



John Brady
Ira Chavis
Matthew Finlayson
Michael Harkins
John P. Mullin
Julie Peet
Sheryl Qualters
Rodrigo Samper
Elisabeth Stahl

Implementing the Poughkeepsie Green Data Center: Showcasing a Dynamic Infrastructure

The IBM® Design Center Data Center in Poughkeepsie, New York, was running out of cooling capacity and needed to address continuing IT growth requirements. Innovative change was required and transformation of this infrastructure was achieved through IBM's smarter planet and dynamic infrastructure vision. Energy efficiency for data centers has become an imperative area of focus as the price of energy increases and systems grow beyond the capacity of current facilities to supply their power and cooling needs.

This IBM Redpaper publication discusses IT growth and energy efficiency for a case study about the IBM Design Center Data Center and its transformation to the *Poughkeepsie Green Data Center*. This paper also discusses the analysis and assessment behind this case study. It reviews the enhancements to the data center facilities that were performed. In addition, this paper highlights energy management techniques and demonstrates measurement and monitoring in the data center. Finally, this paper concludes with recommendations for a dynamic infrastructure data center of the future.

The dynamic infrastructure, IT growth, and energy efficiency

The IBM smarter planet vision brings a new level of intelligence to how the world works and how every person, business, organization, government, natural system, and man-made system interacts. Each interaction represents a chance to do something better, more efficiently, and more productively.

Smarter planet is IBM's vision of how interconnected technologies are changing the way the world works. Examples of these systems and processes include physical goods that are developed, manufactured, bought, and sold; services that are delivered; and people, money, oil, water, and electrons that are moved. The catalyst for this change is the transformation of the world's infrastructure, which is becoming increasingly instrumented, interconnected, and intelligent.

A dynamic infrastructure will bring more intelligence, automation, integration, and efficiencies to the digital and physical worlds. It will enable businesses and governments to better manage challenges presented by today's globally integrated planet.

IBM outlines three critical requirements for a twenty-first century infrastructure:

- ▶ The ability to integrate all of a company's digital and physical infrastructures
- ▶ The capacity to store, manage, and analyze massive amounts of data
- ▶ The wherewithal to reduce inefficiencies in business processes, technology, and costs

This dynamic infrastructure encompasses service management, asset management, virtualization and consolidation, information infrastructure, energy efficiency, security, and business resiliency.

For a data center, smarter can involve server consolidation, space savings, and energy efficiency. Smarter can be a lower total cost of ownership. Smarter can be green.

The traditional data center

The IBM Design Center Data Center in Poughkeepsie, New York, started with a ten-year-old data center floor that was running out of cooling capacity and has continually needed to address IT growth requirements. This center was originally built in 1999 to host a client center whose mission was to enable client IT infrastructures with IBM server and storage technologies. The facility was built in a former working lab space within one of the campus buildings. It consisted of several client meeting rooms and a data center space that was divided between a showcase area and a working IT data center.

The showcase area displayed the leading edge IBM server technologies in a glass enclosure, which provided ease of viewing. The working space contained a working IT infrastructure of servers, storage, and network gear, as illustrated in Figure 1 on page 3. The entire center was equipped with the latest IBM servers from the IBM System z® 900 enterprise server line, IBM Netfinity®, IBM RS/6000®, and IBM Enterprise Storage Server® (ESS).

The Design Center consisted of the following initial configuration:

- ▶ Four IBM RS/6000 B50 Rack Mounted Servers (1 rack)
- ▶ One IBM RS/6000 SP cluster (1 rack)
- ▶ Two IBM Enterprise Storage Server Model E20 (2 racks)
- ▶ Sixteen Netfinity 7000 rack mounted servers (4 racks)
- ▶ One IBM AS/400® (1 rack)
- ▶ One IBM S/390® environment including automated tape library

