data chart, the Celsius range is used for temperature.

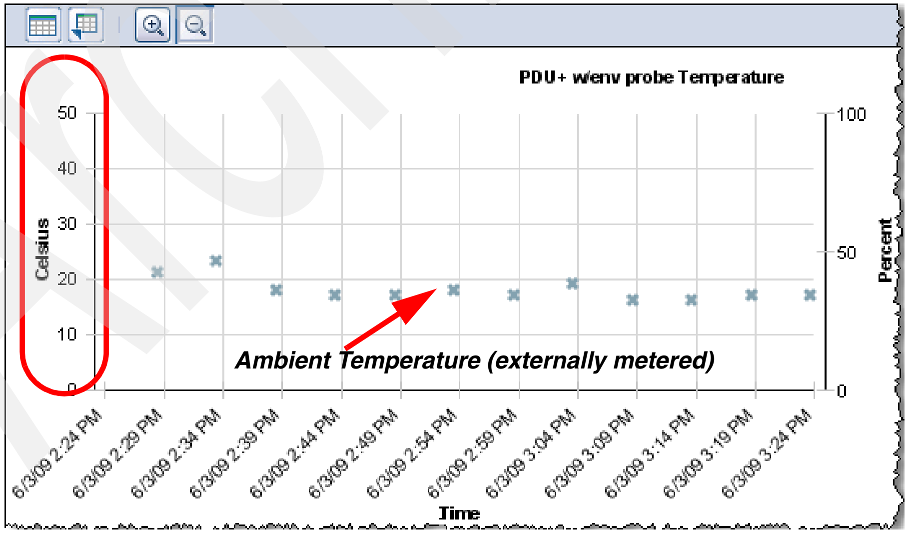


*Figure 8-50 The temperature-trend data for the PDU+ with the temperature probe*

## 8.4.3 Monitoring power data on Raritan PDUs

Active Energy Manager 4.1.1 also integrates with the Raritan and other non-IBM PDUs that support power-trend monitoring.

The supported hardware is the Dominion PX PDU. The Dominion PX DPCR8A-20L6 model includes the following specifications:

- Input
  - Nominal voltage: 208-240V AC, single phase
  - Maximum line current per phase: 20A
  - Rated current: 16A
  - Power: 3.3kVA
  - Frequency: 50-60 Hz
  - Plug type: NEMA L6-20P
- Output
  - Nominal voltage: 208-240V AC
  - Maximum current draw: 20A
  - Switched receptacles: 8 x IEC C-13
- Metering
  - Inlet: Voltage, current, power
  - Outlet: Voltage, current, power
- Rack Environment Management
  - Optional temperature and humidity sensors (RJ-12 connector)

> **Note:** You can find the specification for the Dominion PX DPCR8A-20L6 model at:
>
> http://www.raritan.com/products/power-management/Dominion-PX/DPCR8A-20L6/

The Raritan PDU is listed as a power unit under the Active Energy Manager Resources navigation group, as shown in Figure 8-51.

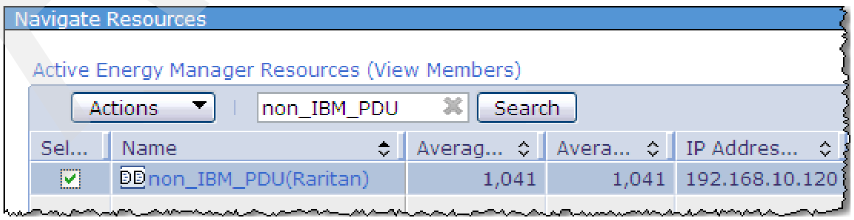


*Figure 8-51 Raritan PDU under the Active Energy Manager Resources navigation group*

The Raritan PDU in this example provides eight different outlet groups, as shown in Figure 8-52. You can associate a logical module with each Outlet Group by configuring a Metering Device. Refer to 8.3.2, “Assigning systems to an outlet group” on page 385 for more information.

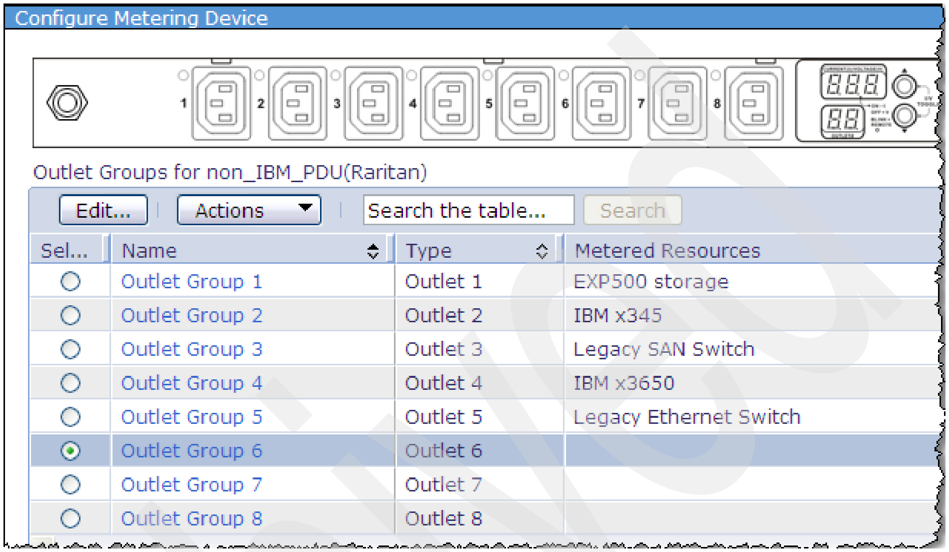


*Figure 8-52 Raritan PDU Outlet Groups*

To view the power-trend data for the Raritan PDU, which in this example is *non_IBM_PDU (Raritan)*, click the PDU+ (in this example **PDU+_no2**) and go to **Energy** → **Trend Data**, as shown in Figure 8-53.

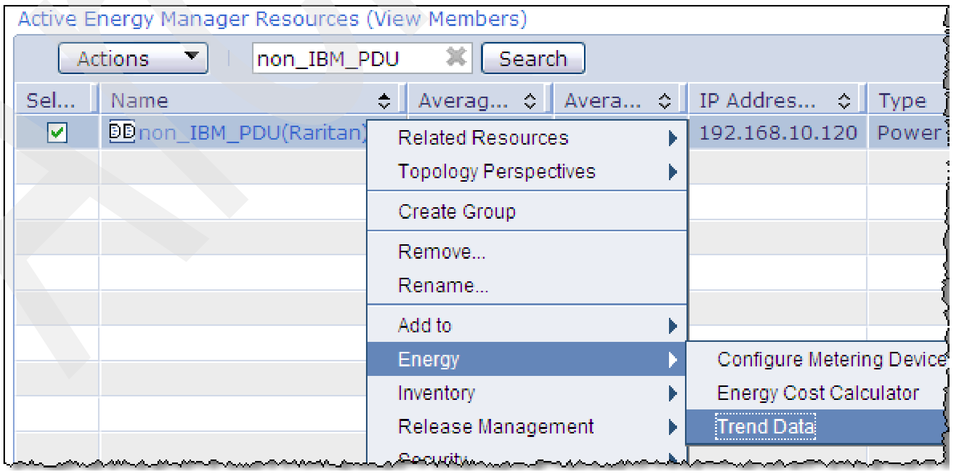


*Figure 8-53 Viewing the power-trend data for the Raritan PDU*

Figure 8-54 shows the power-trend data for the Raritan PDU.

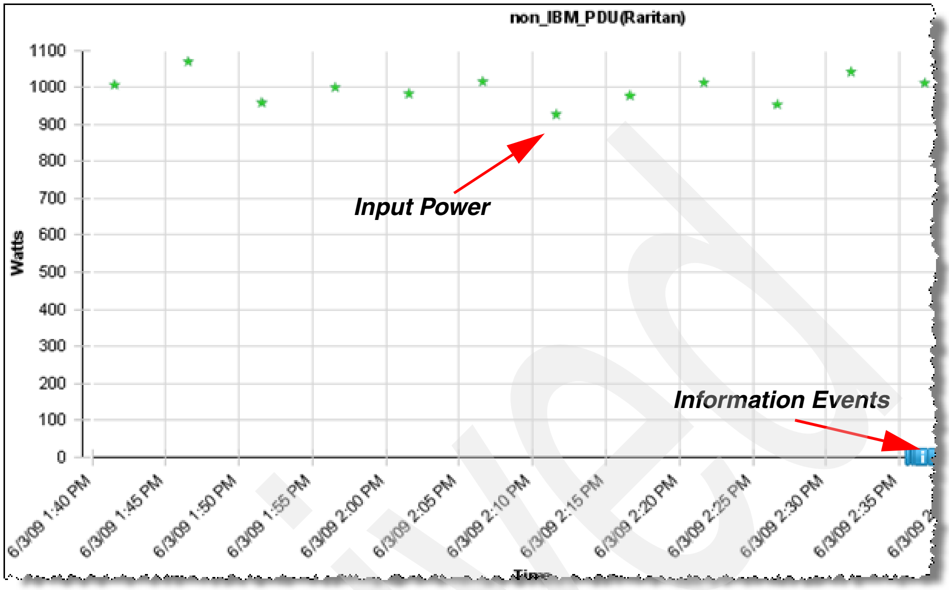


*Figure 8-54 Viewing the power-trend data for the Raritan PDU*

Figure 8-55 shows the ambient temperature and humidity for the environment in which the Raritan PDU operates.

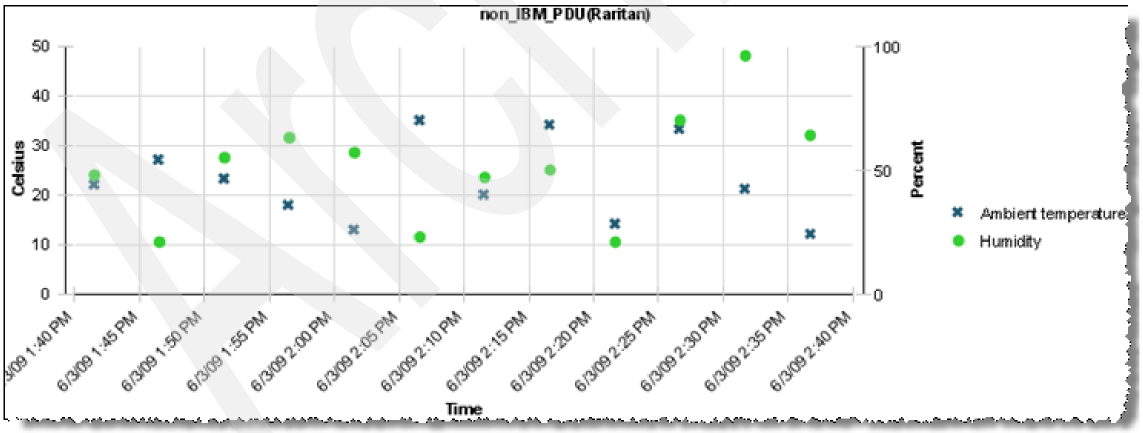


*Figure 8-55 Viewing the temperature and humidity data for the Raritan PDU*

## 8.5 Logical outlet grouping

Using Active Energy Manager, you can integrate PDU+ outlets that are separated physically into one metered resource so that you can collect power data in a power supply redundancy environment.

As shown in Figure 8-56, in the case of a Power Domain redundancy environment, each redundant power supply needs to be connected separately to each redundant PDU+. Thus, all connected outlets need to be combined to monitor the aggregate power trends from the redundant Power Supply.

Here is the summary steps of the logical PDU+ outlet grouping scenario:

1. IBM x3650 PS#1 connects to PDU+#1 Outlet Group1.

   IBM x3650 PS#2 connects to PDU+#2 Outlet Group1.
2. Create x3650 Server PS#1 metered resource to monitor PDU+#1 Outlet Group1. Create x3650 Server PS#2 metered resource to monitor PDU+#2 Outlet Group1.
3. Create x3650 Server(PS1+PS2) Group to monitor aggregated power trends from x3650 Server PS#1 and x3650 Server PS#2 and to monitor x3650 Server (PS1+PS2) Group to view the power trends of the IBM x3650.

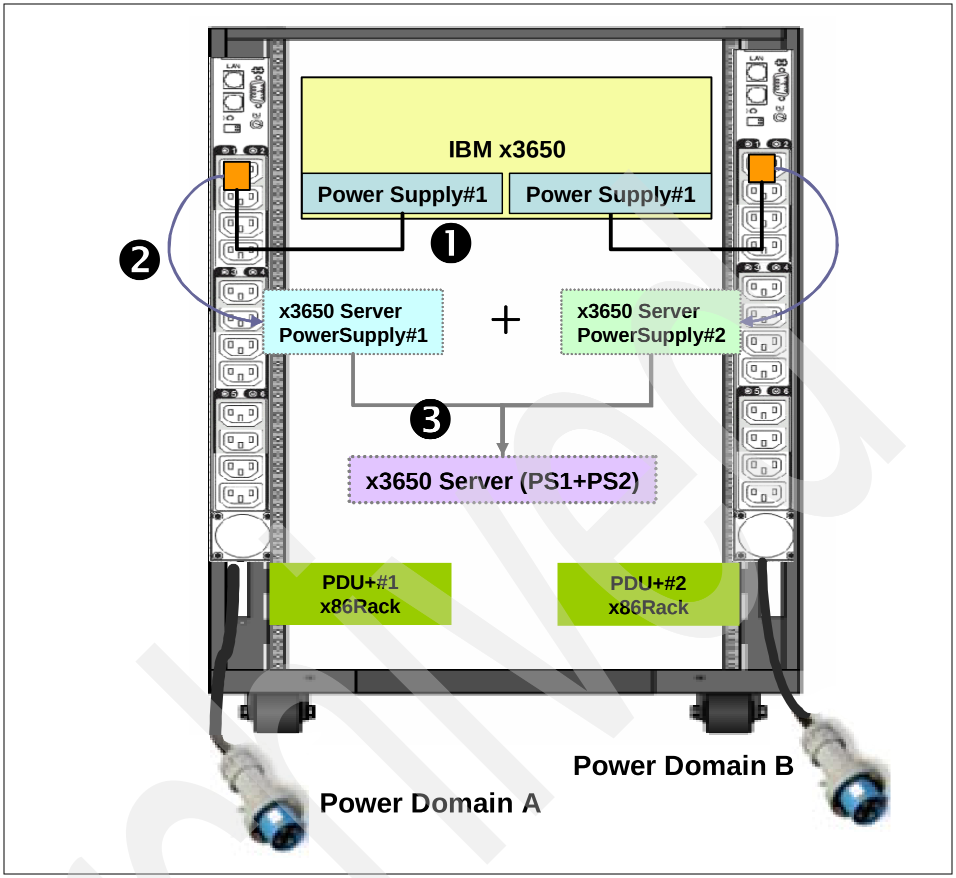


*Figure 8-56 Scenario diagram: Logical PDU+ outlet grouping*

The detailed steps for this scenario are as follows:

1. The PDU+#1_x86Rack and the PDU+#2_x86Rack are installed, as shown in Figure 8-57.

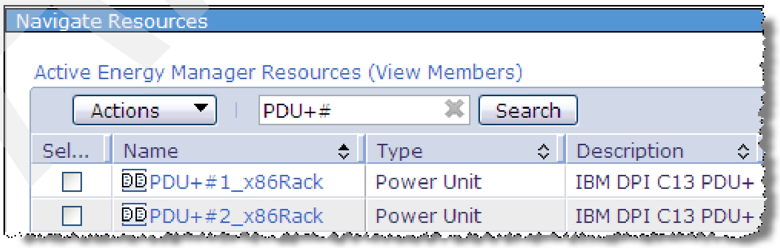


*Figure 8-57 Two PDU+ in a separated power domain environment*

2. The x3650 Server Power Supply#1 metered resource is associated with PDU+#1_x86Rack Outlet Group1, as shown in Figure 8-58. (Refer to “Assigning systems to an outlet group” on page 385 for information about how to configure a metering device.) Also, the x3650 Server Power Suppy#2 metered resource is create following the same procedure.

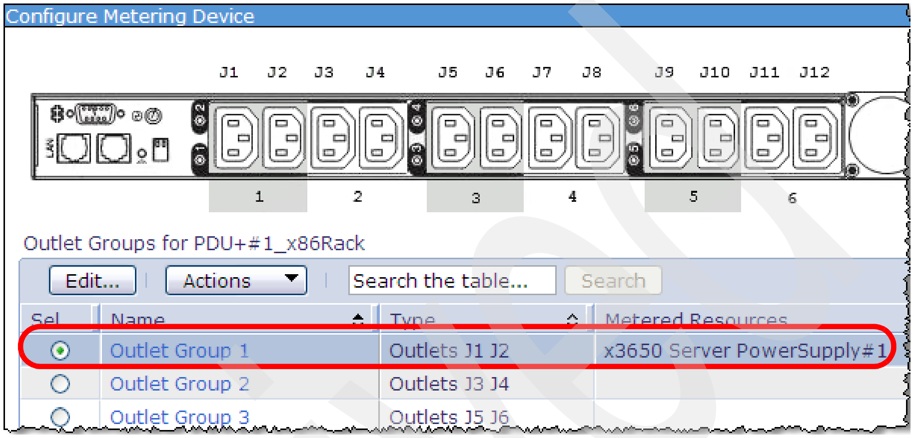


*Figure 8-58 Associating x3650 Power Supply#1 with PDU+#1 OutletGroup1*

The power-trend data for the x3650 Server PowerSupply#1 and x3650 Server PowerSupply#2 metered resource from the each connected Outlet Group displays, as shown in Figure 8-59.

> **Tip:** To add the data for the “Average Input Power (externally metered)” column for a metered resource, change the field setting. Click **Actions**, and go to **Columns**.

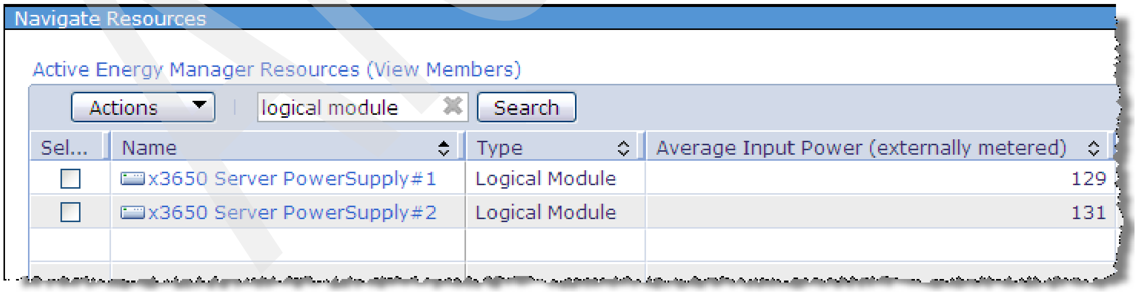


*Figure 8-59 Power-trend monitoring from each separated PDU+ outlet group.*

3. To monitor the integrated power-trend from x3650 Server PowerSupply#1 and x3650 Server PowerSupply#2, you need to combine these two different metered resources by grouping. To create the x3650 Server(PS1+PS2) Group, select **x3650 Server PowerSupply#1** and **x3650 Server PowerSupply#2**, and then **click Actions** → **Create Group**, as shown in Figure 8-60.

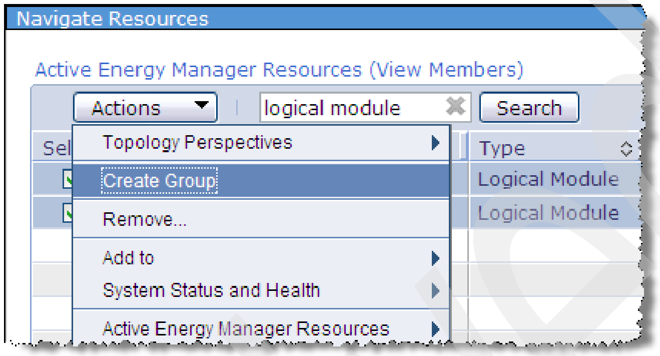


*Figure 8-60 Combine two different metered resources by grouping*

4. In the Group Editor Wizard task, enter `x3650 Server(PS1+PS2)` as shown in Figure 8-61, and click **Next.**

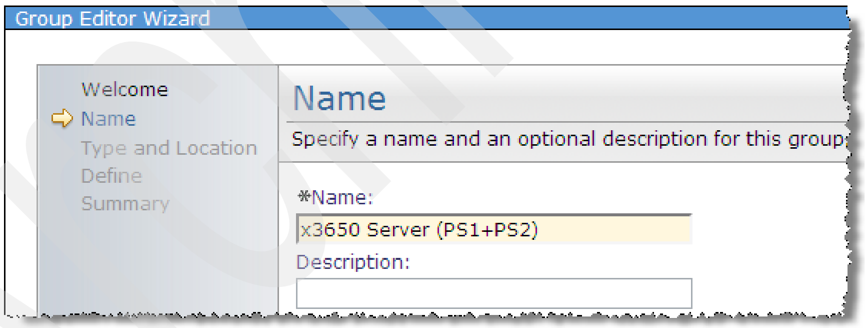


*Figure 8-61 Enter the name of group, x3650 Server (PS1+PS2)*

5. In the Type and Location panel, select **Any** as the Member Type to include any metered resources under the group. In the Group Location, select the group in which the x3650 Server(PS1+PS2) is located (in this example **x86 Server Rack#1**). Refer to Figure 8-62. Click **Next**.

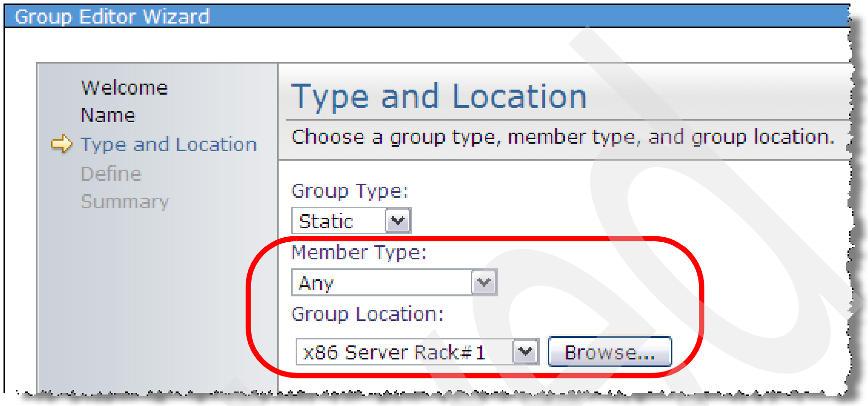


*Figure 8-62 Selecting Member Type and Group Location*

6. In the Define panel, click **Next**.
7. In the Summary panel, verify the information, and then click **Finish** to create the group (Figure 8-63).

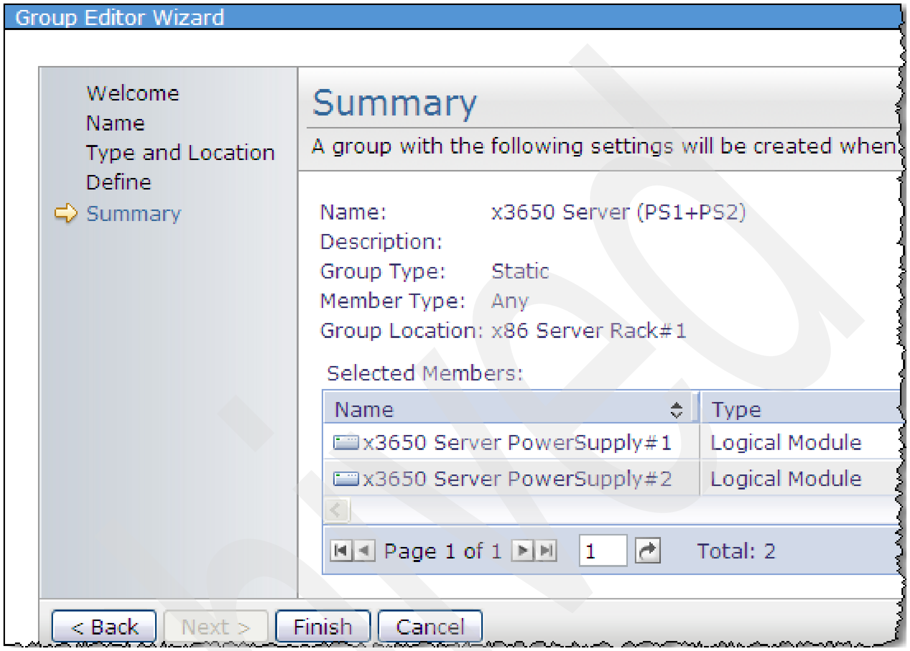


*Figure 8-63 Creating x3650 Server (PS1+PS2) summary task*

The integrated power-trend data for the x3650 Server, which is provided by a separate PDU+ Outlet Group (x3650 Server PowerSuppy#1, x3650 Server PowerSuppy#2) displays, as shown in Figure 8-64.

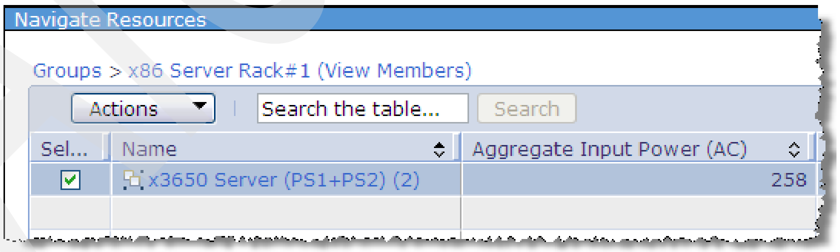


*Figure 8-64 The integrated power-trend data for the x3650 Server (PS1+PS2)*

Figure 8-65 shows the aggregated power-trend data for the x3650 Server (PS1+PS2).

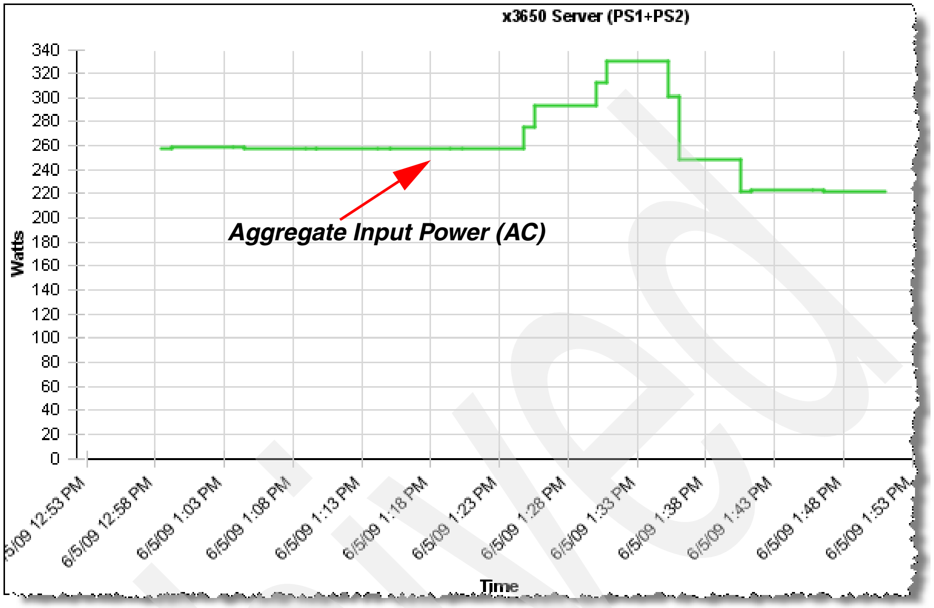


*Figure 8-65 The aggregated power-trend of x3650 Server (PS1+PS2)*

# 9

# Using sensors with Active Energy Manager

In the data center, you need to consider various environmental factors beyond monitoring power. For example, dew point and humidity are indicators of how well the air conditioning system in the data center is operating. By integrating many different types of sensors with Active Energy Manager, you can monitor environment data to reduce risk and to enhance reliability in the data center.

This chapter introduces the sensors that are available for use with Active Energy Manager and then describes how to integrate the use of those sensors into Active Energy Manager.

In this chapter, we discuss the following topics:

## 9.1 An overview of sensors

Beyond monitoring the power consumption of systems, gathering environmental information is also important to ensure the efficiency and reliability of the data center. If humidity is too high for example, a server can become more at risk of failing.

You can verify the health of a specific point in the data center or a specific server to determine whether there is risk. In addition, you can set thresholds for dew point and humidity, so that you can receive notification immediately if the value reaches a certain level. You can then choose to take appropriate action based on that notifications received.

Because Active Energy Manager can monitor power usage through the PDU+, it also can communicate with sensors and gather useful environmental information around the area in which the sensors operate.

Figure 9-1 shows a scenario in our data center that addresses how to monitor a hot zone area using temperature sensors as well as a method for connecting sensors with Active Energy Manager. The numbers in the figure represent the following information:

1. Sensors (in this example humidity and temperature) are installed throughout the data center, and sensors are connected to a gateway (or hub) as a child node using RJ-45, RJ-11, and so forth.
2. Sensor gateway communicates with Active Energy Manager over TCP/IP networks and provides environment data that is gathered from sensor nodes.
3. Active Energy Manager provides various monitoring features such as viewing data, event-viewer, configuring metering device, and topology map as Active Energy Manager does for the PDU+ and the system connected through service processors (BMC, IMM, FSP, and so forth).
4. An Active Energy Manager operator can set thresholds to monitor the hot zone area and can notify you immediately if the temperature reaches a certain level. Then, you can take an action to improve the hot zone area.

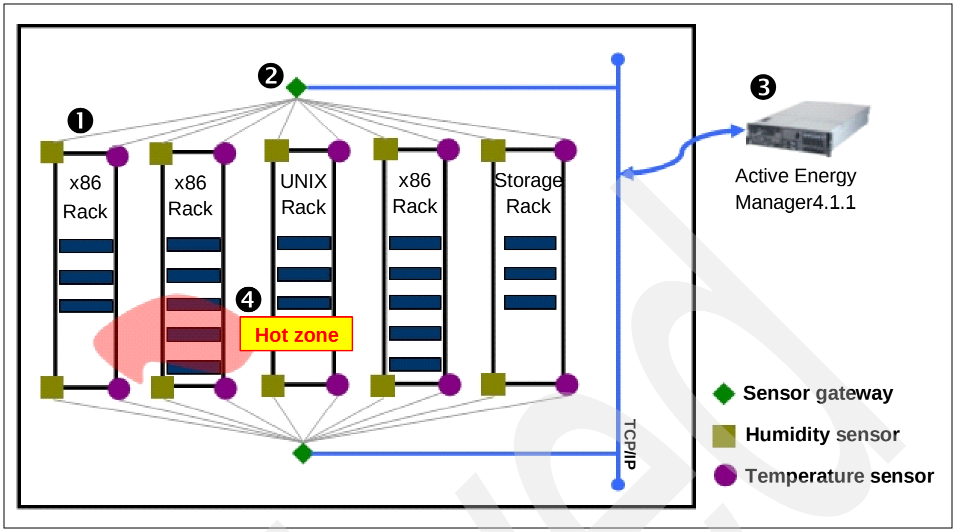


*Figure 9-1 Monitoring environmental data using sensors in data center*

## 9.2 SynapSense sensor products

SynapSense Corporation develops wireless sensor network architectures. It provides sensor network solutions, including system software and hardware for monitoring, analyzing, and collecting data in data center. For more information about SynapSense networks, see:

http://www.synapsense.com

Active Energy Manager can collect data from SynapSense sensors and combine it with many other devices. This data is then aggregated and graphed together using one single Web interface. Active Energy Manager can monitor SynapSense wireless sensor networks that are comprised of agents, gateways, sensor nodes, and sensors.

## 9.2.1 An overview of SynapSense

Figure 9-2 shows a topology of the environment that includes Active Energy Manager and a SynapSense network. Active Energy Manager communicates using SNMP v1, v2c, or v3 with the SynapSense SNMP agent. The SynapSense SNMP agent is software that runs on a separate server and collects data from the wireless sensor nodes at configured intervals.

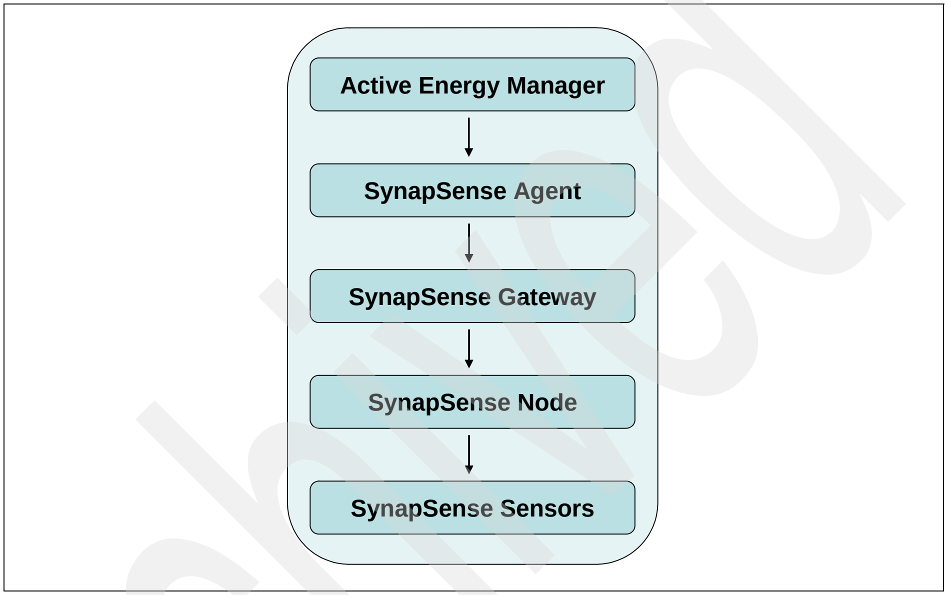


*Figure 9-2 SynapSense architecture integrated with Active Energy Manager 4.1.1*

### SynapSense terminology

Each SynapSense agent communicates with a SynapSense *gateway*, a physical piece of hardware that facilitates communication between the agent and the sensor nodes.

> **Note:** The term *gateway* is used in several different contexts when talking about Active Energy Manager supported hardware. Many devices, such as Emerson-Liebert SiteScan, refer to the device that Active Energy Manager communicates with as a *gateway*. The term *appliance* is also used in this regard. In the SynapSense modeling, Active Energy Manager communicates with the SynapSense agent.

*Sensor nodes* are wireless devices with attached sensors, purchased separately, that can be placed strategically in your data center to monitor things such as humidity, power, and temperature. You can place these nodes in proximity to each other and to a gateway node. It has internal sensors and can have external sensors plugged in to it. There are sensors node that might have a few built-in sensors and up to eight externally-plugged-in sensors. SynapSense has multiple models of nodes with various numbers of possible sensors attached.

The current sensor connects physically to a terminal on the SynapSense constellation node. It has leads that you can place around a power cord that measure the current of that cord.

The topology of the power sensor is a bit different. A SynapSense constellation node contains an internal sensor that returns a power value (valid up to the previous 60 minutes) to the gateway node. The constellation node must be wired from one of its terminals to a power meter (a separate box) that actually measures the power. The power meter is then connected to an energy metering circuit, which calculates power based on measuring current and voltage.

### Supported hardware

Active Energy Manager can monitor SynapSense sensor nodes that are connected to a SynapSense SNMP agent that is discovered by Active Energy Manager. Each sensor node can contain one or more current, power, humidity, temperature, or battery voltage sensors. SynapSense sensor networks at the following versions are supported:

- SynapSoft version 3.0.4 or later
- SynapSoft version 4.x

> **Note:** SynapSoft environments at version 2.x and 3.0.3 and earlier can also work; however, they are not supported officially in Active Energy Manager and might prevent SynapSense SNMP agents from being discovered. In these cases, upgrade to version 3.0.4 or later.

## 9.2.2 Integrating SynapSense with Active Energy Manager

In this section, we discuss how to discover SynapSense with SNMP communication and give an overview of using SynapSense.

### Discovering SynapSense by SNMP communication

To start using the SynapSense network with Active Energy Manager, you need to discover the SynapSense Agent by specifying its IP address in the Advanced System Discovery menu. For more information about discovering sensor nodes, refer to “Discovering SynapSense sensor networks using advanced system discovery” on page 162.

### Using SynapSense

The SynapSense SNMP agent associated with sensor nodes is added to IBM Systems Director. The SynapSense SNMP agent appears as a managed object but not in the Active Energy Manager resources group. Instead, all SynapSense sensor nodes that the agent knows about are displayed in this group.

> **Note:** Active Energy Manager provides support only for the following subset of SynapSense sensors:
>
> - Temperature
> - Humidity
> - Current
> - Power

SynapSense nodes report temperature at the child (individual temperature sensor) level and not at the SynapSense node level because some SynapSense nodes might not have any temperature sensors attached or built-in, while others can have multiple such sensors attached in addition to the built-in sensor. Thus, no one single temperature reading is always available to assign to the SynapSense node itself. Each sensor's readings are reported for that sensor alone.

For power sensors, the sum of the power consumption recorded at each sensor is added together and this is reported as the total power consumption at the node.

After you have discovered the SynapSense agent, it displays in **Groups** → **Groups by System Type** → **Generic Systems** in the IBM Systems Director Navigate Resources panel. Right-clicking the device, and then selecting **Properties** shows a window similar to Figure 9-3.

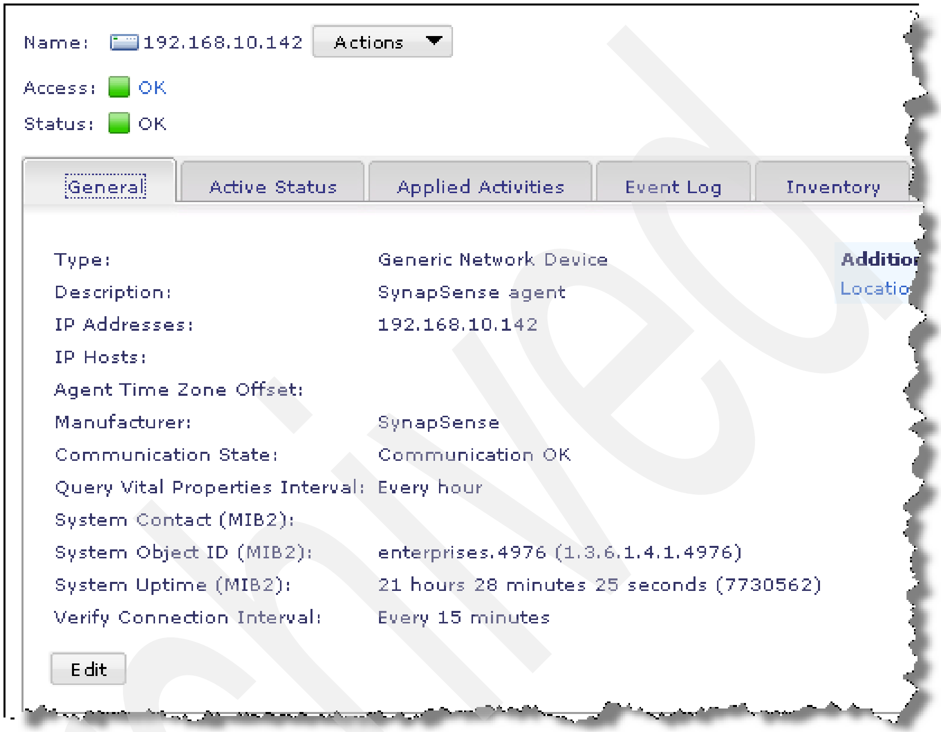


*Figure 9-3 SynapSense agent properties*

You can view the SynapSense nodes that are associated to a SynapSense agent by right-clicking an agent and selecting **Related Resource** → **Port Controller** as shown in Figure 9-4.

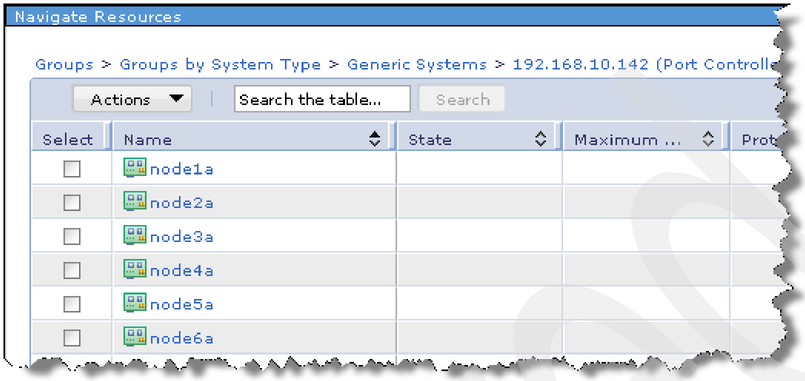


*Figure 9-4 Nodes associated to a specific SynapSense agent*

If you need to see together all the sensors from a particular agent, right-click the agent and select **Related Resource** → **Sensor**. Otherwise, the SynapSense sensors that are affiliated with a particular node can be viewed by right-clicking a sensor node and selecting **Related Resources** → **Sensor**. The view should be similar to Figure 9-5. You can obtain the value for each sensor by right-clicking the sensor and selecting **Properties** → **Active Energy Manager**.

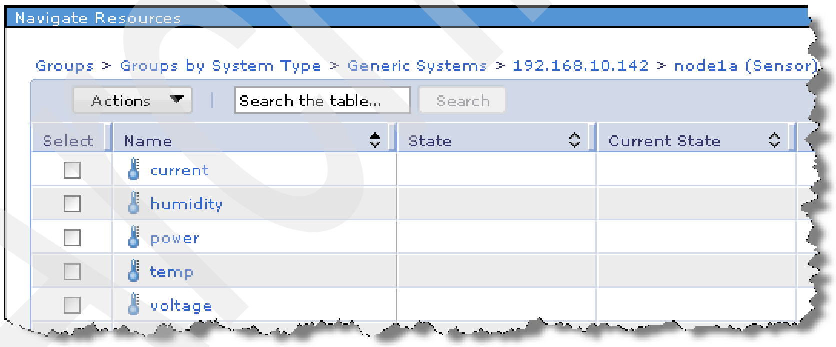


*Figure 9-5 Sensors from a specific sensor node*

**Note:** Active Energy Manager 4.1.1 can show the same information from a topology point of view. For more information, refer to 9.7.2, "Topology map" on page 451.

Figure 9-6 shows the Active Energy tab of the resulting window where you can read the current value for the selected sensor.

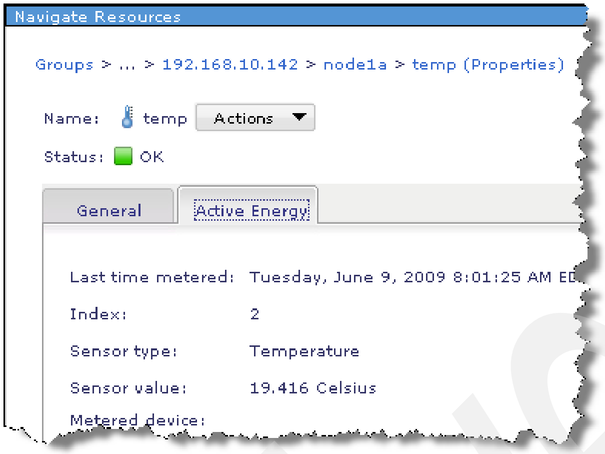


*Figure 9-6 Sensor values*

Figure 9-7 shows trend data from each sensor. You can obtain this chart by right-clicking the sensor, and selecting **Energy** → **Trend Data**.

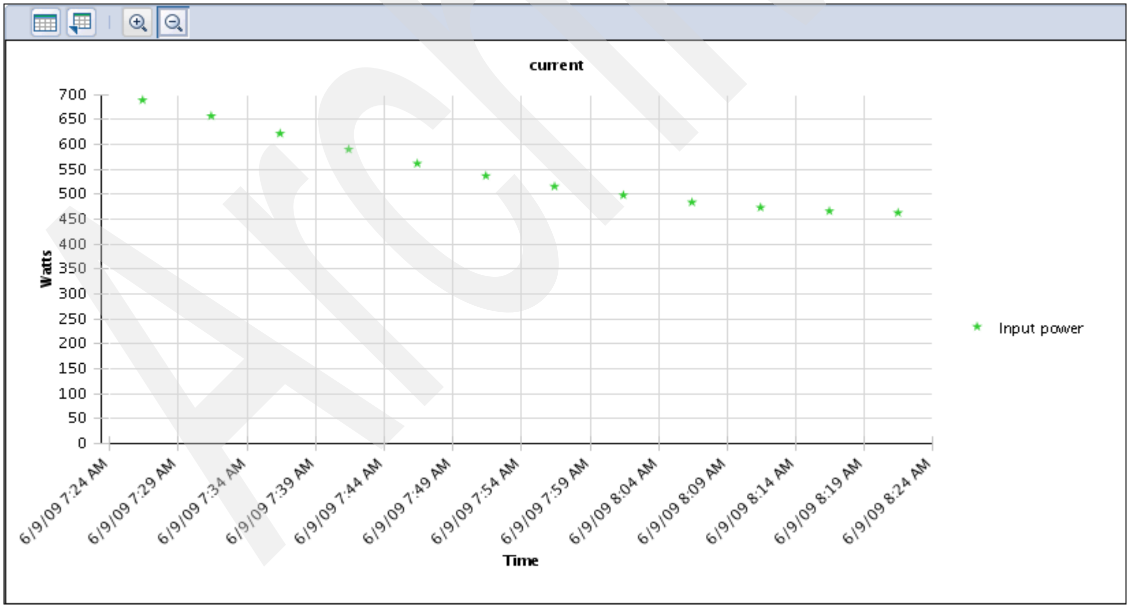


*Figure 9-7 Data from the current sensor*

**Note:** It is possible to trend data only by each sensor. Active Energy Manager does not aggregate SynapSense sensor values at the node (or higher) level.

SynapSense sensor nodes can have current, power, humidity, and temperature sensors attached. They are physically different sensors. The characteristics of each from an Active Energy Manager perspective are as follows:

- Current sensor
  - Senses and reports current in amps.
  - Uses a per-sensor-configurable voltage to convert the amps value to watts (watts = amps x voltage).
  - Trends only the watts value. There is no history for amps stored.
  - Stores only the last value for amps from a sensor, shown only with properties for that sensor on the Active Energy tab and shown as a Sensor Value.
- Power sensor
  - Senses and reports power use in kWh.
  - Active Energy Manager breaks the value into polling-interval-sized chunks and stores or trends one value per polling interval (similar to other devices).
  - Values stored are in watts.
- Battery voltage sensor:
  - Senses and reports sensor node voltage in volts.
  - Measures the life of the battery in the node and as such is not used in conjunction with either the current or power sensors.
  - Is internal to a battery-powered SynapSense sensor node.
  - Store only the last value for volts from a sensor, shown only with properties for that sensor on the Active Energy tab and shown as a Sensor Value.
  - Normally, SynapSense battery voltage sensors take two 1.5 volt batteries. In that case, the range of voltage values in theory is 0 V to 3 V. However, in practice, it is recommended that the battery voltage fall no lower than 2.55 V.

    The voltage used to calculate watts for a sensor can either be the default (configurable from **Energy** → **Active Energy Manager** by clicking **Settings** in the upper-right) or can be set to a sensor-specific value by clicking **Edit** from the Active Energy tab of the sensor's properties. Active Energy Manager does not sense voltage. A current sensor can be sensing current of multiple lines together, so this value must be customer-entered if it differs from the default.

- Humidity sensor
  - Senses and reports Humidity of the resource's surrounding environment.
  - The values are trended and expressed as a percentage.
- Temperature sensor
  - Senses and reports temperature of the resource's surrounding environment.
  - The values are trended and can be expressed in Celsius and Fahrenheit.

## 9.3 Smart Works sensor products

Active Energy Manager can also collect data from the sensors from Smart Works and supports the following products:

- Smart-Watt inline energy meters that monitor power circuits and report the energy (watt-hours) that are consumed by the attached devices. The devices can also determine the load on the circuit over any given period (watts).

  For more information, go to:

  http://www.smart-works.com/smart-watt/

- Smart-Sense environmental and temperature sensors monitor and report the temperature, humidity, and dew point levels in the data center. These sensors are typically placed in air intake points at each rack to ensure air quality.

  For more information, go to:

  http://www.smart-works.com/smart-sense/

- Smart-Net is the networking offering to connect Smart Works products either directly to a computer using a USB or serial port or onto an existing Ethernet network using a Smart-Net Gateway. The sensors are connected in daisy-chain fashion. Each device is uniquely addressable.

  You can add a new device to the network by plugging it into a device that is already connected to the network. The host software recognizes the new device automatically.

  For more information, go to:

  http://www.smart-works.com/smart-net/

### 9.3.1 Smart Works sensors architecture

Figure 9-8 illustrates the architecture for the Smart Works sensors, as follows:

**1.** Active Energy Manager discovers and communicates with the Smart Works sensors gateway using TCP/IP networks. Smart Works sensors have a built-in Web server that includes the IP address so that it can provide trend-data information to Active Energy Manager using HTTP protocol.

**2.** The Smart-Watt sensor has one physical power outlet that is connected to the Smart Works gateway.

**3.** One power supply of a system is plugged in to the Smart-Watt sensor power outlet and Smart-Watt sensor provides its power-trend information to Active Energy Manager using the Smart Works gateway.

**4.** Additional Smart-Watt sensors are connected to the first-child Smart-Watt sensor using daisy chain with an RJ-11 cable.

**5.** Smart-Sense sensors are connected to the last-child Smart-Watt sensor using daisy chain with an RJ-11 cable.

**6.** The Smart-Sense Temp/Humidity sensor, which includes three different types of sensing ability (temperature, dew point, and humidity) is shown as three separated sensors in the Active Energy Manager navigation area.

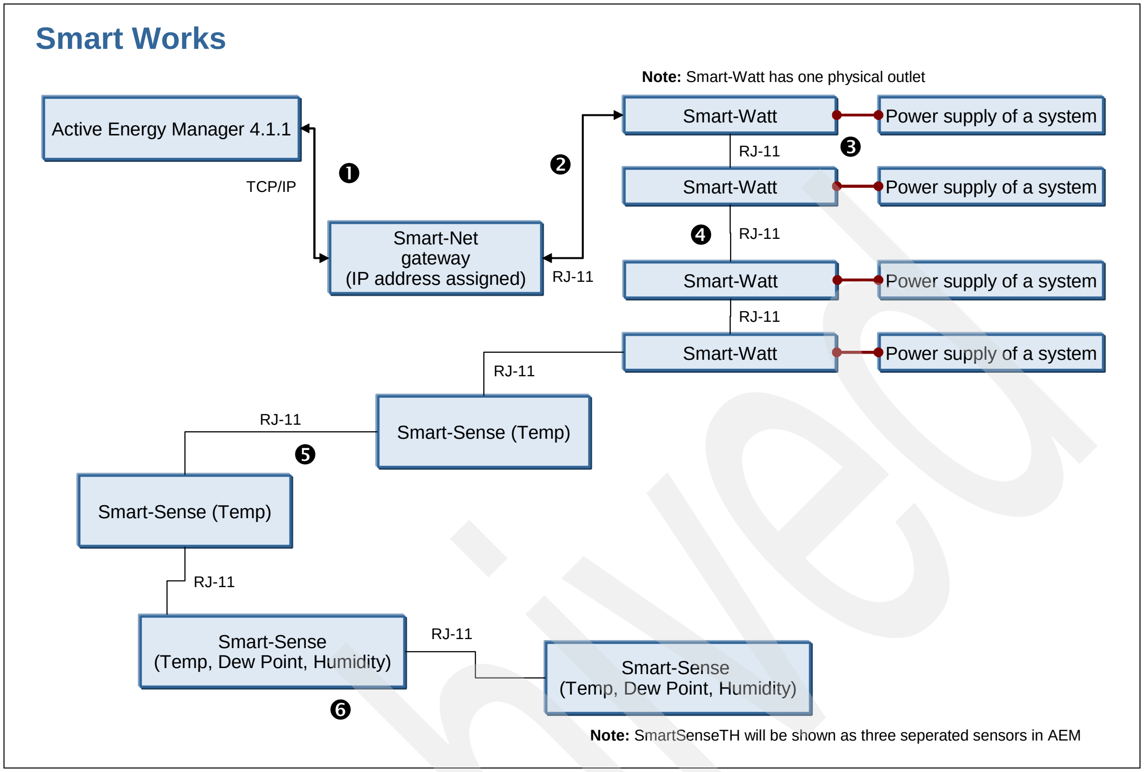


*Figure 9-8 Smart Works architecture integrated with Active Energy Manager 4.1.1*

Active Energy Manager 4.1.1 can monitor Smart Works sensor nodes that are connected to a Smart Works gateway that is discovered by Active Energy Manager. Each sensor node can contain one or more power, temperature, humidity, or dew point sensors. Active Energy Manager 4.1.1 supports the following Smart Works sensors:

- Smart-Net
- Smart-Watt
- Smart-Sense

## 9.3.2 Integrating Smart Works sensors with Active Energy Manager

The Smart Works gateway has a built-in Web server that includes the IP address so that it can be discovered by Active Energy Manager 4.1.1 through HTTP communication. For simplicity, the Smart Works gateway can be discovered using system discovery. If you want to specify the operating system as `type` and specify a protocol of `Appliance`, you can use advanced system discovery.