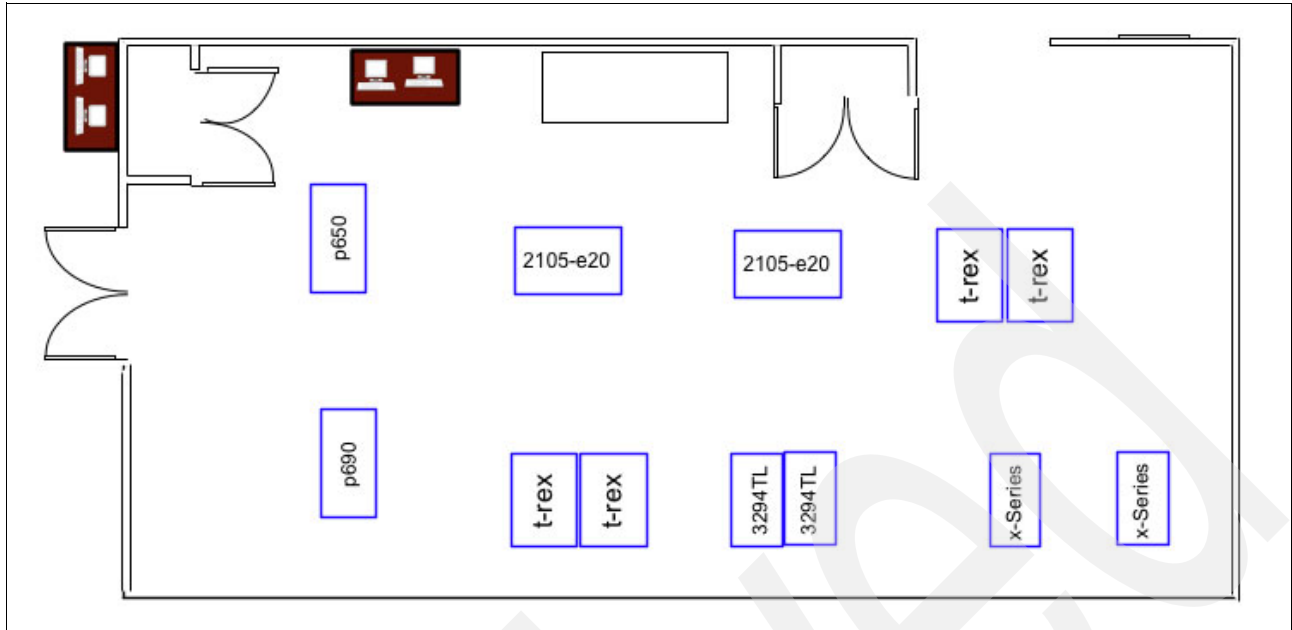


Figure 1 Original IBM Design Center

The power and cooling was planned to be sufficient to serve the IT requirements for the center's mission for years to come. Being built on a raised floor structure, the center was cooled by an Air Conditioning Unit (ACU) from Liebert, and air circulation was provided by randomly placed perforated tiles. Power for the center was delivered by a range of under-floor power configurations (standard 110V/15A to 3 Phase 220V/60A circuits).

Over the years, the server inventory was continually refreshed to the latest technologies, and additional capacity was added as client facing activities demanded more infrastructure. IBM products including IBM BladeCenter®, System p®, System i®, System x®, System z, and TotalStorage® servers were regularly added to the Design Center lab. Each of these servers was a foundation for new technologies such as multi-core processors, extremely dense server enclosures, and virtualization features.

The expansion was routinely handled by introducing new hardware. If needed, new power circuits were added. This trend continued for several years until the observation was made that the ambient temperature of the lab was increasing. Some tuning of the cooling infrastructure was performed, but this work proved to provide a short-term fix. By mid-2008, the server infrastructure in the Design Center had increased as follows:

- ▶ Nine BladeCenter chassis with mixture of 100+ HS/LS/JS/QS BladeCenter servers
- ▶ Fifty System x rack mounted servers
- ▶ Two IBM Power 570 (POWER6™) rack mounted servers
- ▶ One IBM System z9® and z10
- ▶ IBM TotalStorage including the following details
 - XIV®
 - DS4000®
 - DS8000®
 - Brocade SAN switches
 - N Series 3700 and 5200
 - Two TS3100 tape libraries
 - 3494 automated tape library

This configuration occupied a total of eight racks.

It was clear there was a problem: The power and cooling capacities of the facility had been maximized, and there was no opportunity to add power or cooling without significant costs. It was a common occurrence to experience high temperature alarms from the cooling units. Something had to be done to fix this problem.

Initial analysis and assessment

The first step in the Poughkeepsie Green Data Center transformation was to understand the thermals in the existing configuration. Modeling was then used to assist in planning the new data center layout.

Mobile Measurement Technology analysis

Mobile Measurement Technology (MMT) is a thermal and cooling system focused assessment for the entire data center that is based on detailed temperature, humidity, and airflow measurements taken by using IBM developed equipment and software. MMT captures high resolution temperature and humidity data by using a unique IBM sensor array that scans all relevant sections of the data center from floor to just above the tops of the racks. A cart with thermal sensors mounted in a defined three-dimensional (3-D) pattern is rolled through the data center while data is logged and mapped onto a coordinate grid. In addition, detailed airflow measurements are taken. Measurements are also made on the air handling equipment, and IT power conditions for the data center are surveyed.