To discover the Smart Works gateway in IBM Systems Director manually, follow these steps:

1. Open Discovery Manager from the Systems Director welcome page, and click **System discovery** from the Common tasks box. The System Discovery panel opens, as shown in Figure 9-9.
2. Enter the IP address of the Smart Works gateway, and for the resource type, select **Operating System**. Click **Discover**.

   The Smart Works gateway displays as a type of operating system, as shown in Figure 9-9, and credentials to access the gateway are not required for the Smart Works sensors.

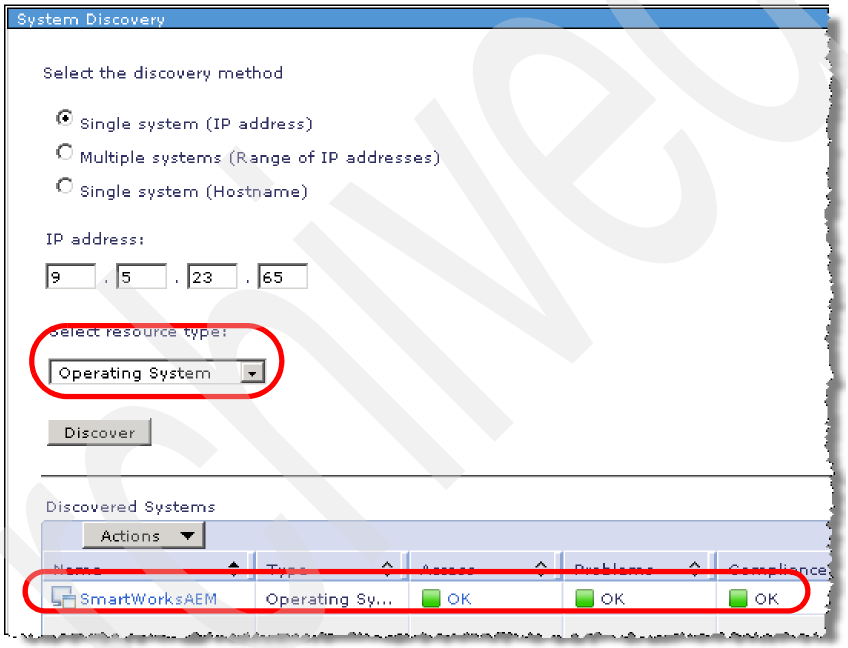


*Figure 9-9 Discovering the Smart Works gateway using system discovery*

To discover the Smart Works gateway using Advanced System Discovery, follow these steps:

1. Open Discovery Manager from the Systems Director welcome page, and click **Advanced system discovery** from the Common tasks box.
2. In the Advanced Discovery Wizard panel, enter a profile name (in this example `SmartWorksAEM`) and select **Operating System** as the System type (as shown in Figure 9-10). Then, click **Next**.

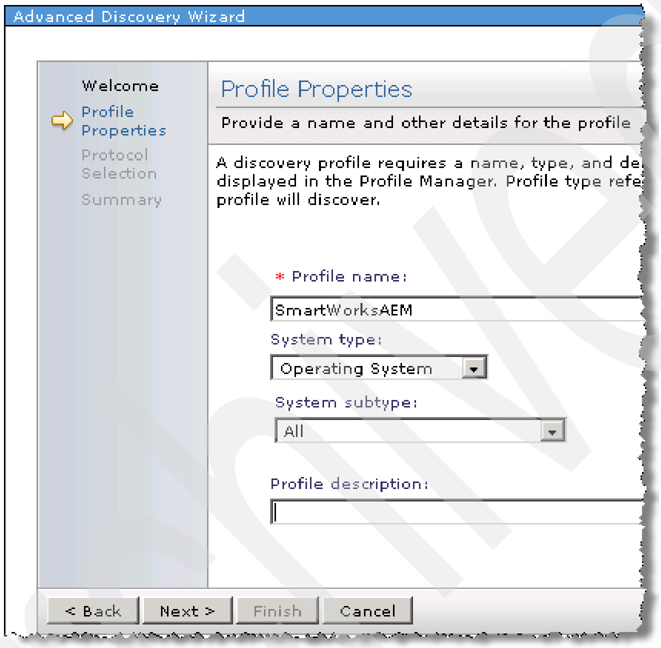


*Figure 9-10 Advanced Discovery Wizard: Profile Properties*

3. You can select many different types of protocols. Select **Appliance Discovery** for the Smart Works gateway, as shown in Figure 9-11, which communicates using HTTP. Then, click **Next**.

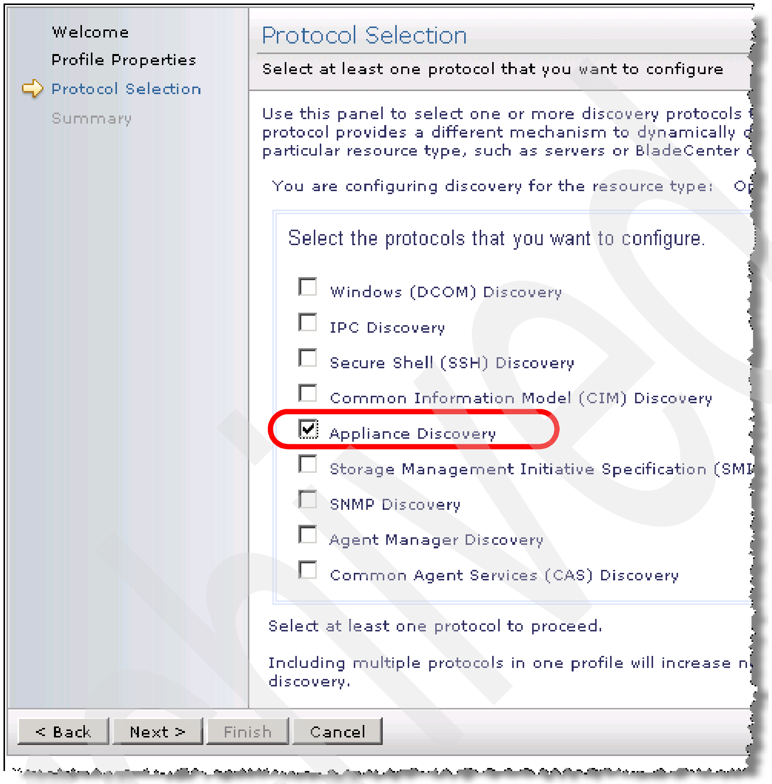


*Figure 9-11 Advanced Discovery Wizard: Protocol Selection*

4. Enter the IP address of the Smart Works gateway in the “Single IP address or beginning range” text box. Then, click **Add** to include that IP address, as shown in Figure 9-12. Click **Next** to continue.

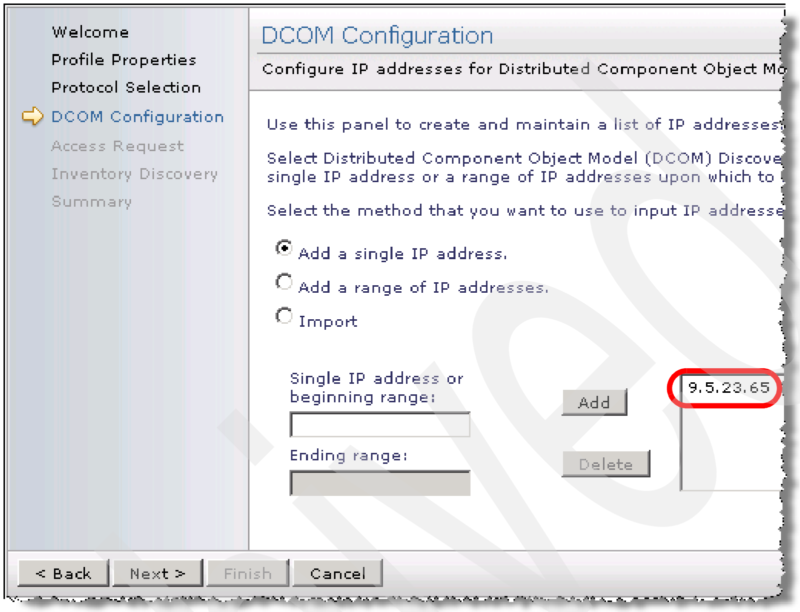


*Figure 9-12 Advanced Discovery Wizard: DCOM Configuration*

5. Because the Smart Works gateway is not required for credential information, select **Deactivate**, and click **Next** as shown in Figure 9-13.

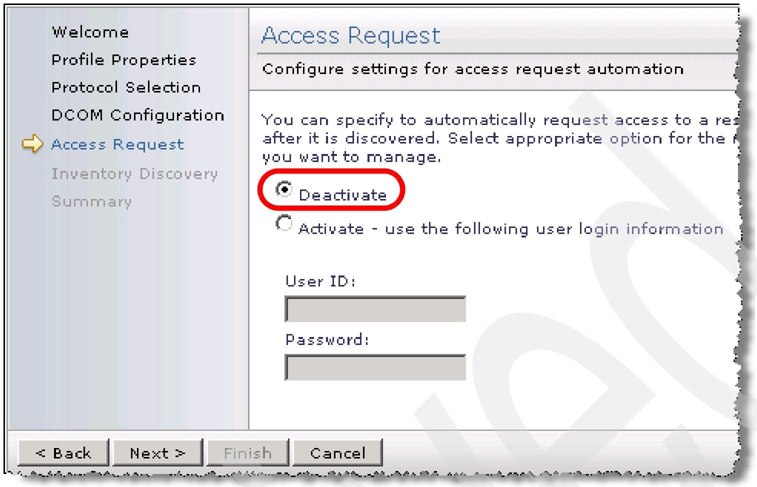


*Figure 9-13 Advanced Discovery Wizard: Access Request*

6. In the Inventory Discovery panel, select **Deactivate**, and click **Next** as shown in Figure 9-14.

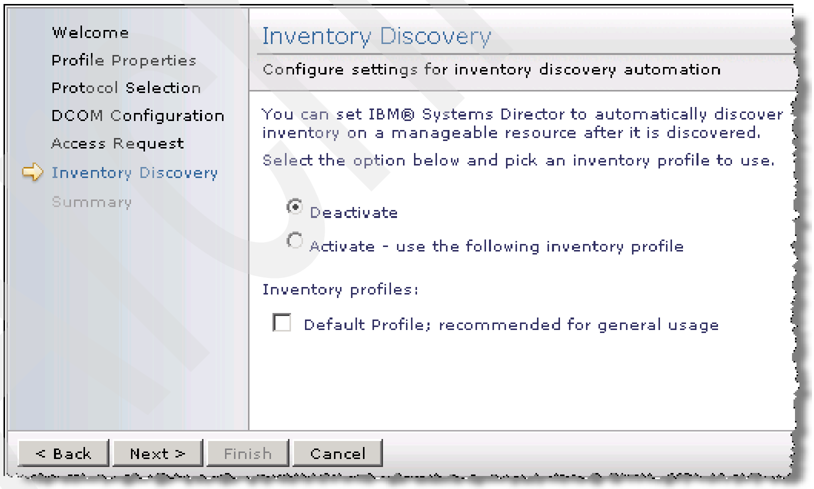


*Figure 9-14 Advanced Discovery Wizard: Inventory Discovery*

7. In the Summary task, click **Finish** after verifying the information.

8. When the steps of advanced system discovery complete, the Profile (in this example *SmartWorksAEM*) is created as shown in Figure 9-15. Click **Run** to discover the Smart Works gateway using the SmartWorksAEM profile.

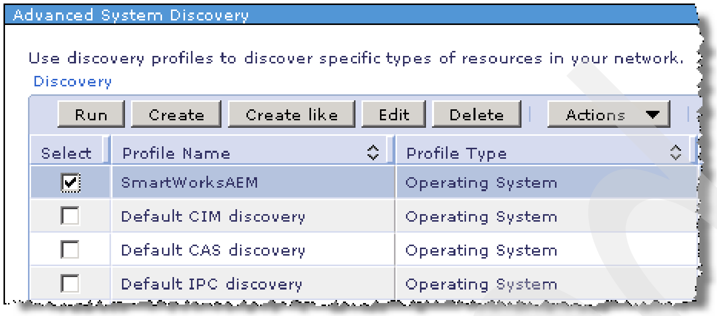


*Figure 9-15 Running Smart Works advanced system discovery profile*

The Smart Works gateway displays in the navigation area, as shown in Figure 9-16.


*Figure 9-16 The Smart Works gateway is displayed in the navigation area*

## Using Smart Works sensors

After the Smart Works gateway is discovered successfully, the Smart Works Node and Smart-Watt sensors are shown in the Active Energy Manager Resources navigation area, as shown in Figure 9-17 on page 436.

> **Note:** Because most of the Active Energy Manager features to use the Smart Works sensors are the same as with other sensors, we cover this information in this section in more detail in 9.7, “Working with sensors in Active Energy Manager” on page 448.

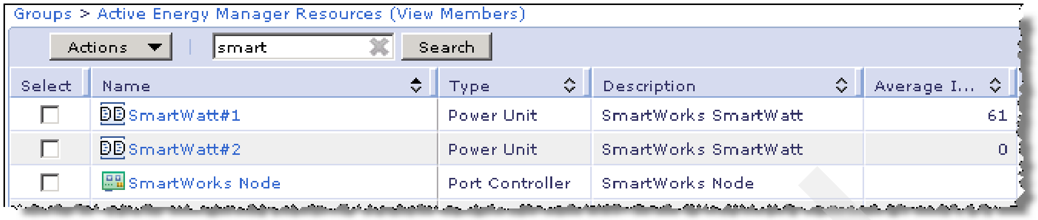


*Figure 9-17 Smart Works Node and Smart-Watt in the Active Energy Manager Resource area*

### *Configuring Smart Works Node*

A Smart-Sense Sensor reports temperature data and a Smart-SenseTH sensor reports temperature, humidity, and dew point data. However, because Active Energy Manager can report only one piece of data for each sensor, the Smart-SenseTH sensor is split into three logical sensors, as shown in Figure 9-18.

> **Note:** You configure metered resources using the same method as configuring a PDU+. For more information, refer to for 8.3, “Configuring the PDU+” on page 381.

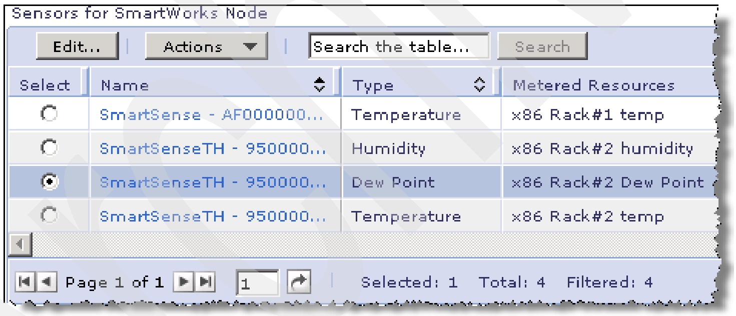


*Figure 9-18 Configuring Smart Works sensor node*

### DIsplaying the Sensor Value data for sensors

In the Smart Works Node (Sensor) navigation area, the trend data provided by different types of sensors (such as the temperature, dew point, and humidity data) displays in the “Sensor Value” column, as shown in Figure 9-19. However, this column is not shown by default. To see this data, click **Actions** → **Columns** and customize the table.

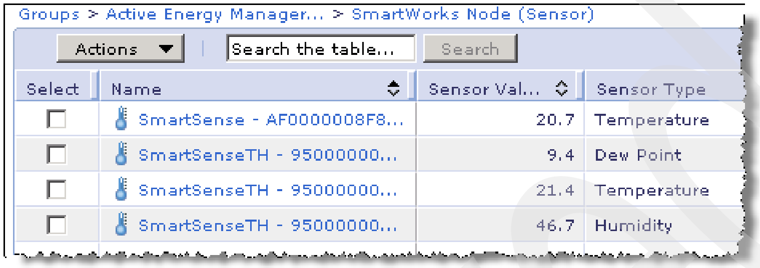


*Figure 9-19 Sensor Value data*

### Configuring the Smart-Watt sensor

You configure the Smart-Watt sensor using a similar method to configuring the PDU+. The Smart-Watt sensor is much like a small PDU+ that provides one outlet. Therefore, by associating a logical module with the Smart-Watt sensor, you can monitor power trends for a system.

To configure the Smart-Watt sensor, right-click the name of Smart-Watt sensor (in this example **SmartWatt#1**), and then click **Energy** → **Configure Metering Device**, as shown in Figure 9-20.

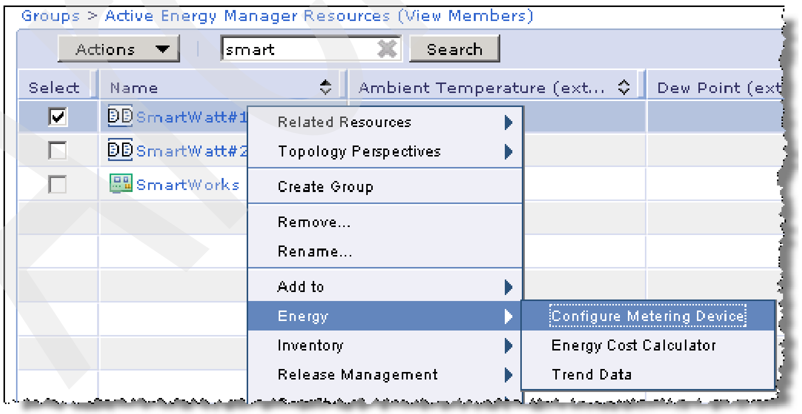


*Figure 9-20 Configuring the Smart-Watt sensor*

The IBM DS3200 Storage is associated with the Smart-Watt sensor as a logical device, as shown in Figure 9-21.

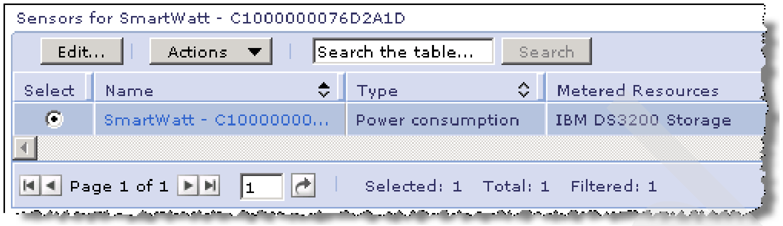


*Figure 9-21 Associating a system with Smart-Watt sensor*

# 9.4 iButtonLink sensor products

Active Energy Manager can also collect data from the sensors from iButtonLink. This solution is compatible with the 1-Wire networking industry standard. The products supported include:

- LinkHub-E

  The LinkHub is a gateway device for sensors. It is designed to operate on long and short 1-Wire busses. It has four inputs for sensors (sensors can be daisy chained together to each of the four inputs) and connects the sensors to an existing Ethernet network. The gateway is assigned an IP address and can be accessed by Active Energy Manager.

  For more information, go to:

  http://www.ibuttonlink.com/linkhube.aspx

- T-Sense

  A small temperature sensor that provides fast and accurate temperature readings. The device has two RJ45 connectors which accepts cat5 (RJ45 connectors) or telephone (RJ11/12 connectors) wiring.

  For more information, go to:

  http://www.ibuttonlink.com/t-sense.aspx

- T-Probe

  The T-Probe is a temperature sensor integrated on the end of a 10-ft cable. The other end of the cable is an RJ45 connector for connectivity to the LinkHub gateway.

  For more information, go to:

  http://www.ibuttonlink.com/tprobe.aspx

- T-String

  The T-String is a temperature sensor that is designed to be installed in server racks and networking cabinets. The T-String has a total length of 8 feet. It has six temperature sensors, spaced 1 foot apart starting 3 feet from the RJ45 connector. The spacing and length are optimized for a 72-inch rack. Each sensor is uniquely addressable.

  For more information, go to:

  http://www.ibuttonlink.com/tstring.aspx

- MS-T

  The MS-T is a temperature sensor with RJ-45 connectors.

  For more information, go to:

  http://www.ibuttonlink.com/ms-t.aspx

- MS-TH

  The MS-TH measures both temperature and humidity.

  For more information, go to:

  http://www.ibuttonlink.com/ms-th.aspx

## Architecture

Figure 9-22 illustrates the LinkHub-E architecture when integrated with Active Energy Manager.

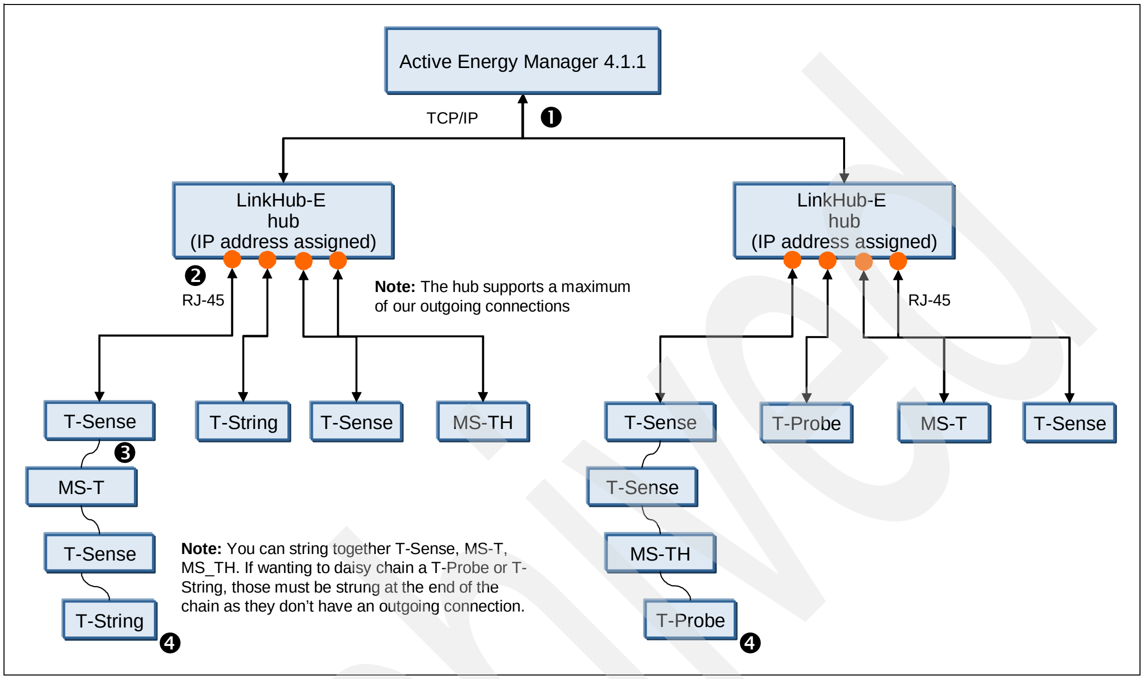


*Figure 9-22 LinkHub-E architecture integrated with Active Energy Manager 4.1.1*

In the figure, the numbers correspond to the following:

1. Active Energy Manager discovers and communicates with the LinkHub-E hub using TCP/IP. LinkHub-E provides trend-data information to Active Energy Manager through a Telnet connection.
2. Up to four “first child” sensors are connected to LinkHub-E hub using the RJ-45 cable.
3. All other sensors are strung together to the first child sensor using daisy chain.
4. T-String and T-Probe sensors are connected to LinkHub-E hub as a “last child.” If they are connected directly to the hub, then no other sensors can go beyond that probe or string.

## 9.4.1 Integrating LinkHub-E sensors with Active Energy Manager

The LinkHub-E sensor does not have a built-in Web server. The LinkHub-E sensor communicates with Active Energy Manager 4.1.1 using Telnet. Although LinkHub-E is communicating internally with Active Energy Manager using Telnet, Active Energy Manager can discover the LinkHub-E sensor in the same manner as discovering the Smart Works sensors, which communicates using HTTP.

### Discovery

The method to discover LinkHub-E and its attached sensors is the same as SynapSense sensor networks. Refer to 9.2.2, “Integrating SynapSense with Active Energy Manager” on page 421 for instructions.

### Using the LinkHub-E sensors

Most of Active Energy Manager features for the LinkHub-E sensors are same as features for other sensors. Refer to 9.7, “Working with sensors in Active Energy Manager” on page 448 for more information.

The following features are specific to the LinkHub-E:

- The number of sensors limitation

  If you see a sensor value of -1 (as highlighted in Figure 9-23), the sensor is located in the daisy chain where it does not receive enough power. This is a limitation noted by the manufacturer. To correct this issue, reduce the number of sensors within your daisy chain.

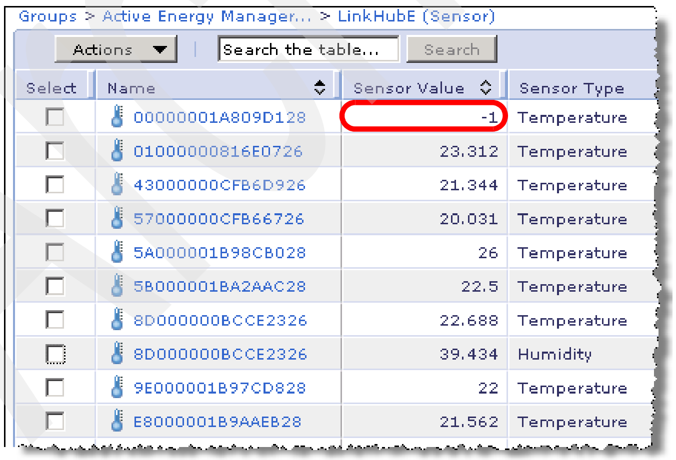


| Select | Name | Sensor Value | Sensor Type |
|---|---|---|---|
| | 00000001A809D128 | -1 | Temperature |
| | 01000000816E0726 | 23.312 | Temperature |
| | 43000000CFB6D926 | 21.344 | Temperature |
| | 57000000CFB66726 | 20.031 | Temperature |
| | 5A000001B98CB028 | 26 | Temperature |
| | 5B000001BA2AAC28 | 22.5 | Temperature |
| | 8D000000BCCE2326 | 22.688 | Temperature |
| | 8D000000BCCE2326 | 39.434 | Humidity |
| | 9E000001B97CD828 | 22 | Temperature |
| | E8000001B9AAEB28 | 21.562 | Temperature |

*Figure 9-23 Sensors in the LinkHub-E resource group*

- The default sensor name

  The name that is assigned to the individual sensor is the identified number that is returned during searches and shown in windows, such as Figure 9-23 on page 441. The name of sensor is not changeable. We recommend that you associate a logical module with each sensor to recognize by the name of logical module.

## 9.5 Embedded Data Systems sensor products

Active Energy Manager can also collect data from the sensors from Embedded Data Systems. This solution is compatible with the 1-Wire networking industry standard. The products supported include:

- HA7Net Ethernet 1-Wire Host Adapter

  The HA7Net is a gateway device connecting the sensor network to your existing Ethernet network. The device has three 1-Wire ports.

  For more information, go to:

  http://www.embeddeddatasystems.com/HA7Net--Ethernet-1-Wire-Host-Adapter_p_22.html

- Sensors and probes

  Active Energy Manager 4.1.1 can monitor HA7Net sensor nodes that are connected to a HA7Net sensor that is discovered by Active Energy Manager. Each sensor node can contain one or more temperature or humidity sensors.

  Active Energy Manager supports many 1-wire compatible sensors:

  – Temperature sensors: T-Sense, T-String, MS-T
  – Temperature probes: T-Probe
  – Temperature and humidity sensors: MS-TH

## Architecture

Figure 9-24 illustrates the HA7Net sensor architecture when integrated with Active Energy Manager.

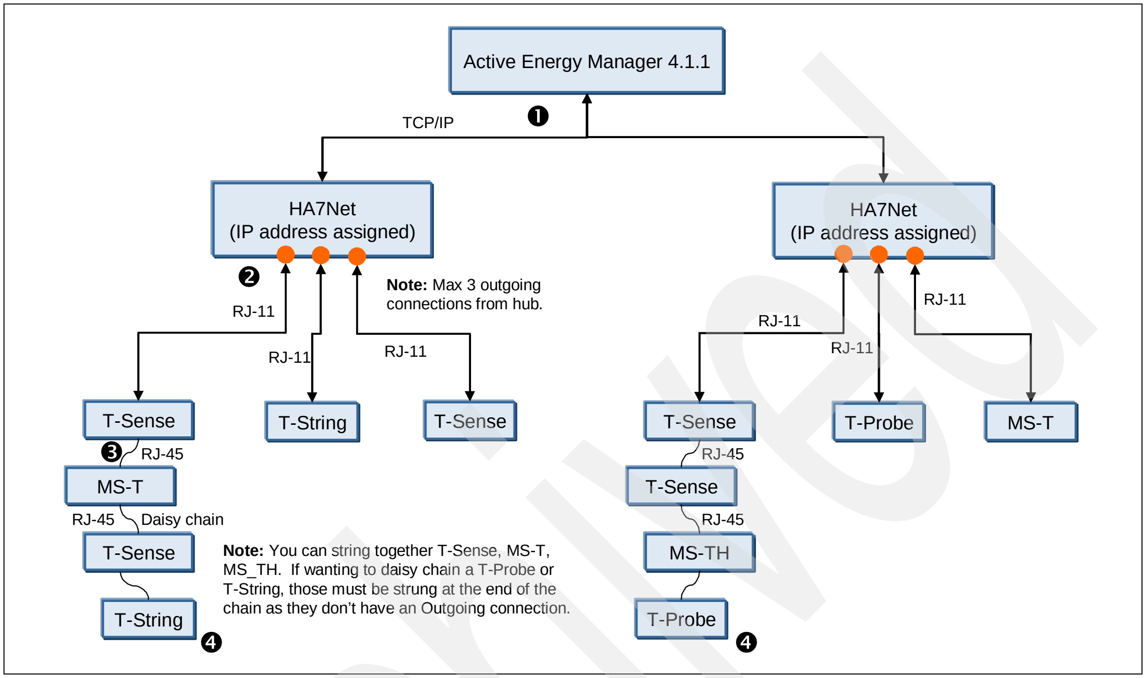


*Figure 9-24 HA7Net architecture integrated with Active Energy Manager 4.1.1*

In the figure, the numbers correspond to the following information:

1. Active Energy Manager discovers and communicates with the HA7Net sensor using TCP/IP networks. The HA7Net has a built-in Web server that includes the IP address so that it can provide trend-data information to Active Energy Manager by HTTP protocol.
2. Up to three first child sensors are connected to HA7Net sensor using a RJ-11 cable.
3. All other sensors are strung together to the first child sensor using daisy chain.
4. The T-String and T-Probe sensors are connected to HA7Net sensor as a last child. If they are directly connected to the sensor, then no other sensors can go beyond that probe or string.

## 9.5.1 Integrating the HA7Net sensor with Active Energy Manager

The HA7Net sensor has a built-in Web server that includes the IP address so that it can be discovered by Active Energy Manager 4.1.1 using HTTP communication. For simplicity, you can discover the HA7Net gateway using system discovery. If you want to specify the operating system as `type` and specify a protocol of `appliance`, you can use advanced system discovery.

You discover the HA7Net sensor in the same manner as discovering the Smart Works sensors, except authentication is required. (Refer to 9.2.2, “Integrating SynapSense with Active Energy Manager” on page 421 for more information.)

HA7Net requires request access for authentication using the Web Services credentials. As shown in Figure 9-25, the HA7Net sensor access information is indicated as “No access” by default. To get an authentication, right-click **No access** and select **Request Access**.

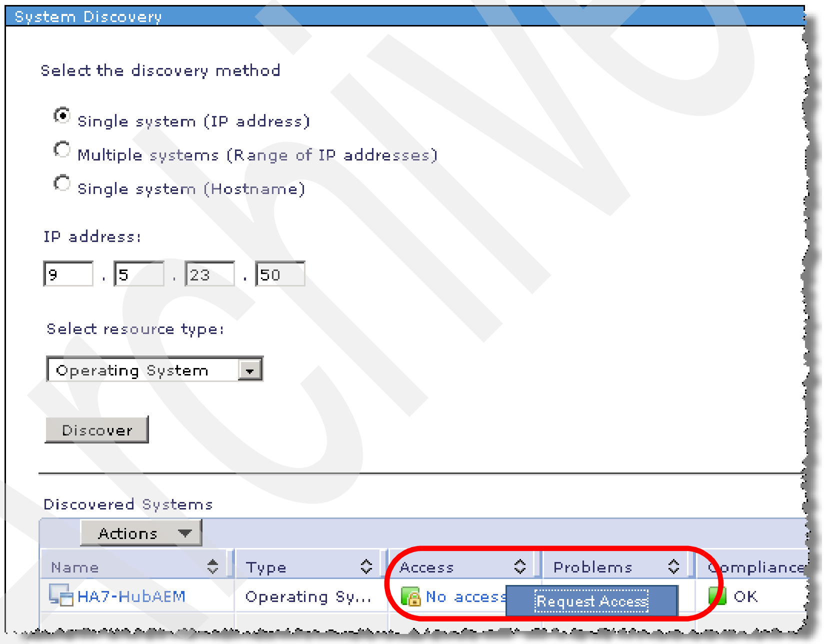


*Figure 9-25 Request Access for HA7Net Web server in system discovery*

In case of advanced system discovery, you need to specify Web Services authentication information in the Access Request wizard (as shown in Figure 9-26). Select **Activate** and enter the credential information. Then, click **Next**.

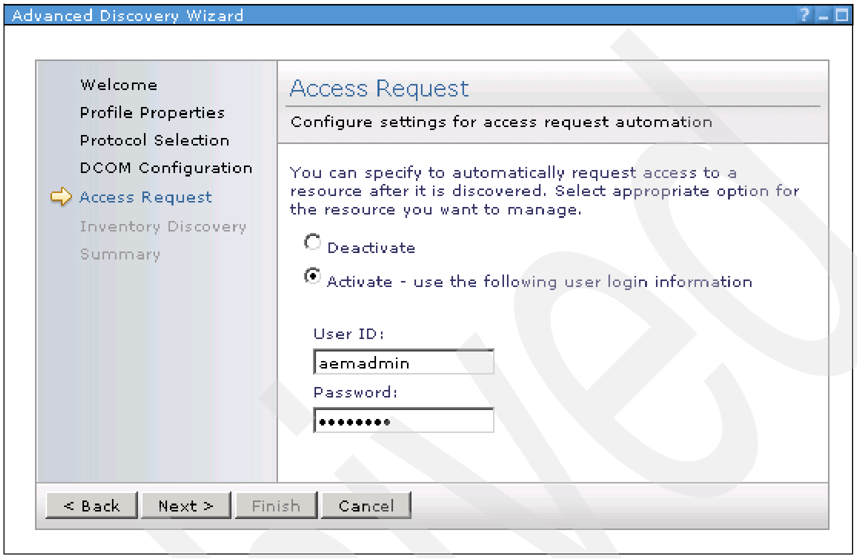


*Figure 9-26 Access Request for the HA7Net Web server in advanced system discovery*

## Using the HA7Net sensor

Most of Active Energy Manager features to use the HA7Net sensor are the same as features for other sensors. For more information, see 9.7, "Working with sensors in Active Energy Manager" on page 448.

Figure 9-27 shows an example of several temperature sensors that are connected to an HA7Net gateway.

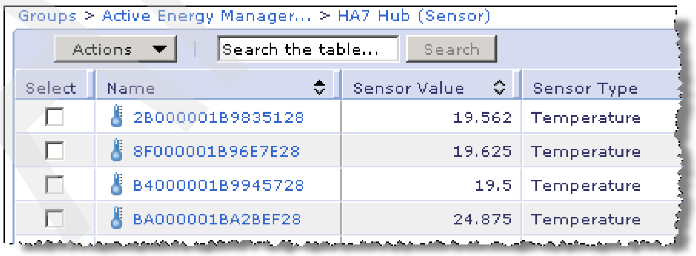


*Figure 9-27 Temperature sensors connected to the HA7Net sensor*

## 9.6 Sensatronics sensor products

Active Energy Manager can also collect data from the sensors from Sensatronics and supports the following products:

- E16 Temperature Monitor

  The E16 lets you connect up to 16 temperature probes to your existing Ethernet network. After you configure the static address of the device using the serial port, you can then manage the device and connected probes using the monitor's built-in Web interface. You can also manage the monitor using SNMP and MIB files that are available for download. The four-port E4 is also available.

  For more information, go to:

  http://www.sensatronics.com/index.php/industrial-monitors/model-e16.html

- U16 Universal Temperature Monitor

  The U16 monitor is similar to the E16 but also supports probe types other than standard temperature sensing, including power presence sensors, magnetic door sensors, dry contact sensors, and high and cryo-low temperature sensing. The four-port U4 is also available.

  For more information, go to:

  http://www.sensatronics.com/index.php/industrial-monitors/model-u16.html

- EM1

  The EM1 is similar in features to the E and U monitor models. It offers four ports and is designed to monitor temperature, humidity, and wetness. It also supports magnetic door sensors and dry contact sensors.

  For more information, go to:

  http://www.sensatronics.com/index.php/industrial-monitors/model-em1.html

- Senturion

  The Senturion is a complete rack-mounted environmental monitoring solution that includes on-board temperature, humidity, airflow, and light level sensors coupled with eight external ports for all types of additional probes. Senturion also supports connecting up to four external camera. The device includes a color LCD screen providing local temperature and humidity information. The device has a built-in Web server for management.

For more information about this sensor, go to:

http://www.sensatronics.com/index.php/it-monitors/senturion.html

- Probes

  Probes are available to connect to the monitors and gateways. Various cable lengths are available from 10 feet to 300 feet (longer lengths are also orderable). Cables are 2-conductor, 22 AWG, and PVC jacketed.

## Architecture

Figure 9-28 illustrates the Sensatronics sensor architecture when integrated with Active Energy Manager. In this example, U16 monitors are used.

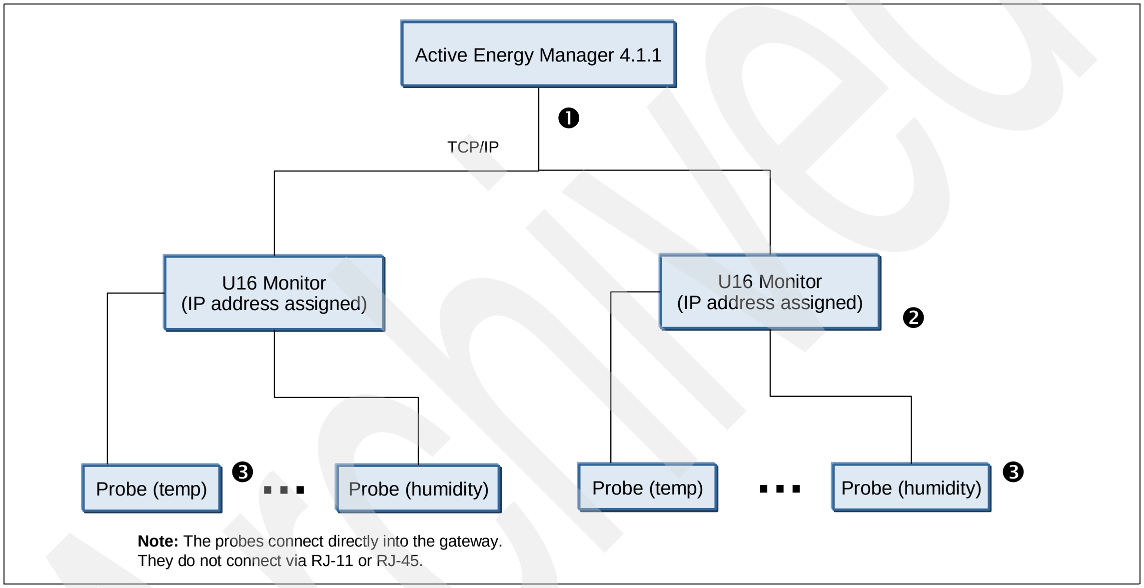


*Figure 9-28 Example Sensatronics architecture integrated with Active Energy Manager 4.1.1*

In this figure, the numbers correspond to the following information:

1. Active Energy Manager discovers and communicates with the Sensatronics gateway or model E, U, or EM monitors using TCP/IP. These monitors each have a built-in Web server that provide trend-data information to Active Energy Manager using HTTP.
2. The number of probes and the types of probes that can be connected depends on the monitor that you use.
3. Probes are all connected directly to the monitor. They are not daisy chained together like 1-wire solutions.

### 9.6.1 Integrating Sensatronics sensors with Active Energy Manager

All Sensatronics monitors have a built-in Web server can be discovered by Active Energy Manager 4.1.1 through HTTP communication. For simplicity, you can discover the Sensatronics sensor using system discovery. If you want to specify the operating system as `type` and specify a protocol of `appliance`, you can use advanced system discovery.

The method for discovering a Sensatronics sensor is the same as the method for discovering a SynapSense sensor networks. For more information, see 9.2.2, "Integrating SynapSense with Active Energy Manager" on page 421.

#### Using Sensatronics sensors

Most of Active Energy Manager features to use Sensatronics are same as the features for other sensors. For more information, see 9.7, "Working with sensors in Active Energy Manager" on page 448.

Figure 9-29 indicates the temperature probe sensor that is connected to Sensatronics gateway.

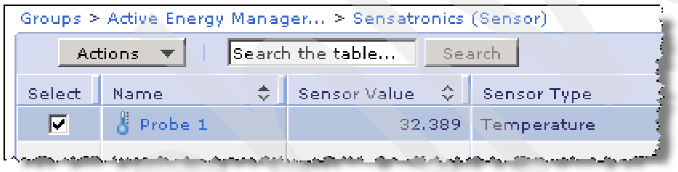


*Figure 9-29 The temperature probe connected to Sensatronics sensor*

## 9.7 Working with sensors in Active Energy Manager

The sensors from different vendors are explored by Active Energy Manager in a similar way. In this section, we cover how to view power and environmental data, the event viewer, and the topology map with sensors.

## 9.7.1 Power and environmental data

This section shows how to view power and environmental data in Active Energy Manager.

### Viewing sensor trend-data in the navigation area

Figure 9-30 shows some logical modules that are associated with different types of sensors, providing trend-data such as temperature, dew point, and humidity.

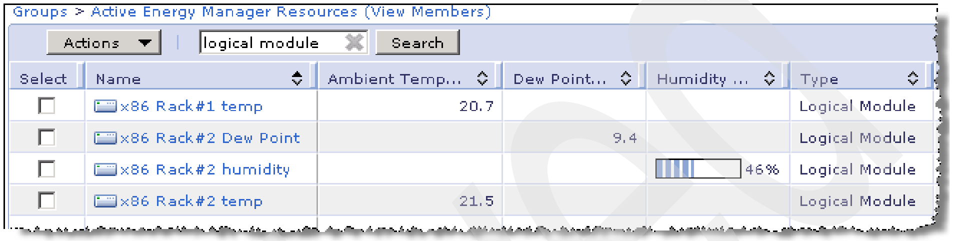


*Figure 9-30 Viewing sensor trend-data in the navigation area*

In the Active Energy Manager Resource area, you can customize field sets to view environmental trend data by changing the columns. Refer to for 8.3, “Configuring the PDU+” on page 381 more information about this procedure.

### Viewing sensor trend data in a chart

To view sensor trend data in a chart, follow these steps:

1. Right-click the name of sensor logical module (in this example **x86 Rack#2 Dew Point**) and then select **Energy** → **Trend Data**, as shown in Figure 9-31.

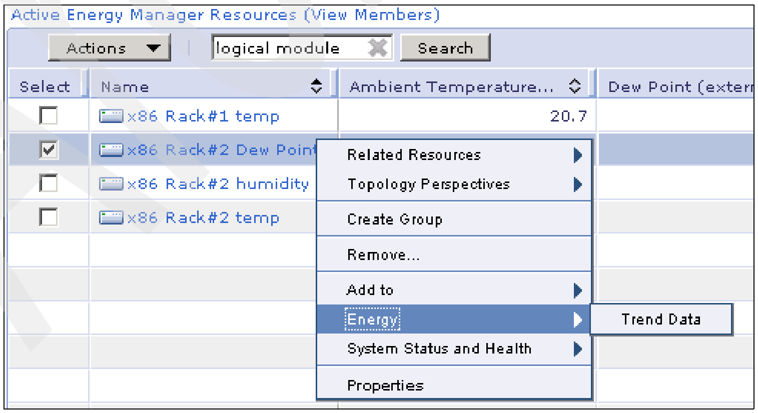


*Figure 9-31 Viewing sensor trend-data in a chart*

The Dew Point for the x86 Rack#2 area displays as a chart, as shown in Figure 9-32 on page 450. In case of Dew Point, refer to percent range at the right side of the chart.

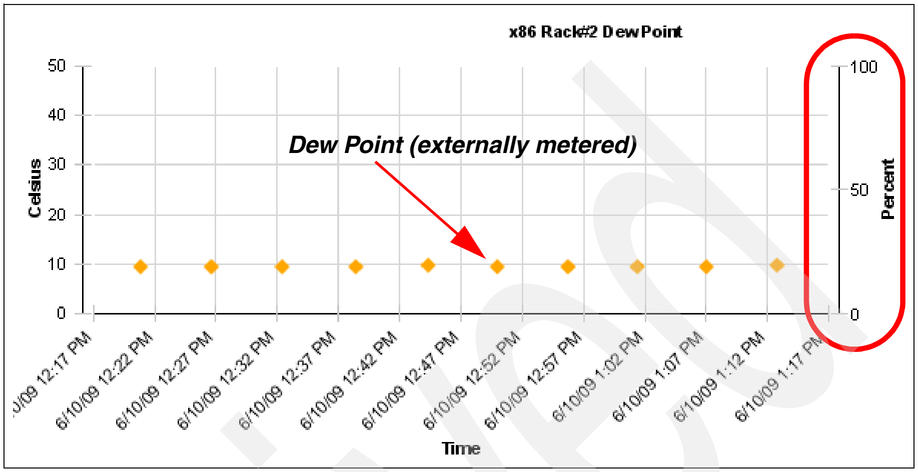


*Figure 9-32 Viewing Dew Point trend-data around the x86 Rack#2 area*

2. By default, the trend data view does not show Dew Point automatically. To add this value, click **Options** in the Trend Data view, as shown in Figure 9-33.

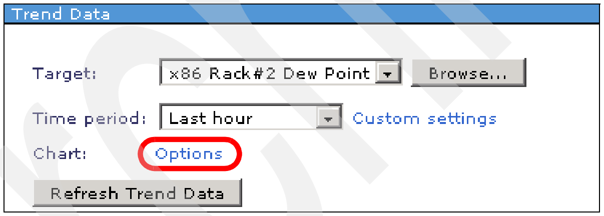


*Figure 9-33 Customizing chart options in Trend Data*

3. Select the “Dew point range” and “Dew point (externally metered)” options in the Environmental and CPU Chart section, as shown in Figure 9-34 on page 451. Then, click **OK** to continue.

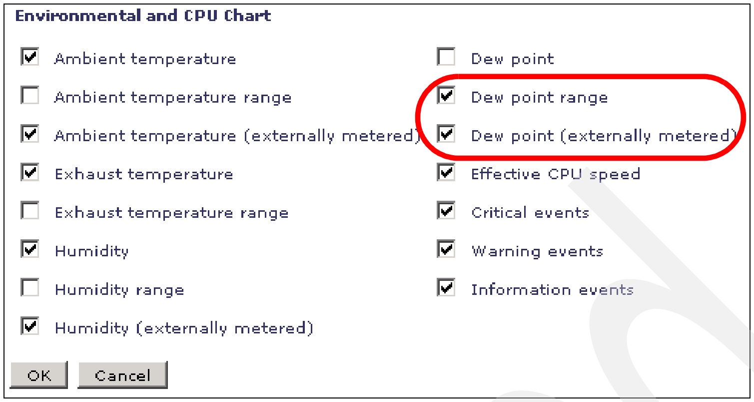


*Figure 9-34 Environmental and CPU Chart section*

## 9.7.2 Topology map

To view the topology, right-click a sensor in the navigation area (in this example **node1a**). Then, select **Topology Perspectives** → **Active Energy**.

For example, Figure 9-35 shows sensors that associated with the SynapSense node as a topology map. You can monitor the event status of each sensor and node using the Status Items icons. To monitor trend information from the topology map view, right-click the sensor and then select **Energy** → **Trend data**.

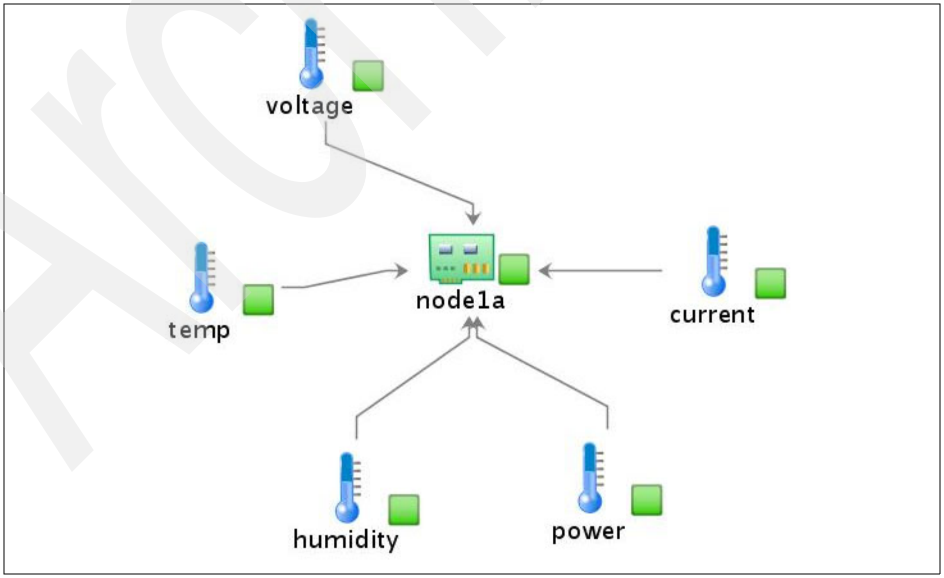


*Figure 9-35 Topology view of each SynapSense Node*

10

# Integrating facility managers with Active Energy Manager

This chapter provides information about how to integrate various facility managers with Active Energy Manager. It includes the following topics:

## 10.1 Overview

Components hosted inside a typical data center can be broadly classified into two categories:

- IT equipment, which includes servers, storage, and network components
- Infrastructure equipment, including:
  - Chillers
  - Computer room air conditioner (CRAC) units
  - Air removal units (ARUs)
  - Environmental monitoring units (EMUs)
  - Power distribution units (PDUs)
  - Automatic transfer switches (ATSes)
  - Uninterruptible power supply units

Active Energy Manager extends the scope of energy monitoring and management by integrating with various facility endpoints to enable a more complete view of energy consumption within the data center. Integrating facility managers with IBM Active Energy Manager enables IT administrators to be alerted about issues with facilities equipment, such as overheating, low battery power on uninterruptible power supplies, or other conditions that might keep IT equipment in a data center from running properly.

Because PDUs are also monitored and measured within the facility manager applications, Active Energy Manager has additional sources of data for measuring power consumption of IT equipment. Active Energy Manager allows the IT administrator to obtain a more accurate and complete view of the power consumption of the servers in the entire data center.

Active Energy Manager 4.1.1 integrates with the following facility managers to discover and manage all power and cooling equipment in the data center:

- Emerson Liebert SiteScan (SSWEB) Version 3.0

  Active Energy Manager can monitor PDUs, uninterruptible power supplies and CRAC units which are monitored and controlled by Emerson Liebert SiteScan (SSWEB) Version 3.0. For more information, see 10.2, “Liebert SiteScan Web” on page 456.