Along with the temperature and humidity data, the airflow and power information that is collected is analyzed by using software developed by IBM. The output of this process is full 3-D temperature and humidity maps for the space, as well as calculated targeted airflow and evaluation of cooling system performance and energy consumption metrics.

For the data center in this study, some of the key measurements that were surveyed included the following details:

- ▶ 1,040 square feet of raised floor
- ▶ Approximately 9,100 temperature measurements, including 250 rack inlet temperatures
- ▶ Measured air flow from all perforated tiles (total of 19 flow measurements)
- ▶ One ACU, the only active ACU on the raised floor
- ▶ All relevant data center power levels

Figure 2 shows the layout of the Design Center Data Center before the changes were made. The boxes in gray indicate the servers (light gray) and storage (dark gray). The darker blue color highlights the inlet air locations for the respective servers. The Xs show the locations of the perforated tiles that allow cool air to enter the data center from the underfloor plenum. The ACU is shown in the light blue box with an arrow showing direction of airflow.

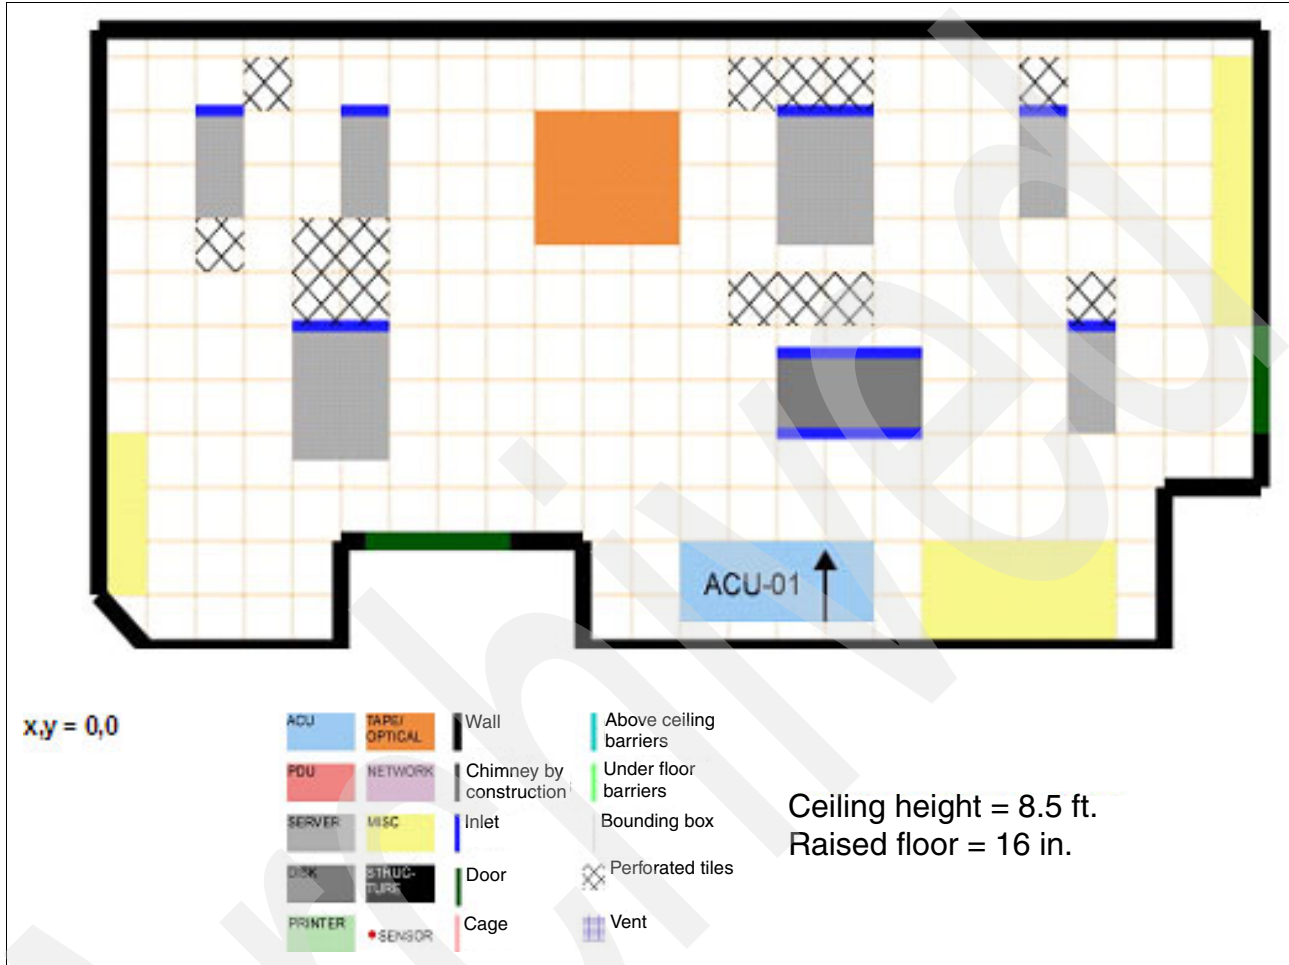


Figure 2 Layout of Design Center Data Center

Figure 3 illustrates the chiller and ACU power consumption and cooling efficiency. The initial cooling efficiency for the data center was found to be quite high. Upon further investigation, it was found that the only ACU in the area is being utilized at 77% of its rated capacity. This relatively high ACU utilization, on the only unit in the space, results in good cooling efficiency for the data center as a whole. Note that the heat density in the data center is only 57 Watts/square foot (W/sq. ft.). That value was considered typical in the past, but low in today's environment where you can find a rack full of blades dissipating 20 kiloWatts (kW) per rack.

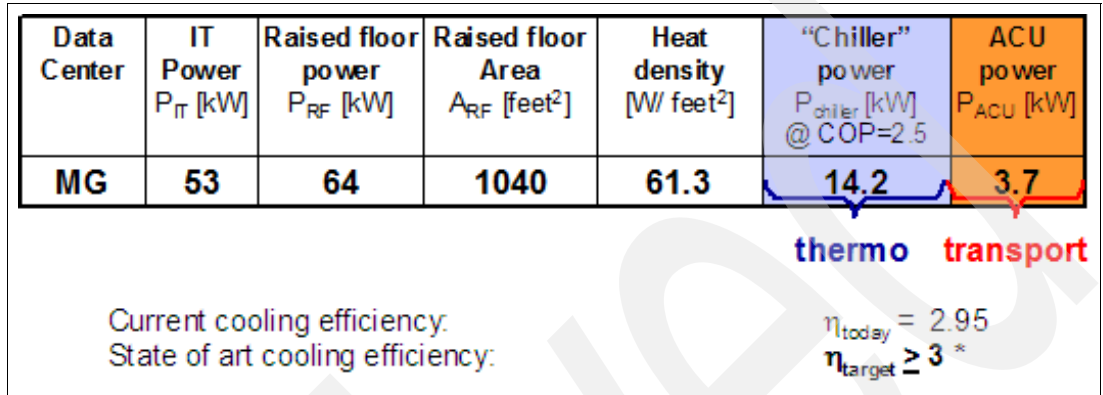


Figure 3 Chiller and ACU power consumption and cooling efficiency

Figure 4 shows an MMT scan at Z=0.5 feet with American Society of Heating, Refrigerating and Air conditioning Engineers (ASHRAE) recommendations.¹