- Eaton Power Xpert Reporting System

  Active Energy Manager can monitor PDUs and uninterruptible power supplies that are defined in an Eaton Power Xpert Reporting Hierarchy database server. For more information, see 10.3, “Eaton Power Xpert” on page 462.

- APC InfraStruXure Central

  Active Energy Manager can monitor data center infrastructure components that are managed by an InfraStruXure Central server. These metering devices can contain one or more current, power, humidity, temperature, or dew point sensors.

  For more information, see 10.4, “APC InfraStruXure Central” on page 472.

Integrating a facility manager with Active Energy Manager provides the following advantages:

- Because data centers rarely contain equipment from just one vendor, integrating Active Energy Manager with various facility managers can provide a more holistic view of data center as well as real-time monitoring of the health status.
- Active Energy Manager integration allows IT personnel to be alerted about any problems with facility equipment inside the data center. For example, instead of waiting for an alert from the facility manager about a data center overheating due to the failure of an air conditioning unit, the IT administrator can determine the issue immediately using Active Energy Manager and can take proactive measures before servers become too hot. This is a large step toward making a view of all data center equipment available.
- The power and thermal data trending capability of Active Energy Manager can be a substantial advantage to IT administrators as well as facility managers. Because Active Energy Manager monitors both IT and facility equipment power and thermal values, the trend graphs of IT and facility equipment over a period of time can be analyzed to arrive at right sizing of the cooling infrastructure for the data center.
- Active Energy Manager allows IT administrators to associate IT equipment to facility equipment, such as power, thermal, and humidity sensors. This feature can help both IT administrators and facility administrators in understanding the happenings in the data center.
- Active Energy Manager can exploit status and monitoring features in IBM Systems Director, which allows IT and facility personnel to monitor the health of IT and facility equipment. The Topology feature in IBM Systems Director provides association maps between IT and facility equipment, which can help determine the following information:

  – How the data center components are associated together

  – How IT equipment is affected if a malfunction or failure of any power and cooling facility equipment occurs

  Resource thresholds can be set on IT and facility equipment inside the data center so that Active Energy Manager can generate an event or alarm that indicates when these thresholds are exceeded.

When critical events occur inside the data center, the IT or facility administrator can be alerted immediately. Useful scenarios for these alerts include:

– When the battery charge values on monitored uninterruptible power supply equipment falls below the thresholds, the facility manager can be alerted about the status of the battery charge.
– If a failure in the cooling system of the data center occurs, IT personnel can be alerted immediately, which allows them to take precautionary measures.
– When monitored IT equipment inlet temperature rises above the threshold, an alert can be sent to the IT and facility administrator with the resource names of IT equipment and the location of that equipment inside the data center. In this scenario, the facility administrator can increase the amount of cooling provided to the location where the respective IT equipment is located if needed.

# 10.2 Liebert SiteScan Web

This section provides an overview of the Liebert SiteScan facility manager from Emerson Network Power as well as information about how to discover and use the solution with Active Energy Manager.

## 10.2.1 Overview of Emerson Network Power and Liebert integration

Liebert SiteScan Web is a facility management software application from Emerson Network Power that monitors and manages the data center infrastructure equipment, including data center cooling equipment, power supply, and power monitoring equipment. It captures energy consumption from the data center infrastructure equipment and provides analysis and reporting for enhanced availability and efficiency.

Integrating Active Energy Manager with Liebert SiteScan Web enables IT administrators to be alerted about issues with facilities equipment. With Liebert SiteScan Web, you can use a Web-based application to gather information, change operating parameters, run reports, and perform similar functions on various types of critical equipment.

For Active Energy Manager to monitor Emerson Network Power and Liebert facility equipment in a data center, you must use Liebert SiteScan version 3.0 or 4.0. Active Energy Manager makes use of the Web services that are defined by this version to get environmental monitoring values as well as events and alert notifications. Refer to 10.2.2, “Supported facility equipment” on page 458 for the list of equipment that Active Energy Manager supports.

Figure 10-1 depicts Liebert SiteScan Web integration with Active Energy Manager 4.1.1.

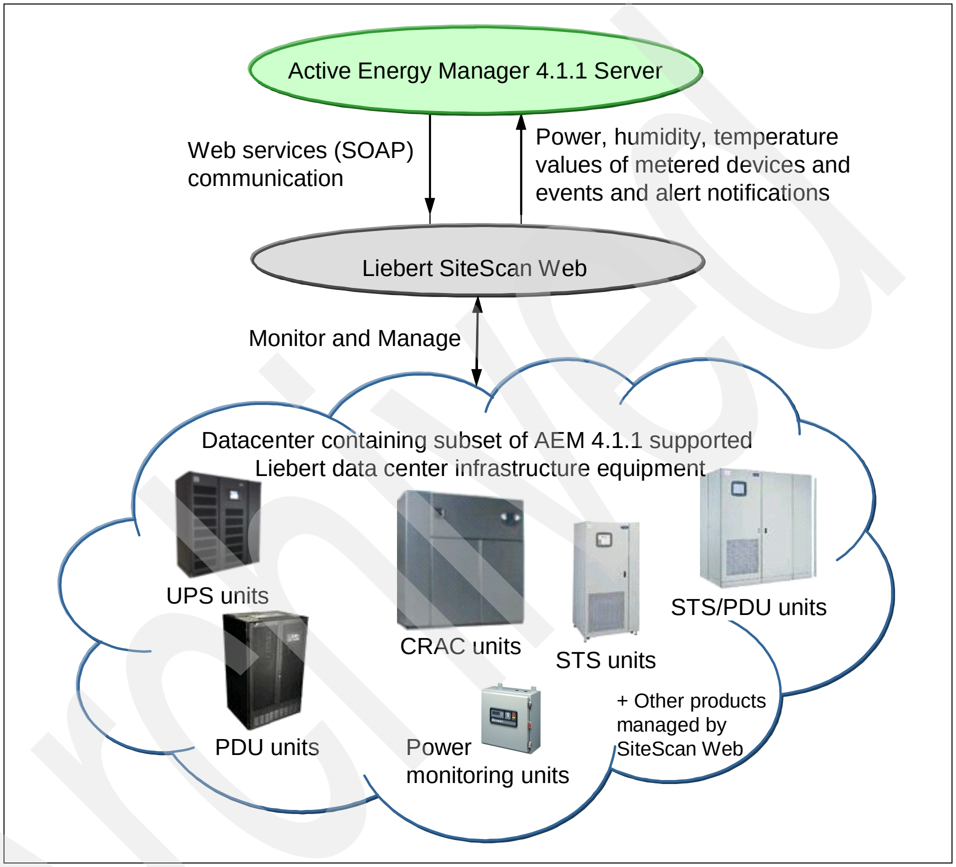


*Figure 10-1 Liebert SiteScan integration with Active Energy Manager*

To discover all Active Energy Manager 4.1.1 supported Emerson Network Power and Liebert facility equipment, you must first discover the facility manager and then authenticate to it using SiteScan Web credentials. Active Energy Manager communicates with SiteScan Web using Web services (SOAP) to get information about all supported facility equipment that is monitored by the SiteScan Web management software. Active Energy Manager polls SiteScan Web every 5 minutes, by default, to retrieve environmental values such as temperature, humidity, and power for the supported facility equipment.

To add and work with Emerson Network Power and Liebert devices in Active Energy Manager, the user ID used for accessing SiteScan Web must have “Remote Data Access - SOAP” functional privilege.

For more information, about how to discover Emerson Network Power and Liebert SiteScan Web facility equipment in Active Energy Manager, refer to 10.2.3, “Discovering Liebert SiteScan Web-managed facility equipment in Active Energy Manager” on page 459.

For more information about the tasks performed by Active Energy Manager on Emerson Network Power and Liebert facility equipment, refer to 10.5, “Active Energy Manager monitoring tasks” on page 477.

## 10.2.2 Supported facility equipment

Active Energy Manager v4.1.1 can retrieve the power and environmental values as well as the event and alert notifications of the following facility equipment inside the data center:

- Emerson Network Power and Liebert power units
  - Breaker Cabinet Power Monitoring (FPC)
  - Power Monitoring Panel (Ext. Protocol) (PM2)
  - Power Monitoring Panel (PMP)
  - Static Transfer Switch PDU Dual Output (STS)
  - Static Transfer Switch PDU Dual Output (STS-2)
  - Static Transfer Switch PDU (EDS)
  - Voltage-Current-Frequency Monitor Panel (VCF)
  - Voltage-Current Monitoring Panel (VCM)
- Emerson Network Power and Liebert uninterruptible power supplies:
  - Multi Module UPS – S600/610 Extended Prot. – MM4
  - Multi Module Series – MMS
  - System Control Cabinet – S600600/610 Ext. Prot. – SC4
  - System Control Cabinet – SCC
  - Single Module Series – SMS
  - Single Module UPS – S600600/610 Ext. Prot. – SM4
  - Multi-Module SICE 7200 and HiPulse UPS – SMM
  - Systems Cabinet SICE 7200 UPS – SSC
  - Single Module SICE 7200 and HiPulse UPS – SSM
  - UPStation S3 – US3
  - Single Module Series AP301/302 – SM3
  - Single Module UPS – NPower—IMP
  - Single Module UPS – NX—PNX

   - HiNet – PHN
   - NFinity – PNF
   - GXT – PGX
   - PSI – PPS
- Computer room air conditioning (CRAC) units

  Emerson Network Power and Liebert hardware that is monitored and controlled by Liebert SiteScan

> **Note:** Active Energy Manager can retrieve the power and environmental values as well as the events and alert notifications of these facility equipment in the data center *only* if the facility equipment are managed by SiteScan Web version 3.0. Support for SiteScan Web 4.0 is planned for a future version of Active Energy Manager.

### 10.2.3 Discovering Liebert SiteScan Web-managed facility equipment in Active Energy Manager

You can discover facility equipment managed by SiteScan Web by following these steps:

1. In the IBM Systems Director navigation area, click to expand **Inventory**, and then click **System Discovery**.
2. On the System Discovery page, specify the IP address of the operating system on which Liebert SiteScan Web is located.

3. In the Select resource type field, select **Operating System** from the list. Then, select **Discover** as shown in Figure 10-2.

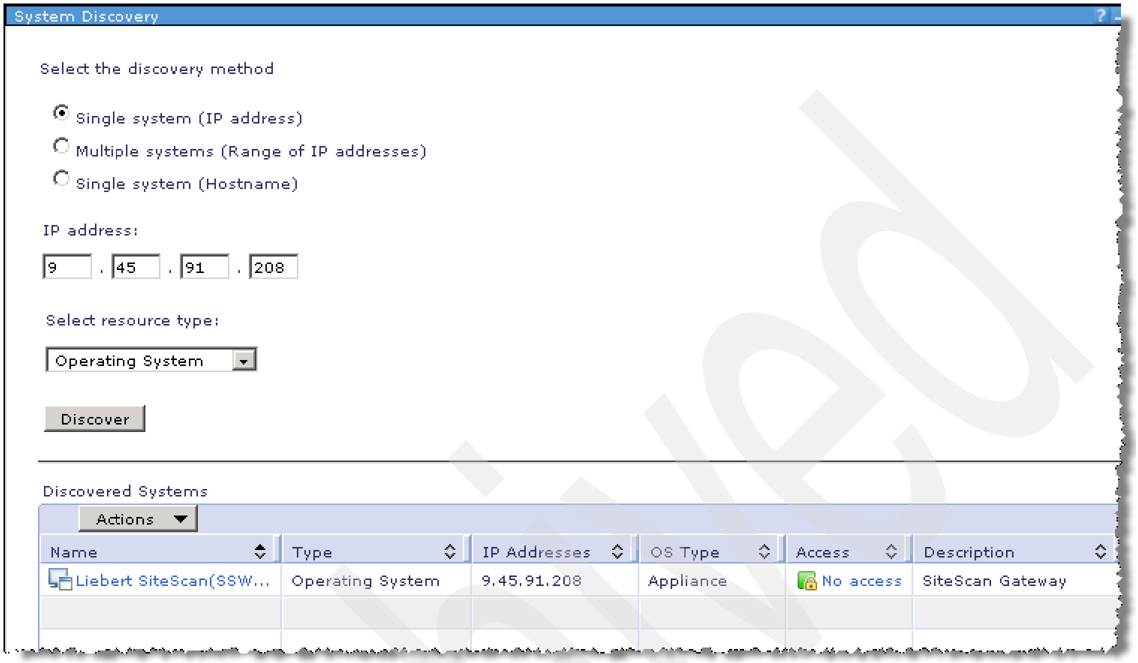


*Figure 10-2 Discovery of Liebert SiteScan Web management server in IBM Systems Director*

4. After the operating system is discovered, right-click the operating system, and select **Security** → **Request Access**.

5. On the Request Access page, enter the Liebert SiteScan Web credentials, and then select **Request Access**. Click **Close** after the access is granted, as shown in Figure 10-3.

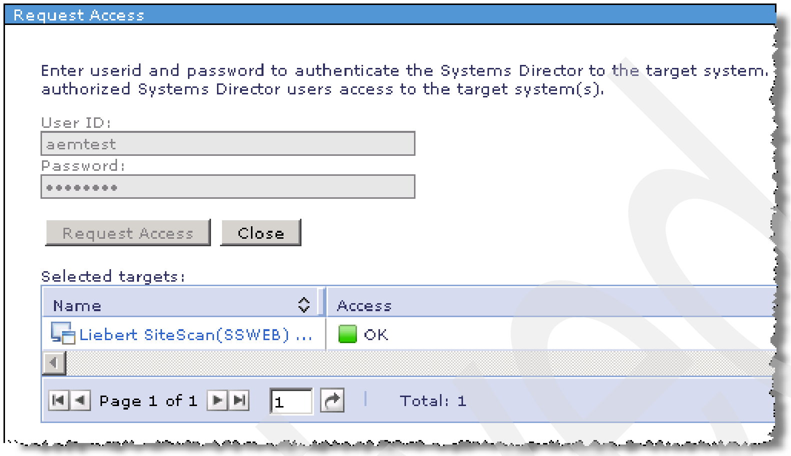


*Figure 10-3 Request access to Liebert SiteScan Web server*

6. After access moves to `Submitted` or `OK`, Liebert SiteScan Web devices are being discovered and should begin displaying in the Discovered Systems pane and also in the Active Energy Manager resources group, as shown in Figure 10-4.

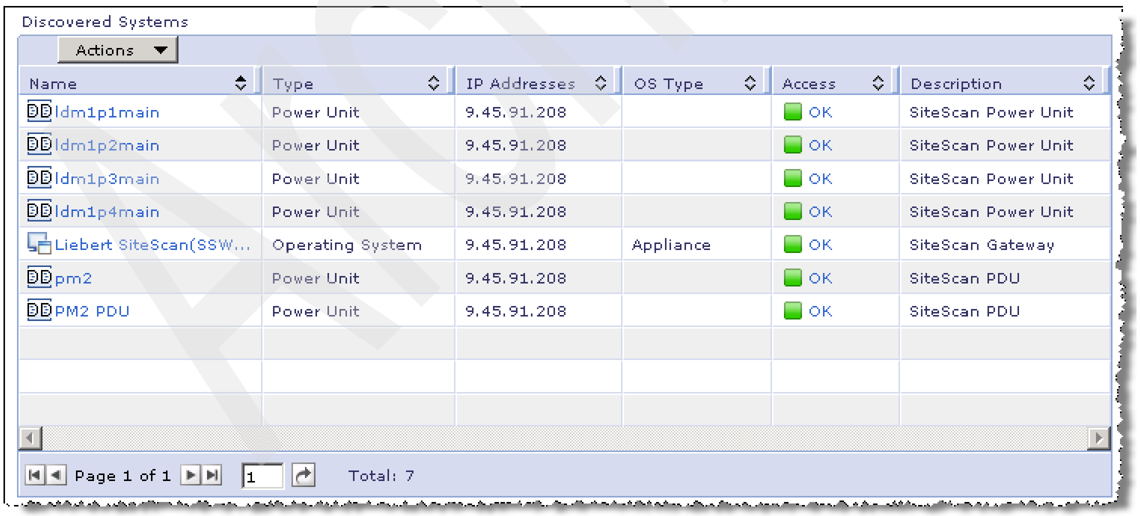


*Figure 10-4 Liebert SiteScan managed facility equipment listed under Discovered Systems*

7. In the IBM Systems Director navigation area, click **Navigate Resources**, and then click the **Active Energy Manager Resources group** to view the discovered SiteScan devices that Active Energy Manager can monitor as shown in Figure 10-5.

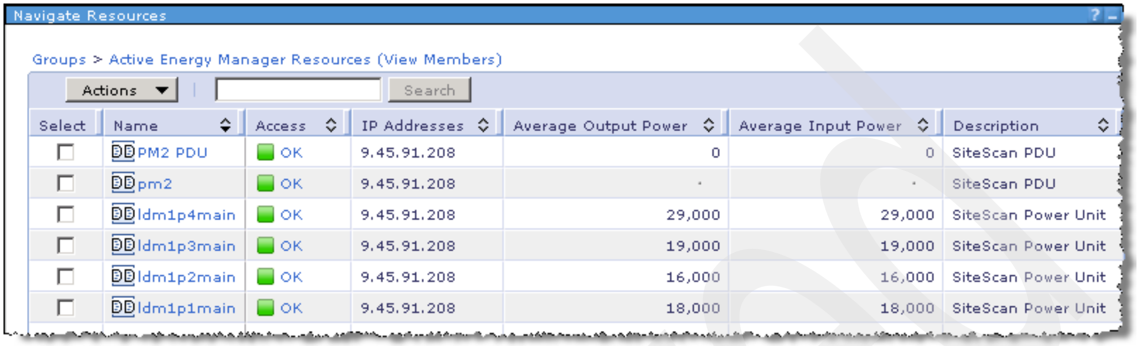
Navigate Resources

Groups > Active Energy Manager Resources (View Members)

Actions | Search

| Select | Name | Access | IP Addresses | Average Output Power | Average Input Power | Description |
|---|---|---|---|---|---|---|
| ☐ | PM2 PDU | OK | 9.45.91.208 | 0 | 0 | SiteScan PDU |
| ☐ | pm2 | OK | 9.45.91.208 | - | - | SiteScan PDU |
| ☐ | ldm1p4main | OK | 9.45.91.208 | 29,000 | 29,000 | SiteScan Power Unit |
| ☐ | ldm1p3main | OK | 9.45.91.208 | 19,000 | 19,000 | SiteScan Power Unit |
| ☐ | ldm1p2main | OK | 9.45.91.208 | 16,000 | 16,000 | SiteScan Power Unit |
| ☐ | ldm1p1main | OK | 9.45.91.208 | 18,000 | 18,000 | SiteScan Power Unit |

*Figure 10-5 Discovered Liebert SiteScan facility equipment*

> **Note**: An *appliance* is a computer systems or computer device where the operating system is not exposed to the user. For the appliance device discovery, you need to enter the user credentials of the service that is running on the target server.
>
> For example, the Liebert SiteScan Web server that you discovered in the previous scenario is an *Appliance* type object, where you requested access to the server on which Liebert SiteScan Web was installed using the user credentials of the SiteScan Web utility.

# 10.3 Eaton Power Xpert

This section provides an overview of the Power Xpert facility management solution from Eaton as well as information about how to use the features of this solution with Active Energy Manager.

## 10.3.1 Overview of Power Xpert integration

The integration of Active Energy Manager 4.1.1 and Power Xpert allows Active Energy Manager to integrate Eaton power equipment monitoring along with IT equipment monitoring, which is a first step towards integrating IT and facility equipment monitoring and management inside a data center. Active Energy Manager communicates with Power Xpert Reporting (PXR) and Power Xpert

Software (PXS) to discover and monitor Eaton facility equipment in the data center.

- Power Xpert Reporting

  Power Xpert Reporting is a reporting interface that enables administrators to define a hierarchical view of the power equipment infrastructure in the data center and to select and view the power attributes of devices for the purpose of monitoring.

- Power Xpert Software

  Power Xpert Software is a facility equipment management software suite. Power Xpert Software is responsible for real-time monitoring that includes reception of events and alarms from Eaton data center facility equipment, such as uniterruptible power supply units, PDUs, and so forth.

Active Energy Manager 4.1.1 can discover and monitor the uniterruptible power supply units and PDU devices that are listed in 10.3.2, “Supported equipment” on page 465 and that are monitored by Power Xpert Software management software.

Figure 10-6 shows the integration of Active Energy Manager 4.1.1 with Power Xpert.

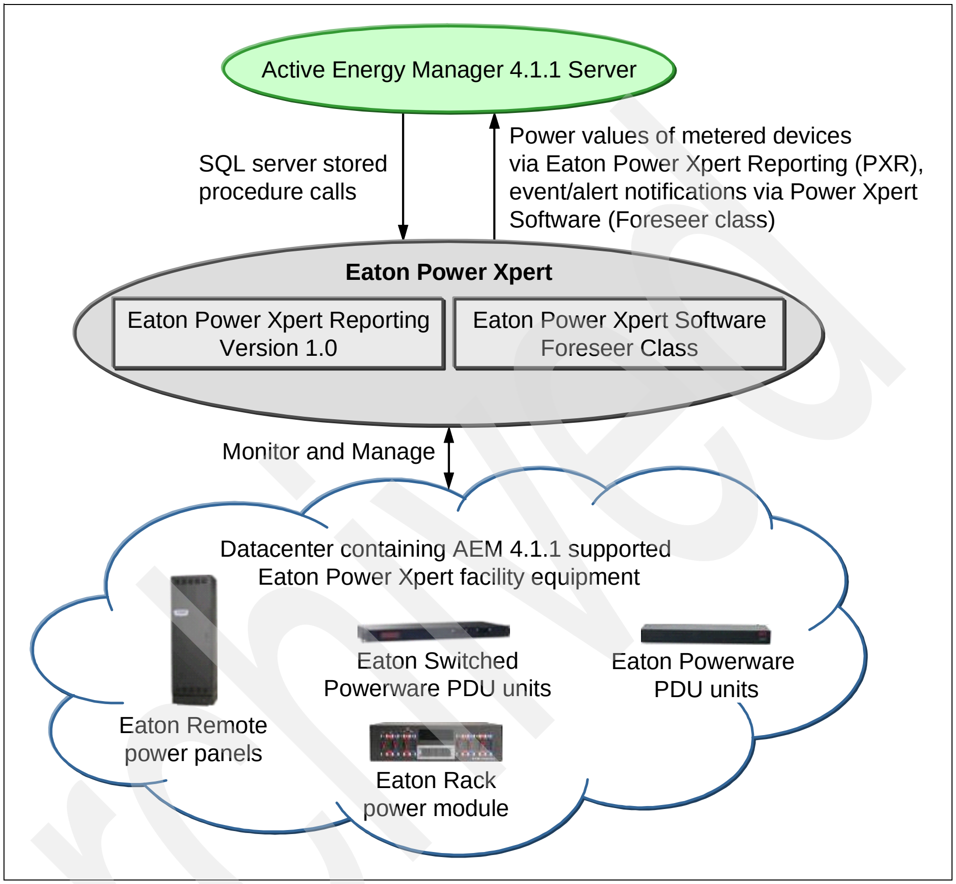


*Figure 10-6 Eaton Power Xpert integrated with Active Energy Manager*

Active Energy Manager communicates with the Power Xpert Reporting system using a set of SQL Server stored procedure calls to retrieve the power attribute values of Eaton power units in the data center. Active Energy Manager can also receive events and alert notifications from Eaton facility equipment using Eaton Power Xpert Software. For receiving events and alert notifications, you can configure the Systems Director IP address as one of the SNMP trap receivers in the Eaton Power Xpert Software system. Active Energy Manager supports communication with one instance of Eaton Power Xpert Reporting at a time. In turn, this instance of Power Xpert Reporting could have multiple Power Xpert Software configured.

Active Energy Manager polls the Power Xpert Reporting management server once every 15 minutes instead of every 5 minutes, which is the default polling interval for Active Energy Manager. If the Active Energy Manager default is set to anything under 15 minutes, devices based on Power Xpert Reporting are still polled every 15 minutes. If the Active Energy Manager polling interval is set to anything greater than 15 minutes, devices based on Power Xpert Reporting are polled at that greater interval.

Active Energy Manager can automatically associate Power Xpert Reporting Managed uninterruptible power supply devices to all the PDUs that are drawing power from those uninterruptible power supply devices.

Active Energy Manager v4.1.1 allows logical association of multiple resources such as servers, storage, and other IT equipment with the sensors supported in the facility manager equipment, such as PDU units, which allows the Active Energy Manager to externally meter the power and environmental monitoring values of the IT equipment that are present in the data center.

When any of the uninterruptible power supply devices fail in the data center, Active Energy Manager receives a critical event about the uninterruptible power supply failure. Using the automation plan feature of Systems Director, an alert e-mail can be sent to IT and facility administrators alerting them about the failure of the uninterruptible power supply devices and the IT equipment that can be affected from the failure of the device.

For more information about how to discover Power Xpert managed facility equipment, refer to 10.3.4, “Discovering Power Xpert in Active Energy Manager” on page 468.

## 10.3.2 Supported equipment

Active Energy Manager can monitor the following Power Xpert reporting managed facility equipment:

- Eaton Monitored Powerware PDUs
  - PW102SW0U150
  - PW102SW0U151
  - PW103SW0U152
  - PW103SW0U153
  - PW105SW0U154
  - PW306SW0U155
  - PW306SW0U156

- Eaton Switched Powerware PDUs
  - PW105MI0U096
  - PW105MI0U097
  - PW105MI0U098
  - PW105MI0U099
  - PW306MI0U113
  - PW309MI0U114
  - PW309MI0U115
  - PW110MI0U116
  - PW110MI0U117
  - PW110MI0U118
- Eaton Remote Power Panels
  - Models with a Power Xpert Gateway Series 1000 card
- Eaton models with a ConnectUPS-X, ConnectUPS-BD, or ConnectUPS-E Web or SNMP UPS Connectivity Device

### 10.3.3 Prerequisites before discovering Power Xpert

The following requirements need to be met before discovering Power Xpert managed facility equipment in Active Energy Manager:

- Power Xpert Reporting stores all the information that is related to the managed facility equipment in a Microsoft SQL Server database. To access the Power Xpert Reporting database to retrieve information about the monitored equipment, Active Energy Manager requires a Microsoft SQL Server JDBC driver. The SQL Server JDBC driver is downloadable free from the Microsoft Web site. It is not ship as part of the Active Energy Manager installation package because it is cannot be redistributed.

  To download and use the drive, follow these steps:

  a. Download the Microsoft SQL Server JDBC driver v2.0 from the following Microsoft Web site:

     http://msdn.microsoft.com/en-us/data/aa937724.aspx

  b. Install the downloaded Microsoft SQL Server JDBC driver per Microsoft instructions. Then copy file `sqljdbc.jar` into the following directory:

     ```
     <installation_path>/Director/ActiveEnergyManager/eclipse/plugins/
     com.ibm.aem.vendor_4.1.1/
     ```

     In this path, *<installation_path>* is the path to the installation directory for IBM Systems Director v6.1.1

  c. Restart IBM Systems Director so that Active Energy Manager can detect the JDBC driver.

- For Active Energy Manager to receive alerts from Power Xpert managed facility power equipment, you need to configure the IP address of the system on which Active Energy Manager is installed as an SNMP trap receiver on the system where the Foreseer version of Power Xpert Software is installed. This enables Active Energy Manager to associate these alerts with the power equipment inventory that is discovered in the Power Xpert Reporting database and to generate an appropriate Systems Director event for the device.

> **Note:** The Foreseer version of Power Xpert Software can be installed only on the system that is running Microsoft Windows. Power Xpert Foreseer uses the built-in SNMP Service of the Windows operating system to process the SNMP traps.

To set Systems Director as trap receiver on the Microsoft Windows system where the Eaton Foreseer version of Power Xpert Software is installed, follow these steps:

a. Click **Start** → **Control Panel** → **Administrative Tools** → **Services**.

b. Locate and right-click **SNMP Services**. Select **Properties** as shown in Figure 10-7.

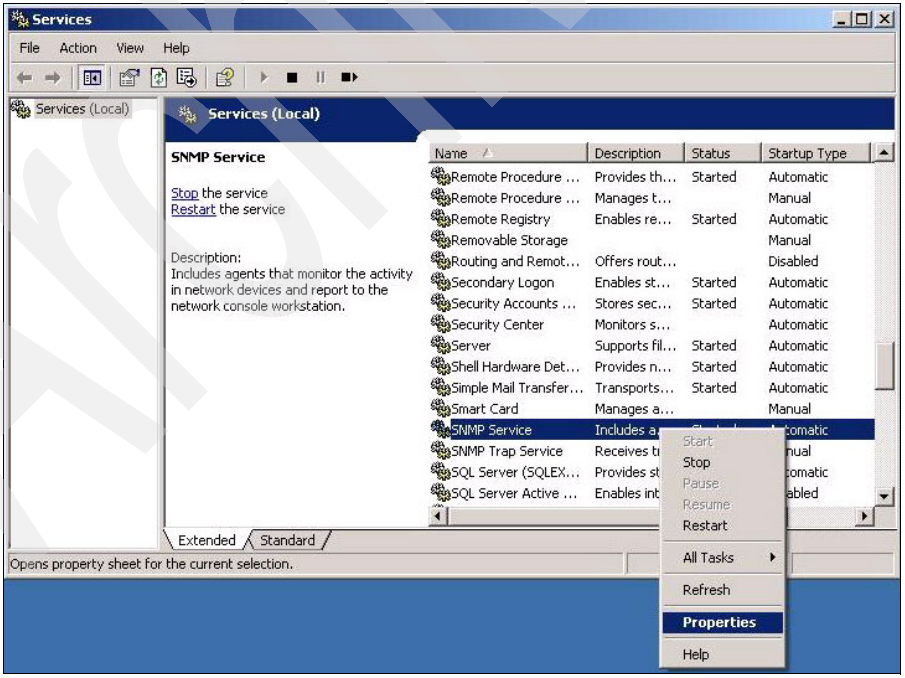


*Figure 10-7 Properties menu*

c. Enter the IP address of the IBM Systems Director Server to the “Trap destinations” list on the Traps tab of the SNMP Service properties page as shown in Figure 10-8. Click **OK**.

d. Restart the SNMP service.

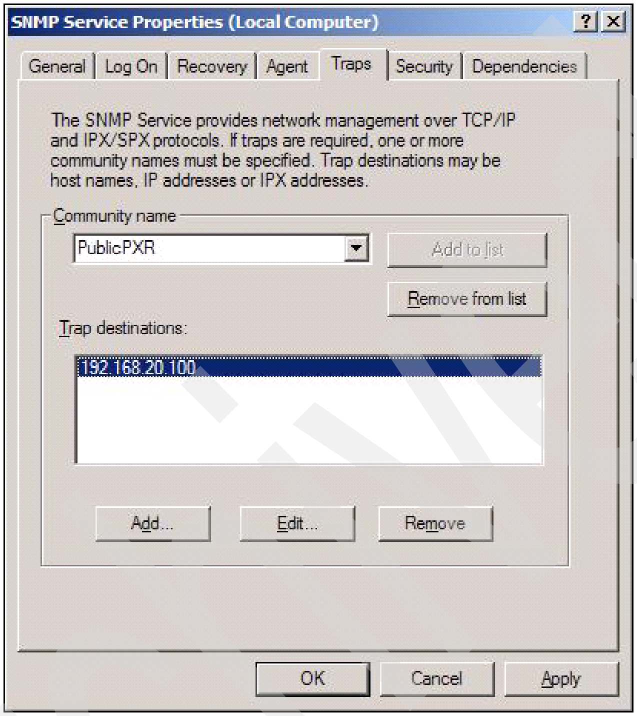


*Figure 10-8 Entering the IP address*

## 10.3.4 Discovering Power Xpert in Active Energy Manager

You can discover the Power Xpert Reporting managed facility equipment in Active Energy Manager by following these steps:

1. In the IBM Systems Director navigation area, click to expand **Inventory**, and then click **System Discovery**.
2. On the System Discovery page, specify the IP address of the operating system on which the Power Xpert Reporting management software is located.

3. In the Select resource type field, select **Operating System** from the list. Then select **Discover** as shown in Figure 10-2.

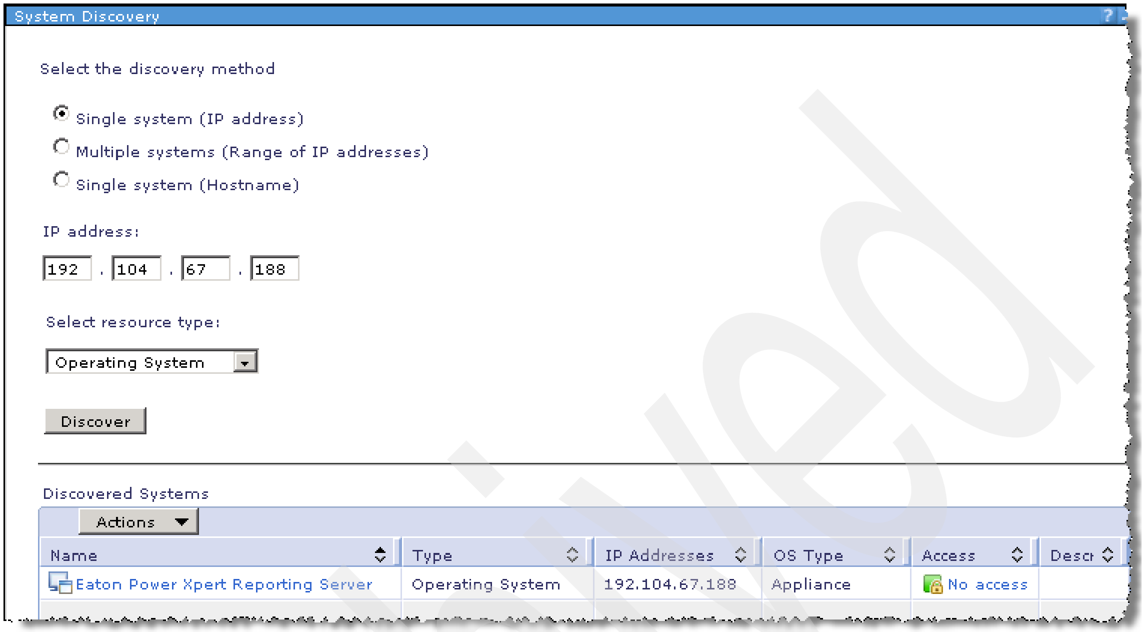


*Figure 10-9 Discovery of Power Xpert Reporting management server in IBM Systems Director*

4. When the operating system is discovered, right-click the operating system and select **Security** → **Request Access**.

5. On the Request Access page, enter the Eaton Power Xpert Reporting SQL Server user credentials, then select **Request Access**. Click **Close** after access is granted as shown in Figure 10-10.

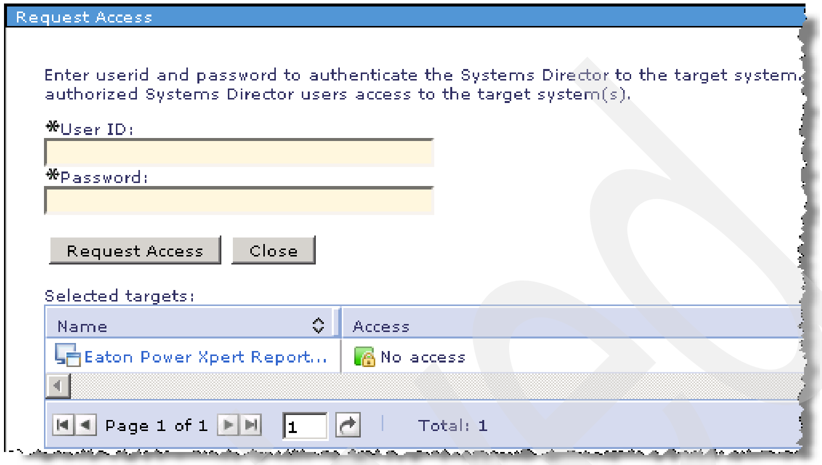


*Figure 10-10 Request access to Power Xpert Reporting server*

6. After access moves to `Submitted` or `OK`, Eaton Power Xpert Reporting managed facility equipment will be discovered and should display in the Discovered Systems pane (as shown in Figure 10-11 on page 471) and also in the Active Energy Manager resources group.

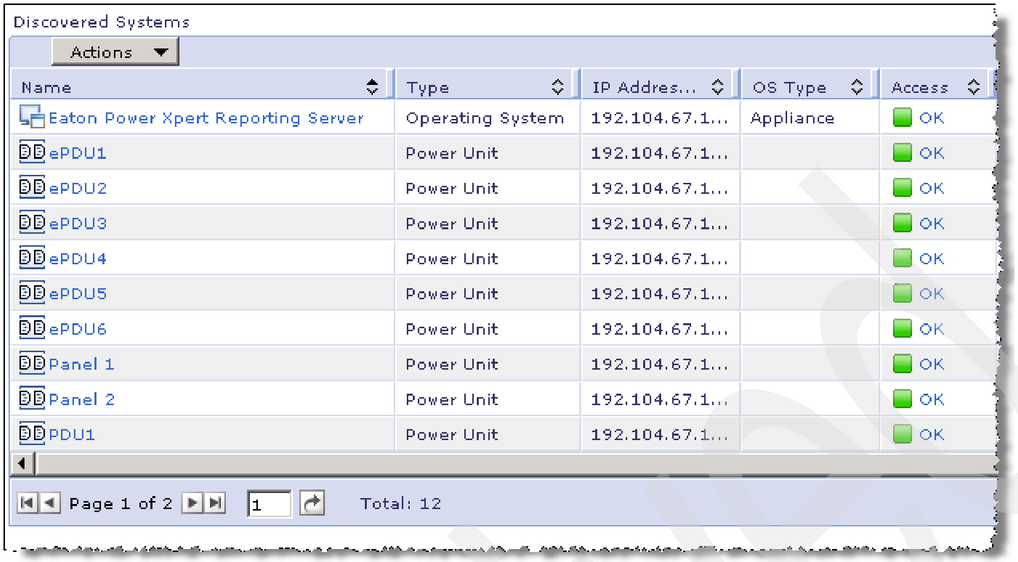


*Figure 10-11 Power Xpert Reporting managed facility equipment*

7. In the IBM Systems Director navigation area, click **Navigate Resources**, and then click the **Active Energy Manager Resources group** to view the discovered Eaton Power Xpert Reporting managed facility equipment that Active Energy Manager can monitor as shown in Figure 10-12.

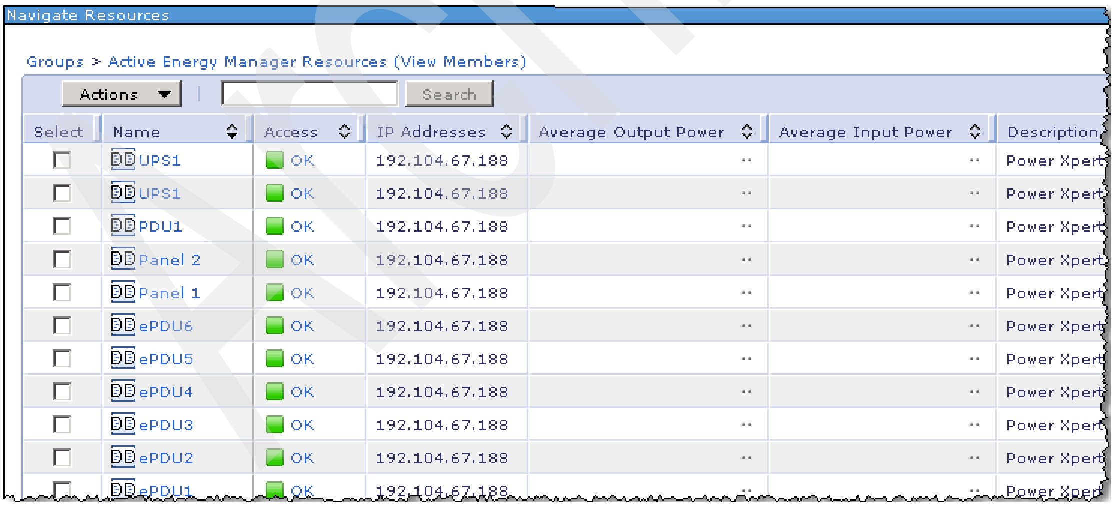


*Figure 10-12 Power Xpert managed facility equipment listed under Active Energy Manager resource group*

## 10.4 APC InfraStruXure Central

This section provides an overview of the InfraStruXure Central (ISXC) facility manager from APC as well as the Active Energy Manager 4.1.1 tasks that you can perform with this solution.

### 10.4.1 Overview of InfraStrucXure Central

InfraStruXure Central (ISXC) from APC is a popular facility management application that manages the APC power and cooling equipment in a data center. Active Energy Manager v4.1.1 can integrate with APC ISXC v5.1 to discover and integrate the monitoring of APC facility equipment along with IT equipment, giving a holistic view of data center health and status.

Figure 10-3 illustrates the integration of Active Energy Manager and the APC ISXC server.

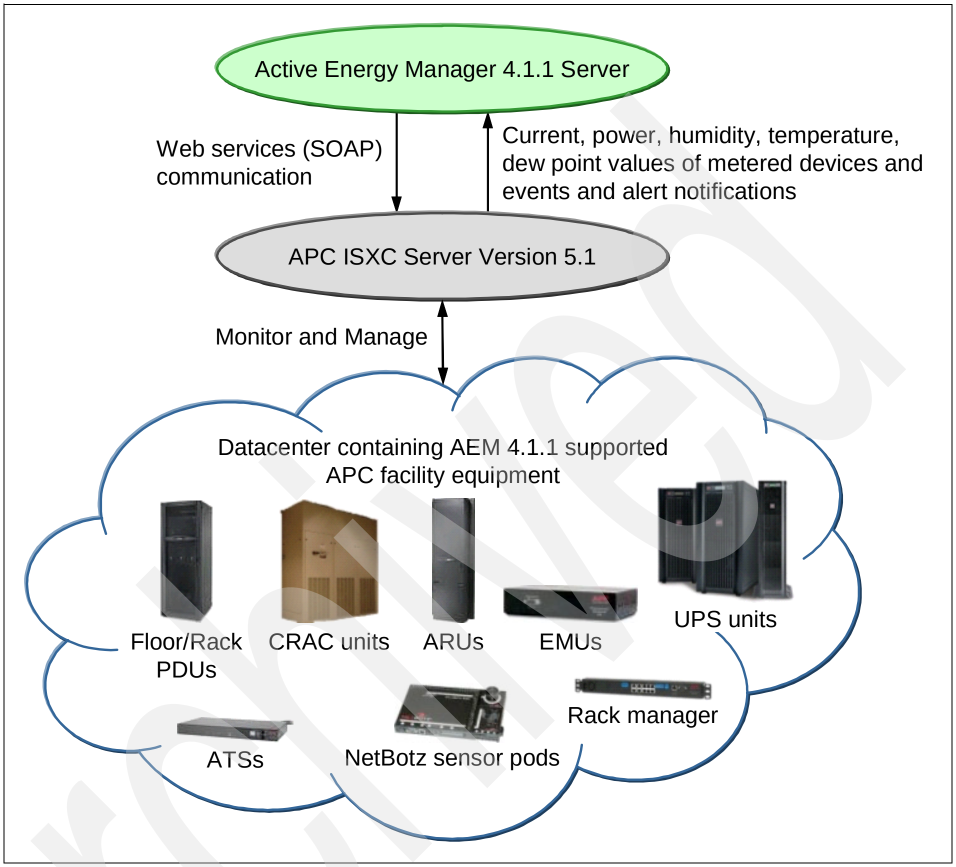


*Figure 10-13 APC ISXC facility manager integrated with Active Energy Manager*

Active Energy Manager can discover APC facility equipment that is managed and monitored by the APC ISXC v5.1 server. Active Energy Manager communicates to APC ISXC using Web services (SOAP) to discover and monitor all facility equipment. Active Energy Manager retrieves current (amps), power, humidity, temperature, and dew point values of the supported power, cooling, and environmental monitoring devices as well as the events and alert notifications from the APC ISXC v5.1 server.

For more information about APC ISXC v5.1 managed power, cooling, and environmental monitoring devices that can be discovered and monitored by Active Energy Manager, refer to 10.4.2, “Supported APC ISXC monitored facility equipment” on page 474.

Active Energy Manager polls the APC ISXC management server once every 5 minutes, which is the Active Energy Manager default polling interval. If the Active Energy Manager polling interval is set to anything greater than 5 minutes, the APC ISXC management server is polled at that greater interval.

For more information about discovery of APC ISXC managed facility equipment in Active Energy Manager, refer to 10.4.3, “Discovering APC ISXC managed facility equipment in Active Energy Manager” on page 474.

For more information about tasks that can be performed by Active Energy Manager on APC ISXC managed facility equipment, refer to 10.4, “APC InfraStruXure Central” on page 472.

## 10.4.2 Supported APC ISXC monitored facility equipment

Active Energy Manager supports the following facility equipment that are monitored by APC ISXC v5.1 server:

- APC Computer room air condition units (CRAC)
- Air Removal Unit (ARUs)
- PDUs and Rack PDUs
- Uninterruptible Power Supply Devices
- Rack Managers
- NetBotz sensor pods
- Environmental Monitoring Unit (EMUs)
- Automatic Transfer Switch (ATSs)

These metering devices can contain one or more current, power, humidity, temperature, or dew point sensors.

## 10.4.3 Discovering APC ISXC managed facility equipment in Active Energy Manager

You can discover APC ISXC managed facility equipment in Active Energy Manager by following these steps:

1. In the IBM Systems Director navigation area, click to expand **Inventory**, and then click **System Discovery**.
2. On the System Discovery page, specify the IP address of the ISXC server appliance on which the APC ISXC management software is located.

3. In the Select resource type field, select **Operating System** from the list. Then, select **Discover** as shown in Figure 10-2.

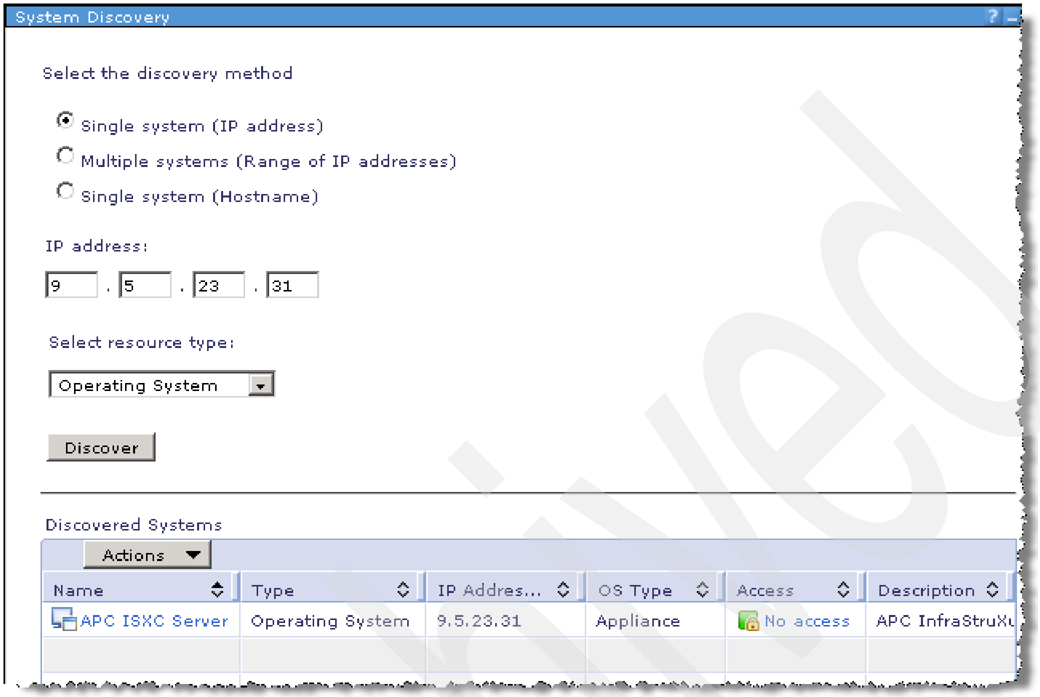


*Figure 10-14 Discovering the APC ISXC management server in IBM Systems Director*

4. When the operating system has been discovered, right-click the operating system and select **Security** → **Request Access**.

5. On the Request Access page, enter the APC ISXC web user credentials, then select **Request Access** as shown in. Click **Close** after the access grant as shown in Figure 10-3.

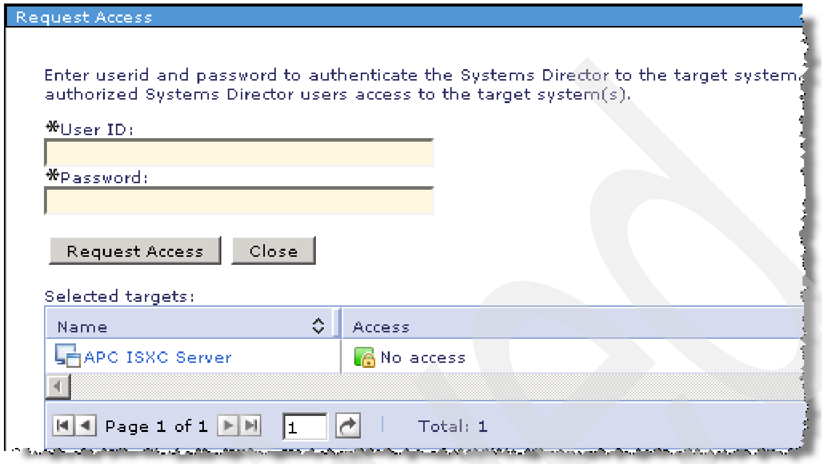


*Figure 10-15 Request access to APC ISXC server*

6. After access moves to `Submitted` or `OK`, the APC ISXC managed facility equipment is discovered and should display in the “Discovered Systems” pane as well as in the Active Energy Manager resources group. See Figure 10-16.

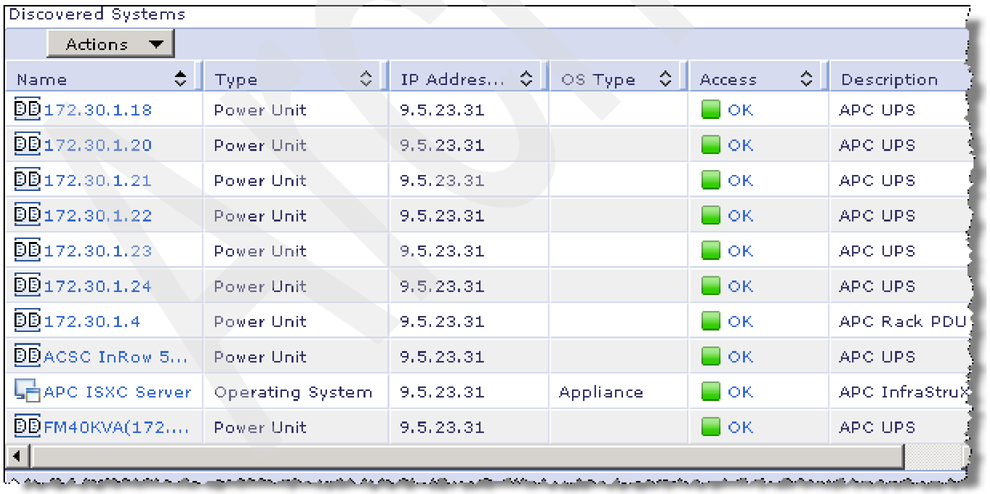


| Name | Type | IP Addres... | OS Type | Access | Description |
|---|---|---|---|---|---|
| 172.30.1.18 | Power Unit | 9.5.23.31 | | OK | APC UPS |
| 172.30.1.20 | Power Unit | 9.5.23.31 | | OK | APC UPS |
| 172.30.1.21 | Power Unit | 9.5.23.31 | | OK | APC UPS |
| 172.30.1.22 | Power Unit | 9.5.23.31 | | OK | APC UPS |
| 172.30.1.23 | Power Unit | 9.5.23.31 | | OK | APC UPS |
| 172.30.1.24 | Power Unit | 9.5.23.31 | | OK | APC UPS |
| 172.30.1.4 | Power Unit | 9.5.23.31 | | OK | APC Rack PDU |
| ACSC InRow 5... | Power Unit | 9.5.23.31 | | OK | APC UPS |
| APC ISXC Server | Operating System | 9.5.23.31 | Appliance | OK | APC InfraStruX |
| FM40KVA(172.... | Power Unit | 9.5.23.31 | | OK | APC UPS |

*Figure 10-16 APC ISXC managed facility equipment in Discovered Systems*

## 10.5 Active Energy Manager monitoring tasks

This section provides information about the following IBM Systems Director and Active Energy Manager features, which you can use to monitor the supported vendors facility equipment:

- Power and environmental values
- Event viewer
- Topology map viewer
- Configuring metering devices

### 10.5.1 Power and thermal data trending

You can use the Active Energy Manager trending feature to show the trends of the power and environmental sensor values that are reported by various vendor supported facility equipment. Active Energy Manager stores the trend data from 7 days to 3650 days.

This section provides the procedure to use the trending feature of Active Energy Manager by taking an example of the Eaton Power Xpert Reporting Server managed facility equipment:

1. Click **Navigate Resources** as shown in Figure 10-17.

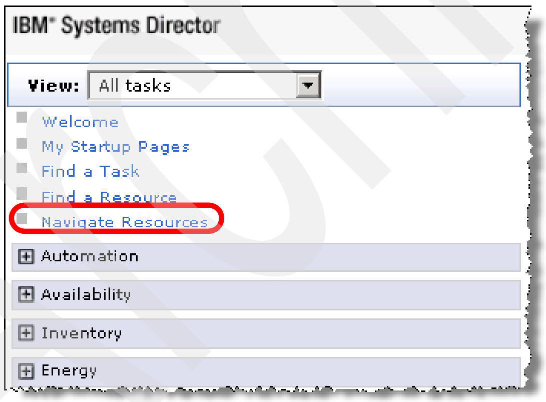


*Figure 10-17 IBM Systems Director Tasks pane*

2. Select **Active Energy Manager Resources** in the Navigate resource section as shown in Figure 10-18.

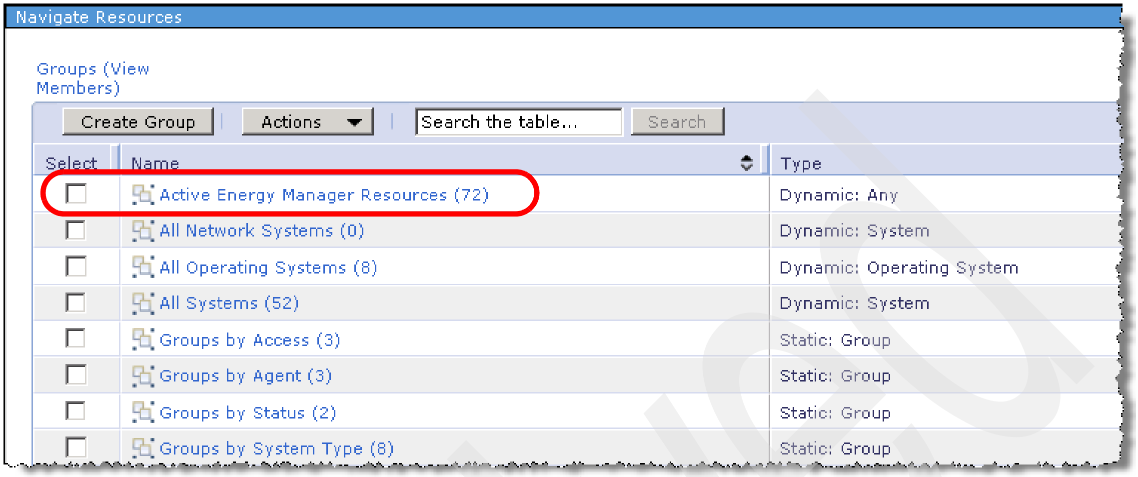


*Figure 10-18 IBM Systems Director Navigate resources*

3. Right-click any of the monitored facility equipment. Select **Energy** → **Trend data** as shown in Figure 10-19.

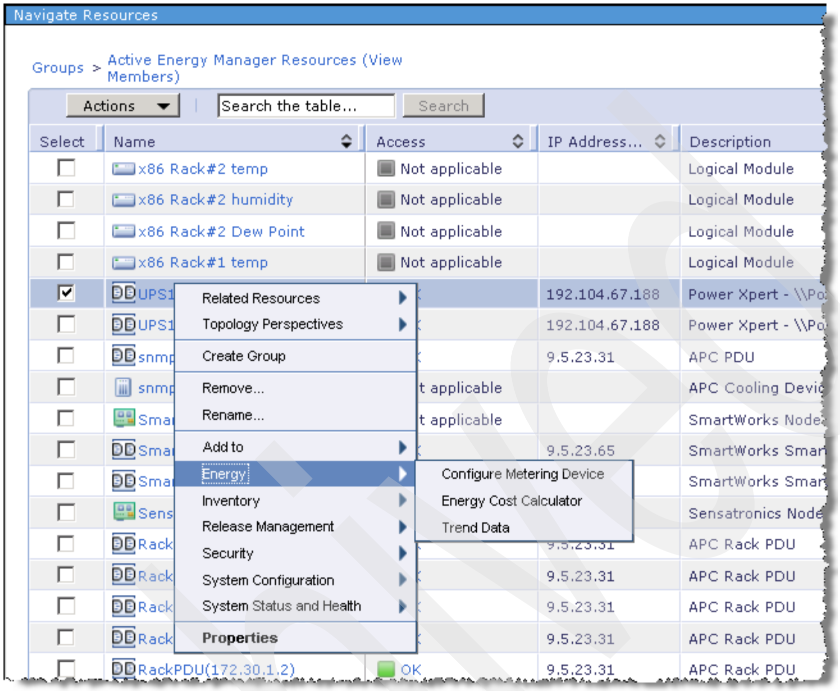


*Figure 10-19 Active Energy Manager Resources (view Members)*

Trend data chart displays as shown in Figure 10-20. You can view trend data in the chart view or the table view.

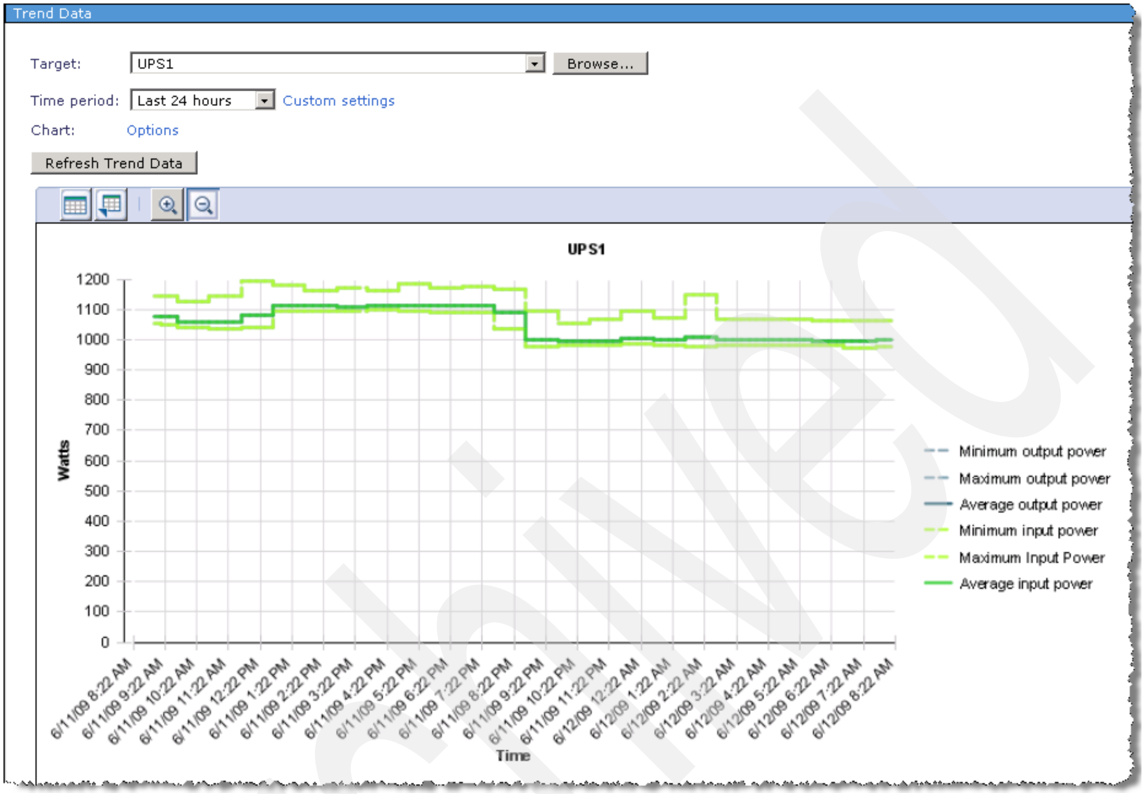


Figure 10-20 Eaton PXR managed uninterruptible power supply facility equipment input power trend

Active Energy Manager displays the events and alert notifications for the Power Xpert Reporting and the APC ISXC managed facility equipment as shown in Figure 10-21.

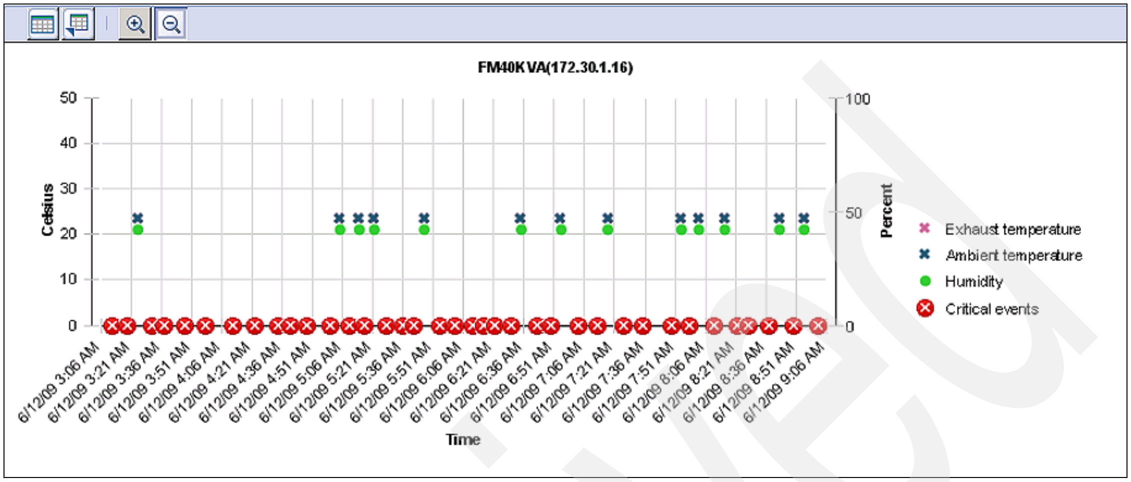


*Figure 10-21 APC ISXC server managed cooling facility equipment trend data*

## 10.5.2 Event viewer

Active Energy Manager can receive all events and alert notification from the supported vendors facility equipment, which allows IT administrators to be notified and to take necessary actions. All alert notifications from the facility equipment are also shown in the trend graphs.

All events and alert notifications are reported in **System Status and Health → Event Log** as shown in Figure 10-22.

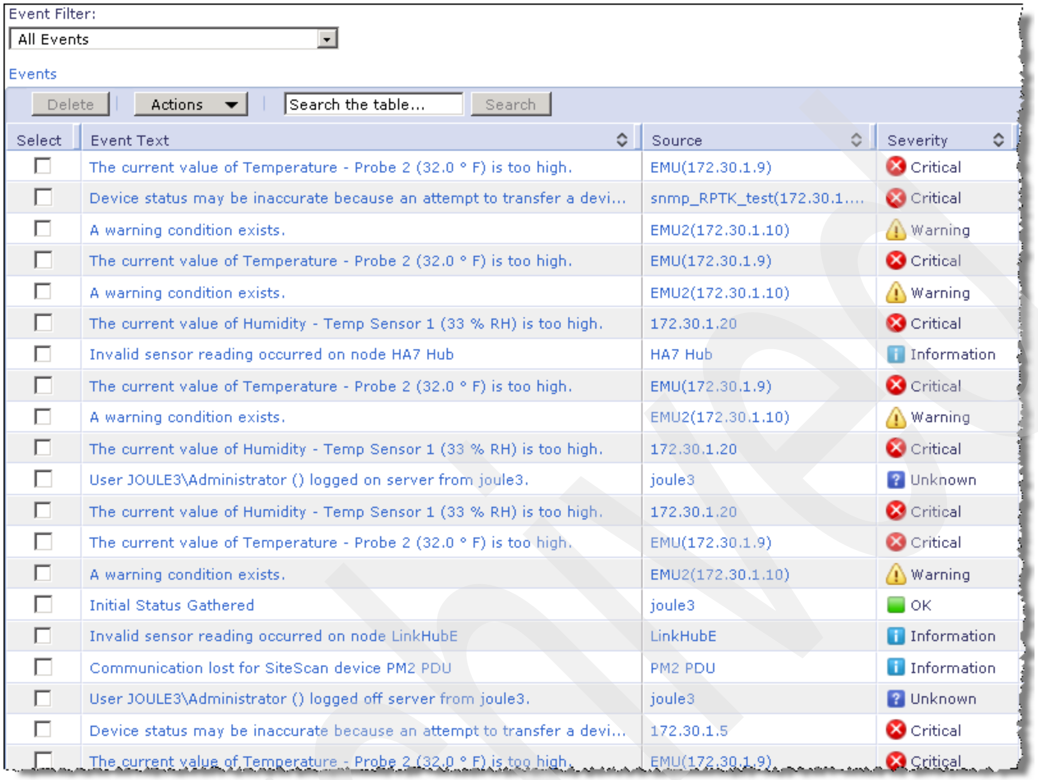

Event Filter:
All Events

Events

Delete | Actions | Search the table... | Search

| Select | Event Text | Source | Severity |
|---|---|---|---|
| ☐ | The current value of Temperature - Probe 2 (32.0 ° F) is too high. | EMU(172.30.1.9) | Critical |
| ☐ | Device status may be inaccurate because an attempt to transfer a devi... | snmp_RPTK_test(172.30.1.... | Critical |
| ☐ | A warning condition exists. | EMU2(172.30.1.10) | Warning |
| ☐ | The current value of Temperature - Probe 2 (32.0 ° F) is too high. | EMU(172.30.1.9) | Critical |
| ☐ | A warning condition exists. | EMU2(172.30.1.10) | Warning |
| ☐ | The current value of Humidity - Temp Sensor 1 (33 % RH) is too high. | 172.30.1.20 | Critical |
| ☐ | Invalid sensor reading occurred on node HA7 Hub | HA7 Hub | Information |
| ☐ | The current value of Temperature - Probe 2 (32.0 ° F) is too high. | EMU(172.30.1.9) | Critical |
| ☐ | A warning condition exists. | EMU2(172.30.1.10) | Warning |
| ☐ | The current value of Humidity - Temp Sensor 1 (33 % RH) is too high. | 172.30.1.20 | Critical |
| ☐ | User JOULE3\Administrator () logged on server from joule3. | joule3 | Unknown |
| ☐ | The current value of Humidity - Temp Sensor 1 (33 % RH) is too high. | 172.30.1.20 | Critical |
| ☐ | The current value of Temperature - Probe 2 (32.0 ° F) is too high. | EMU(172.30.1.9) | Critical |
| ☐ | A warning condition exists. | EMU2(172.30.1.10) | Warning |
| ☐ | Initial Status Gathered | joule3 | OK |
| ☐ | Invalid sensor reading occurred on node LinkHubE | LinkHubE | Information |
| ☐ | Communication lost for SiteScan device PM2 PDU | PM2 PDU | Information |
| ☐ | User JOULE3\Administrator () logged off server from joule3. | joule3 | Unknown |
| ☐ | Device status may be inaccurate because an attempt to transfer a devi... | 172.30.1.5 | Critical |
| ☐ | The current value of Temperature - Probe 2 (32.0 ° F) is too high. | EMU(172.30.1.9) | Critical |

*Figure 10-22 IBM Systems DIrector event log with notifications from various vendor facility equipment*

Active Energy Manager associates all events and alert notifications that it receives from Eaton Power Xpert Reporting and APC ISXC management servers to the appropriate managed facility equipment automatically.

To view the events and alert notification you should perform following steps

1. Navigate to Navigate resource → Active Energy Manager resources (view Members). Right-click any Eaton Power Xpert Reporting managed or APC ISXC managed facility equipment, as shown in Figure 10-23.

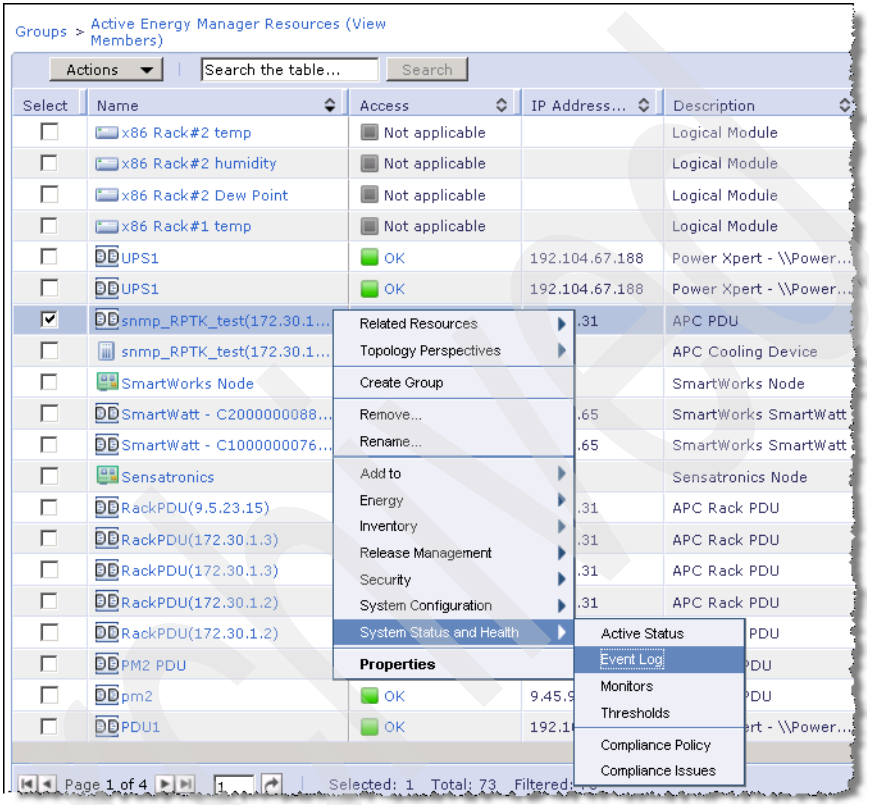


*Figure 10-23 Active Energy Manager Resources (view Members)*

2. All events and alert notifications pertaining of the selected APC PDU can be viewed as shown in Figure 10-24.

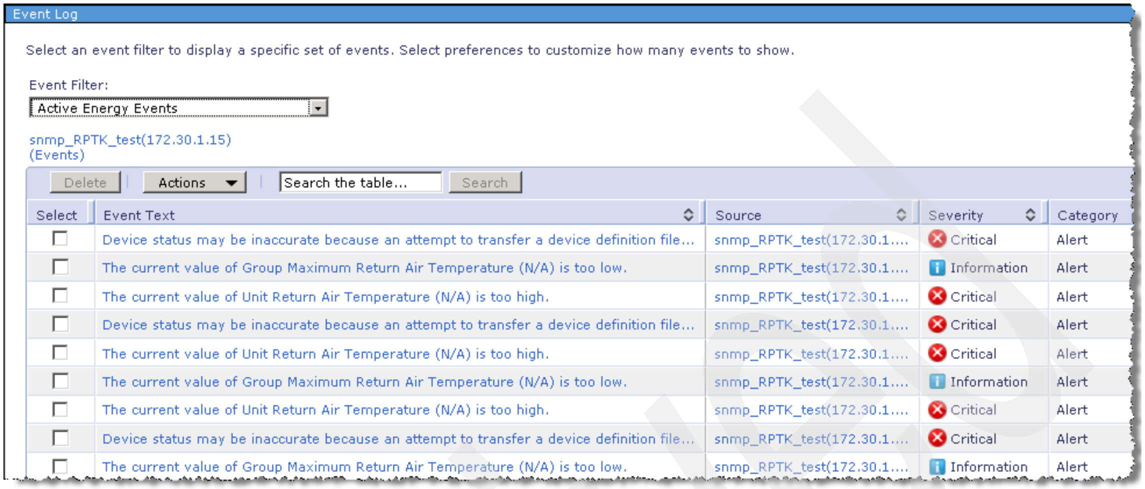


Figure 10-24 APC ISXC managed PDU events and alert log

## 10.5.3 Topology map viewer

The topology map viewer is one of the features of Active Energy Manager that can help the facility administrators to understand the types of sensors that are supported in the facility equipment as well as the connectivity between the facility equipment. For example, the facility manager can determine which power distribution units are getting power from which uninterruptible power supply units in a data center.

The following example explains how to navigate to the topology map viewer of a facility manager:

1. Click **Navigate Resources** in the IBM Systems Director **All Tasks** frame. Select **Active Energy Manager Resources** in Navigate Resources pane.
2. Right-click any facility management equipment that is managed by Eaton Power Xpert Reporting or APC ISXC server, and select **Topology Perspectives** → **Active Energy** as shown in Figure 10-25.

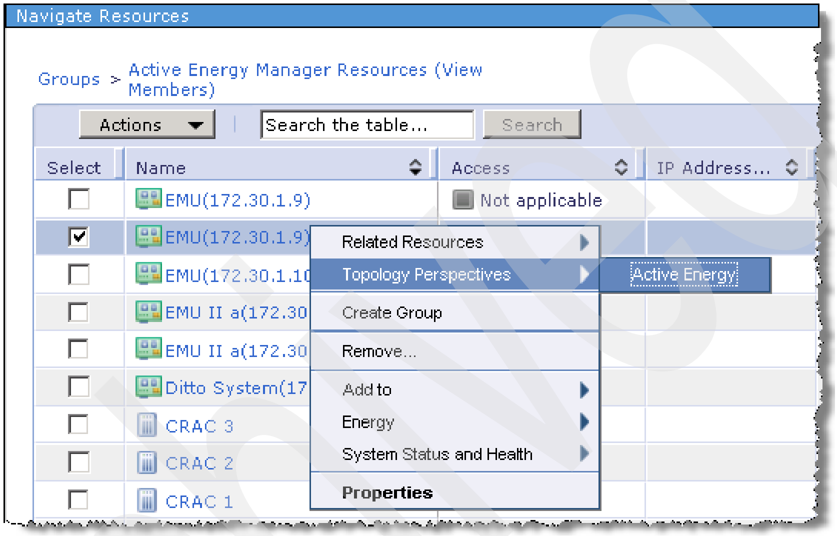


*Figure 10-25 Navigating to Topology view of APC ISXC server managed facility equipment*

The topology map shows the associated diagram of the selected facility equipment, any sensors if present, and any IT equipment that is logically associated with the sensors that are supported by the particular facility equipment. Figure 10-26 shows an APC ISXC managed Environmental Monitoring unit that consists of Temperature and Humidity sensors.

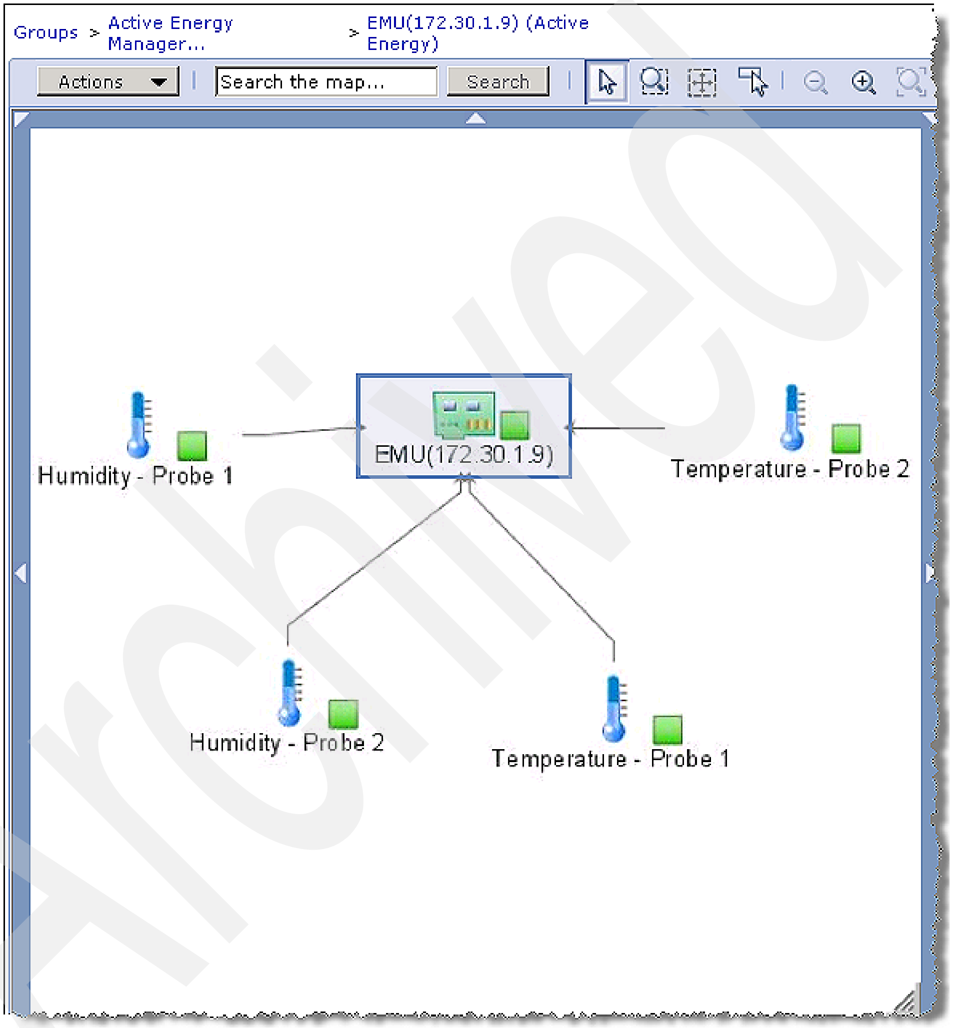


*Figure 10-26 Topology map of an APC ISXC server managed EMU*

## 10.5.4 Configuring metering devices

Configuring metering devices allows the IT administrator to associate the IT equipment to the sensors that are supported on the Power Xpert and APC ISXC server managed facility equipment. This feature allows logical association of IT equipment with power and environmental sensors that are present in supported facility equipment.

The following example shows how to configure metering devices for Power Xpert Reporting Server managed facility equipment:

1. Click **Navigate Resources** in IBM Systems Director **All Tasks** frame. Select **Active Energy Manager Resources** in Navigate Resources pane.
2. Right-click any facility management equipment managed by Eaton PXR or APC ISXC server, select **Energy** → **Configure Metering Device** as shown in Figure 10-27.

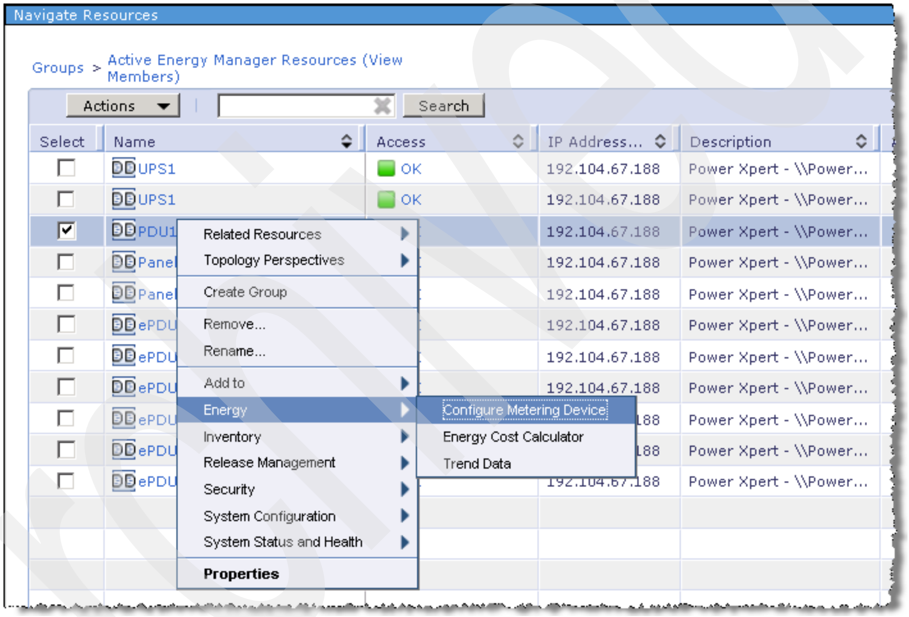


*Figure 10-27 Active Energy Manager Resources (view Members)*

3. If the selected facility equipment contains any type of sensors and power outlets, they display as shown in Figure 10-28. Right-click the sensor that you want to configure to associate it with any IT equipment that is present in the data center. Select **Edit**.

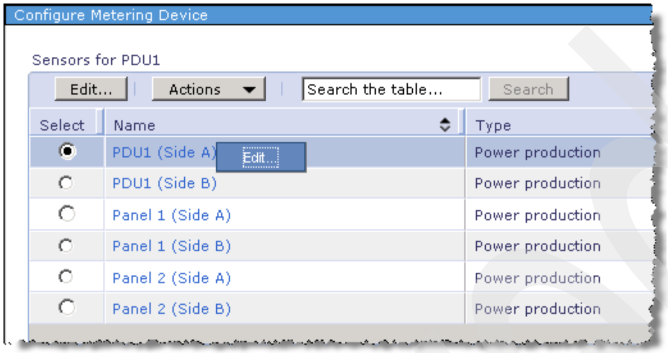


*Figure 10-28 Sensors present in the selected facility equipment*

4. Edit the properties of the sensor that you selected, as shown in Figure 10-29. Select the resources that you need to associate with this sensor:
   a. Select **None** if no resources are associated with this sensor.
   b. If IBM Systems Director does not support discovery of the resource, use the “Add Resource” link to define the logical device in IBM Systems Director.
   c. Associate the device to the sensor for external metering by clicking **Browse** and selecting the device from the appropriate group as shown in Figure 10-30.

> **Note:** If you have created logical device as explained in step b. The newly created logical device will be appearing in Active Energy Manager group.

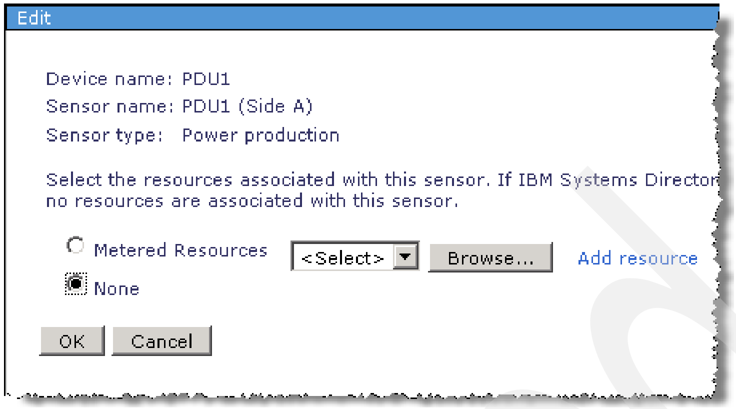


*Figure 10-29 Configuring metering device*

Active Energy Manager allows you to associate multiple devices with sensors present in Power Xpert Reporting managed and APC ISXC managed facility equipment, thus allowing facility managers to get a clearer idea of the topology and real time status of the data center.

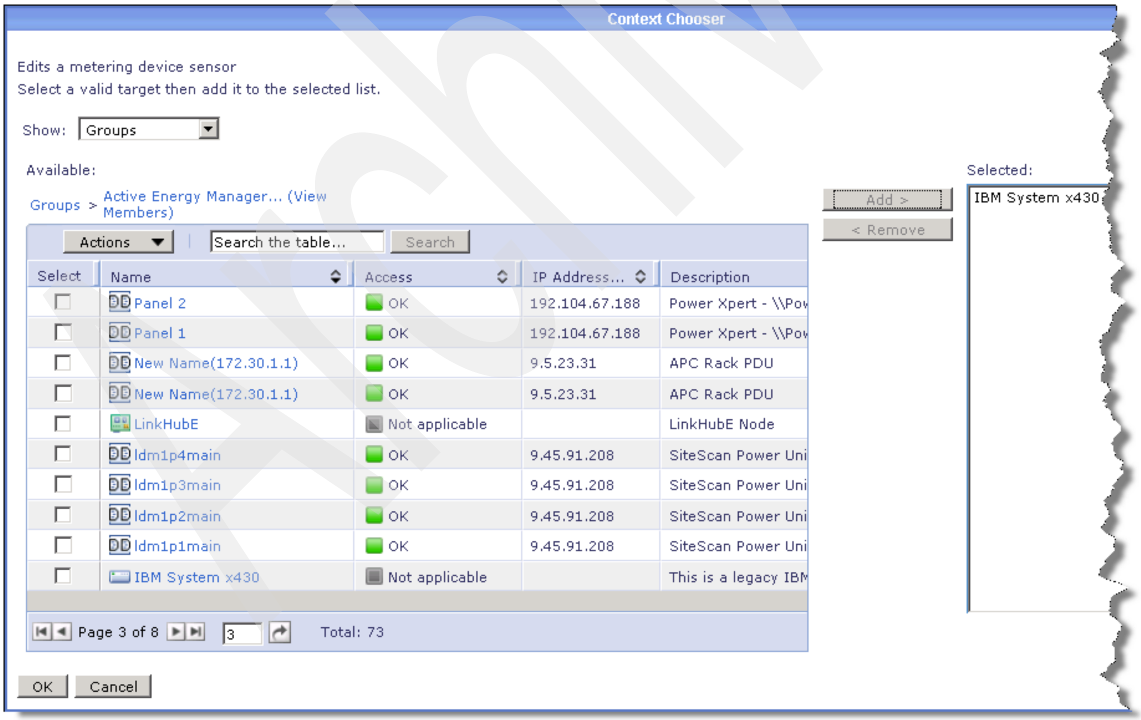


*Figure 10-30 Select devices for associating with the sensors present in the facility manager equipment*

5. After selecting resources need to be metered, select **Metered Resources** as shown in Figure 10-31. Click **OK**.

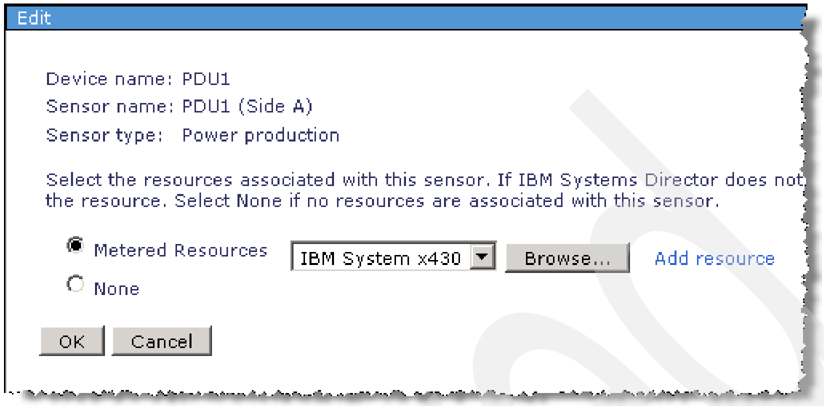


*Figure 10-31 Configuring metering device*

6. To verify the association, navigate to the topology map of the facility equipment. The associated metered device displays in the topology map as shown in Figure 10-32. The metered device has a logical connection to the sensor in the facility equipment. The logical device *IBM System x 430* is externally metered by the PDU1 side A power sensor.

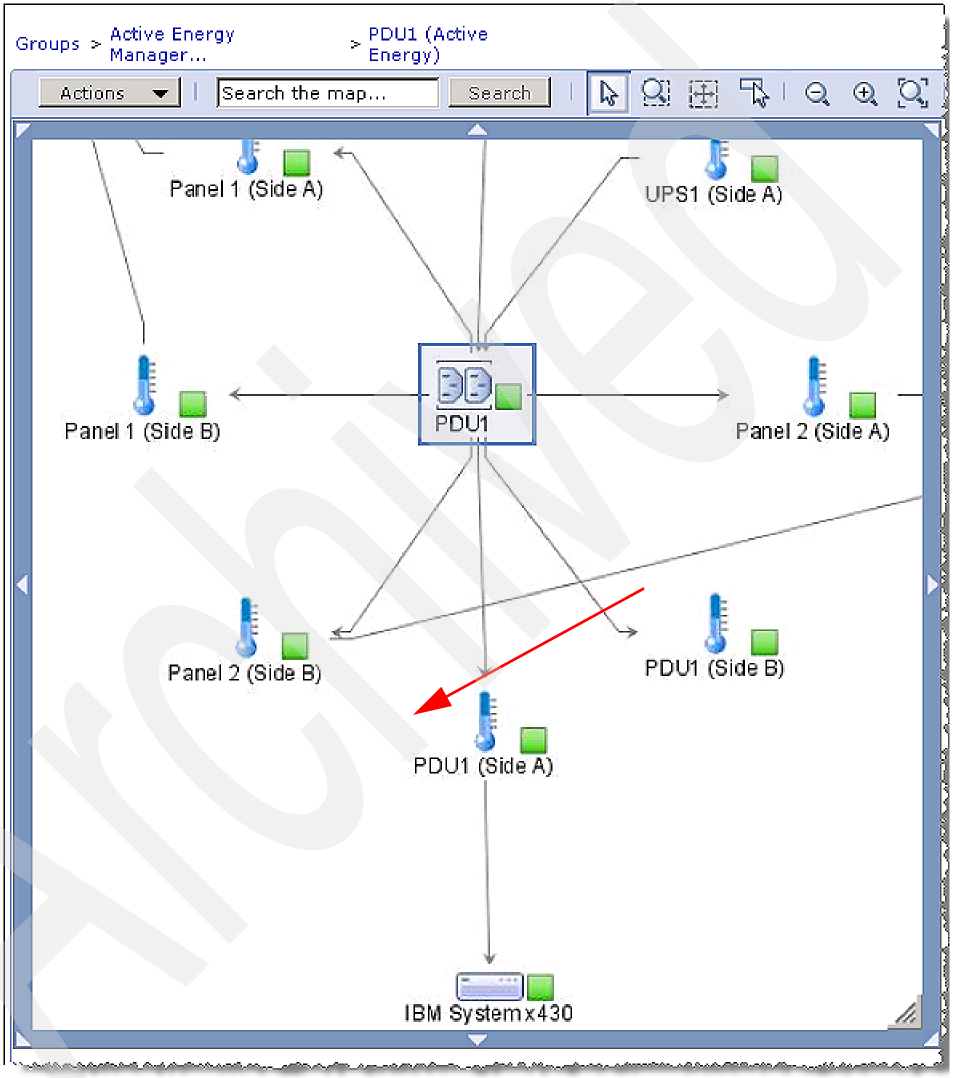


*Figure 10-32   Topology view associating the metered device and the PDU outlet*

# Abbreviations and acronyms

| | |
|---|---|
| **AC** | alternating current |
| **AEM** | Active Energy Manager |
| **AMD** | Advanced Micro Devices |
| **AMM** | Advanced Management Module |
| **APC** | American Power Conversion |
| **API** | application programming interface |
| **ASCII** | American Standard Code for Information Interchange |
| **ASF** | Alert Standard Format |
| **ASMI** | Advanced Systems Management Interface |
| **AWG** | American wire gauge |
| **BC** | BladeCenter |
| **BIOS** | basic input output system |
| **BMC** | Baseboard Management Controller |
| **CA** | Certification Authority |
| **CD-ROM** | compact disc read only memory |
| **CDE** | Common Desktop Environment |
| **CIM** | Common Information Model |
| **CIMOM** | Common Information Model Object Manager |
| **CLI** | command-line interface |
| **COM** | Component Object Model |
| **CPC** | Central Processor Complex |
| **CPU** | central processing unit |
| **CRAC** | Computer room air conditioner |
| **CSV** | comma separated variable |
| **DB** | database |
| **DBS** | demand-based switching |
| **DC** | domain controller |
| **DCIE** | Data Center Infrastructure Efficiency |
| **DCOM** | Distributed Component Object Model |
| **DFS** | Distributed File System |
| **DPI** | Distributed Power Interconnect |
| **DVFS** | Dynamic voltage/frequency scaling |
| **EMU** | Environmental Monitoring Unit |
| **FPC** | Flexible Power Conditioner |
| **FRU** | field replaceable unit |
| **FSP** | Flexible Service Processor |
| **GB** | gigabyte |
| **GUI** | graphical user interface |
| **HMC** | Hardware Management Console |
| **HP** | Hewlett Packard |
| **HTTP** | Hypertext Transfer Protocol |
| **I/O** | input/output |
| **IBM** | International Business Machines |
| **ID** | identifier |
| **IEC** | International Electrotechnical Commission |
| **IMM** | integrated management module |
| **IP** | Internet Protocol |
| **IPMI** | Intelligent Platform Management Interface |
| **ISXC** | InfraStruXure Central |

| | |
|---|---|
| **IT** | information technology |
| **ITSO** | International Technical Support Organization |
| **IVM** | Integrated Virtualization Manager |
| **JDBC** | Java™ Database Connectivity |
| **KB** | kilobyte |
| **KDE** | K Desktop Environment |
| **KW** | kilowatt |
| **LAN** | local area network |
| **LCD** | liquid crystal display |
| **LPI** | Linux Professional Institute |
| **MAC** | media access control |
| **MB** | megabyte |
| **MCSE** | Microsoft Certified Systems Engineer |
| **MDS** | Multi-layer Director Switch |
| **MIB** | management information base |
| **MM** | Management Module |
| **MMS** | Multi Module Series |
| **NEMA** | National Electrical Manufacturers Association |
| **NIM** | Network Installation Manager |
| **NSM** | Network and Systems Management |
| **NVRAM** | non-volatile random access memory |
| **OS** | operating system |
| **PCI** | Peripheral Component Interconnect |
| **PDU** | power distribution unit |
| **PET** | Platform Event Trap |
| **PMP** | Power Monitoring Panel |
| **PTF** | program temporary fix |
| **PURR** | processor utilization resource register |
| **PVC** | Polyvinyl chloride |
| **PXR** | Power Xpert Reporting |
| **PXS** | Power Xpert Software |
| **RSA** | Remote Supervisor Adapter |
| **SAN** | storage area network |
| **SCC** | System Control Cabinet |
| **SE** | Support Element |
| **SMS** | System Management Services |
| **SMT** | Simultaneous Multithreading |
| **SNMP** | Simple Network Management Protocol |
| **SOAP** | Simple Object Access Protocol |
| **SPURR** | scaled processor utilization resource register |
| **SQL** | Structured Query Language |
| **STG** | Server & Technology Group |
| **STS** | Static Transfer Switch |
| **TCP/IP** | Transmission Control Protocol/Internet Protocol |
| **UL** | Underwriters Laboratories |
| **UPS** | uninterruptible power supply |
| **URL** | Uniform Resource Locator |
| **USB** | universal serial bus |
| **VCF** | Voltage-Current-Frequency |
| **VCM** | Voltage-Current Monitoring |
| **VIOS** | Virtual I/O Server |
| **WMI** | Windows Management Instrumentation |
| **WW** | world wide |

# Related publications

We consider the publications that we list in this section particularly suitable for a more detailed discussion of the topics that we cover in this book.

## IBM Redbooks publications

For information about ordering these publications, see "How to get IBM Redbooks publications" on page 499. Note that some of the documents referenced here might be available in softcopy only.


- *Going Green with IBM Systems Director Active Energy Manager 3.1*, REDP-4361
- *Implementing IBM Director 5.20*, SG24-6188
- *Implementing IBM Systems Director 6.1*, SG24-7694


## Other publications

The following publications are also relevant as further information sources:


- IBM Information Center - Active Energy Manager 4.1.1

  http://publib.boulder.ibm.com/infocenter/director/v6r1x/index.jsp?topic=/aem_410/frb0_main.html
- Whitepaper: *IBM EnergyScale for POWER6 Processor-Based Systems*

  http://ibm.com/systems/power/hardware/whitepapers/energyscale.html
  ftp://ftp.software.ibm.com/common/ssi/sa/wh/n/pow03002usen/POW03002USEN.PDF
- DPI C13 PDU+, DPI C13 3-phase PDU+,DPI C19 PDU+, and DPI C19 3-phase PDU+ Installation and Maintenance Guide

  ftp://ftp.software.ibm.com/systems/support/system_x_pdf/43v6030.pdf
- Advanced Configuration and Power Interface Specification revision 3.0b

  http://www.acpi.info/DOWNLOADS/ACPIspec30b.pdf




 

## IBM Active Energy Manager Information Center

- Home page

  http://publib.boulder.ibm.com/infocenter/director/v6r1x/index.jsp?topic=/aem_410/frb0_main.html

- Operating systems supported by IBM Systems Director 6.1.0

  http://publib.boulder.ibm.com/infocenter/director/v6r1x/index.jsp?topic=/director.plan_6.1/fqm0_r_os_supported_by_ibm_director_61.html

- Troubleshooting and support

  http://publib.boulder.ibm.com/infocenter/director/v6r1x/index.jsp?topic=/aem_410/frb0_r_tbs_solving_problems.html

- Supported Managed systems

  http://publib.boulder.ibm.com/infocenter/director/v6r1x/index.jsp?topic=/aem_410/frb0_r_HW_reqs_managed_systems.html

- Migrating from 3.1.x or later

  http://publib.boulder.ibm.com/infocenter/director/v6r1x/index.jsp?topic=/aem_410/frb0_t_migrating.html

- Updating IBM Systems Director

  http://publib.boulder.ibm.com/infocenter/director/v6r1x/index.jsp?topic=/director.updates_6.1/fqm0_t_um_updating_director_features.html

- Upgrading and migrating IBM Systems Director

  http://publib.boulder.ibm.com/infocenter/director/v6r1x/index.jsp?topic=/director.upgrade_6.1/fqm0_t_upgrading_and_migrating.html

- Topology map view

  http://publib.boulder.ibm.com/infocenter/eserver/v1r2/index.jsp?topic=/eica7/eica7_concept_topology_map_view.htm

- Event actions

  http://publib.boulder.ibm.com/infocenter/director/v6r1x/index.jsp?topic=/director.automate_6.1/fqm0_c_ea_actions.html

- Working with active energy properties

  http://publib.boulder.ibm.com/infocenter/director/v6r1x/topic/aem_410/frb0_t_viewing_aem_properties.html

- Metering devices

  http://publib.boulder.ibm.com/infocenter/director/v6r1x/topic/aem_410/frb0_r_HW_reqs_metering_devices.html

- Active Energy Manager commands

  http://publib.boulder.ibm.com/infocenter/director/v6r1x/topic/aem_410/frb0_r_ref_commands.html

- Alphabetical listing of all **smcli** commands

  http://publib.boulder.ibm.com/infocenter/director/v6r1x/index.jsp?topic=/director.cli_6.1/fqm0_r_cli_a_to_z.html

- Global preferences commands

  http://publib.boulder.ibm.com/infocenter/director/v6r1x/index.jsp?topic=/aem_410/frb0_r_ref_commands_preferences.html

# Online resources

The following Web sites are also relevant as further information sources:

- Active Energy Manager home page

  http://www.ibm.com/systems/management/director/plugins/actengmgr/

- Tivoli Monitoring for Energy Management

  http://www.ibm.com/software/tivoli/products/monitor-energy-management

- IBM Fix Central

  http://www.ibm.com/support/fixcentral/

- Announcement letter: IBM Systems Director Active Energy Manager for x86

  http://www.ibm.com/common/ssi/cgi-bin/ssialias?infotype=dd&subtype=ca&&htmlfid=897/ENUS208-332

- Announcement letter: IBM Systems Director Active Energy Manager for Power Systems

  http://www.ibm.com/common/ssi/cgi-bin/ssialias?infotype=dd&subtype=ca&&htmlfid=897/ENUS208-284

- Announcement letter: IBM Systems Director Active Energy Manager for Linux on System z

  http://www.ibm.com/common/ssi/cgi-bin/ssialias?infotype=dd&subtype=ca&&htmlfid=897/ENUS208-344

- IBM DPI PDU enhancements

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5073205

- Dominion PX DPCR8A-20L6 Remote Power Management Solution

  `http://www.raritan.com/products/power-management/Dominion-PX/DPCR8A-20L6/`
- Sensatronics Model EM1 Environmental Monitor

  `http://www.sensatronics.com/index.php/industrial-monitors/model-em1.htm`
- Sensatronics Model E16 Temperature Monitor

  `http://www.sensatronics.com/index.php/industrial-monitors/model-e16.html`
- Sensatronics Model U16 Universal Temperature Monitor

  `http://www.sensatronics.com/index.php/industrial-monitors/model-u16.html`
- Sensatronics Senturion Environmental Monitor

  `http://www.sensatronics.com/index.php/it-monitors/senturion.html`
- Smart-Works: Smart-Net daisy-chain network

  `http://www.smart-works.com/smart-net/`
- Smart-Works: Smart-Sense

  `http://www.smart-works.com/smart-sense/`
- Smart-Works: Smart-Watt Energy Meters

  `http://www.smart-works.com/smart-watt/`
- iButtonLink: LinkHubE Ethernet Hub

  `http://www.ibuttonlink.com/linkhube.aspx`
- iButtonLink: T-Sense Temperature Sensor

  `http://www.ibuttonlink.com/t-sense.aspx`
- iButtonLink: T-Probe - Temperature Sensor with 10 foot cable

  `http://www.ibuttonlink.com/tprobe.aspx`
- iButtonLink: T-String - Temperature Sensor

  `http://www.ibuttonlink.com/tstring.aspx`
- iButtonLink: MS-T - Temperature Sensor

  `http://www.ibuttonlink.com/ms-t.aspx`
- iButtonLink: MS-TH - Temp/Humidity Sensor

  `http://www.ibuttonlink.com/ms-th.aspx`

- EDS HA7Net - Ethernet 1-Wire Host Adapter

  http://www.embeddeddatasystems.com/HA7Net--Ethernet-1-Wire-Host-Adapter_p_22.html
- SynapSense

  http://www.synapsense.com
- Microsoft SQL Server JDBC Driver

  http://msdn.microsoft.com/en-us/data/aa937724.aspx

## How to get IBM Redbooks publications

You can search for, view, or download Redbooks, Redpapers, Technotes, draft publications and Additional materials, as well as order hardcopy Redbooks publications, at this Web site:

**ibm.com**/redbooks

## Help from IBM

IBM Support and downloads

**ibm.com**/support

IBM Global Services

**ibm.com**/services

# Index

## D

## E



## F

IBM

Redbooks

# Implementing IBM Systems Director Active Energy Manager 4.1.1

(1.0" spine)
0.875"<->1.498"
460 <-> 788 pages

# Implementing IBM Systems Director Active Energy Manager 4.1.1

**Introduces energy management concepts and technologies**

**Discusses power and energy efficiency on IBM systems**

**Includes a step-by-step guide on implementing and using Active Energy Manager**

Energy efficiency is a critical priority for IT managers because energy and power costs can be a significant portion of IT costs. Thus, understanding and investing in energy management is critical. With IBM Systems Director Active Energy Manager, an extension of IBM Systems Director, you can monitor and manage the power usage of systems. Originally designed to support IBM BladeCenter and System x, Active Energy Manager now supports the power management of additional IBM systems, including POWER6 processor-based systems, as well as storage devices using an intelligent Power Distribution Unit (PDU+).

Active Energy Manager helps determine the proper power allocation for each system in the data center. It can assist in determining how to allocate power to existing systems more efficiently so that additional systems can be accommodated without the need for additional power and cooling. When power is constrained, chargeable optional features of Active Energy Manager allow power to be rationed on a system-by-system basis, enabling available processing power to match current workload closely.

This IBM Redbooks publication can help system administrators effectively monitor and manage the power usage of systems in a data center. This book introduces energy management concepts and technologies, and then provides a step-by-step guide to implementing and using Active Energy Manager.



SG24-7780-00 ISBN 0738433500