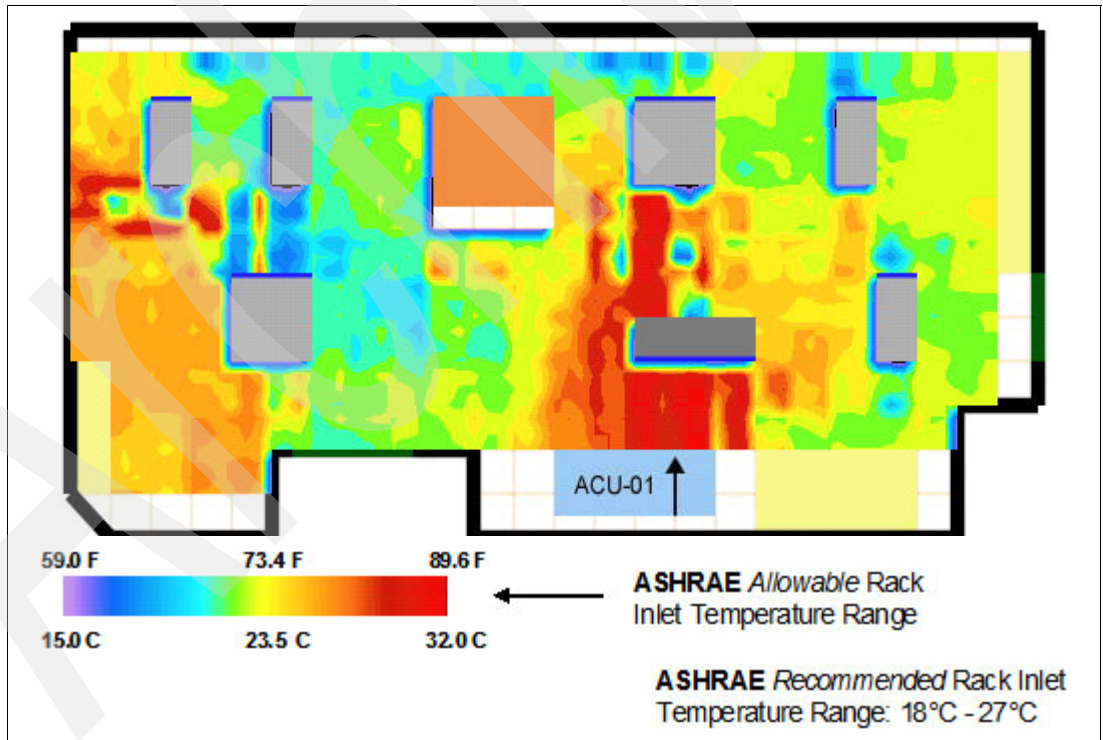


Figure 4 MMT scan at Z=.5 feet

At a height of just 0.5 feet, there are already high inlet temperatures on the storage unit directly in front of the ACU. The IBM System Storage™ DS8000 server requires inlet air both in front and back, which makes it difficult to cool in an open area without isolation.

¹ Z denotes the vertical height in the raised floor of the data center.

Figure 5 shows an MMT scan at Z=5.5 ft.²

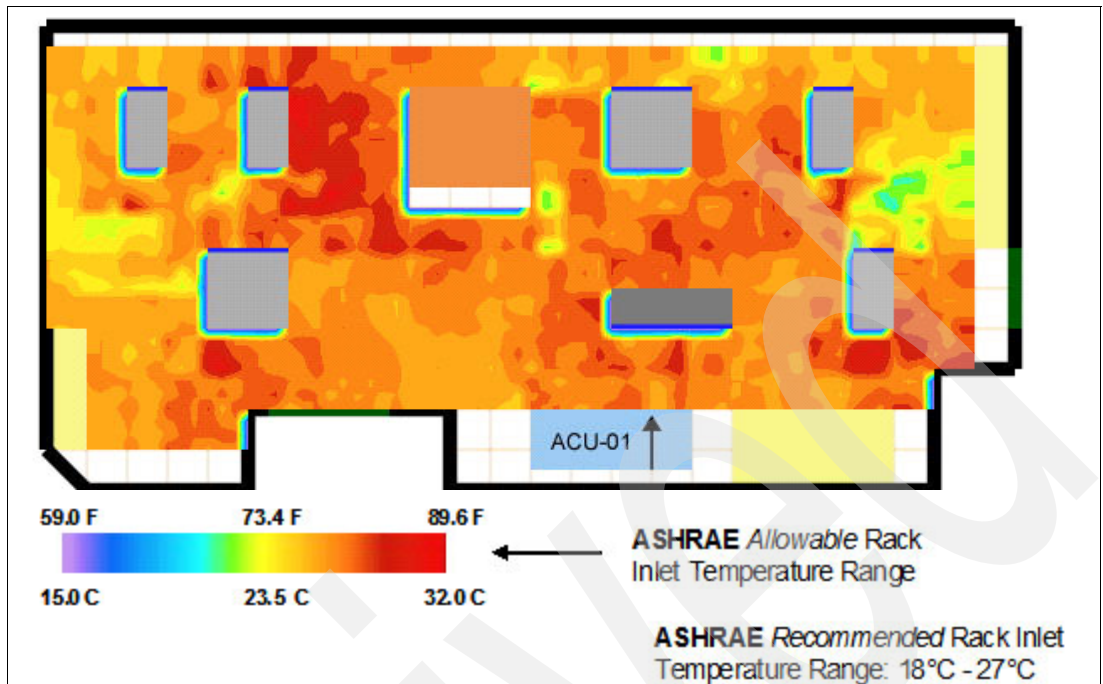


Figure 5 MMT scan at Z=5.5 feet

The MMT performs scans in height intervals of 1 foot. Note that there is a significant temperature increase in the rack inlet temperatures when comparing Figure 4 on page 6 and Figure 5. In addition, note the significant recirculation on the servers in the top left corner. The inlet area in some areas of the scan has exceeded the ASHRAE recommended rack inlet temperatures.

² See note 1 on page 6.

Figure 6 shows inlet temperatures greater than 27°C (80.6°F). Three of the server racks and the storage unit in front of ACU-01 have exceeded ASHRAE recommended temperatures.

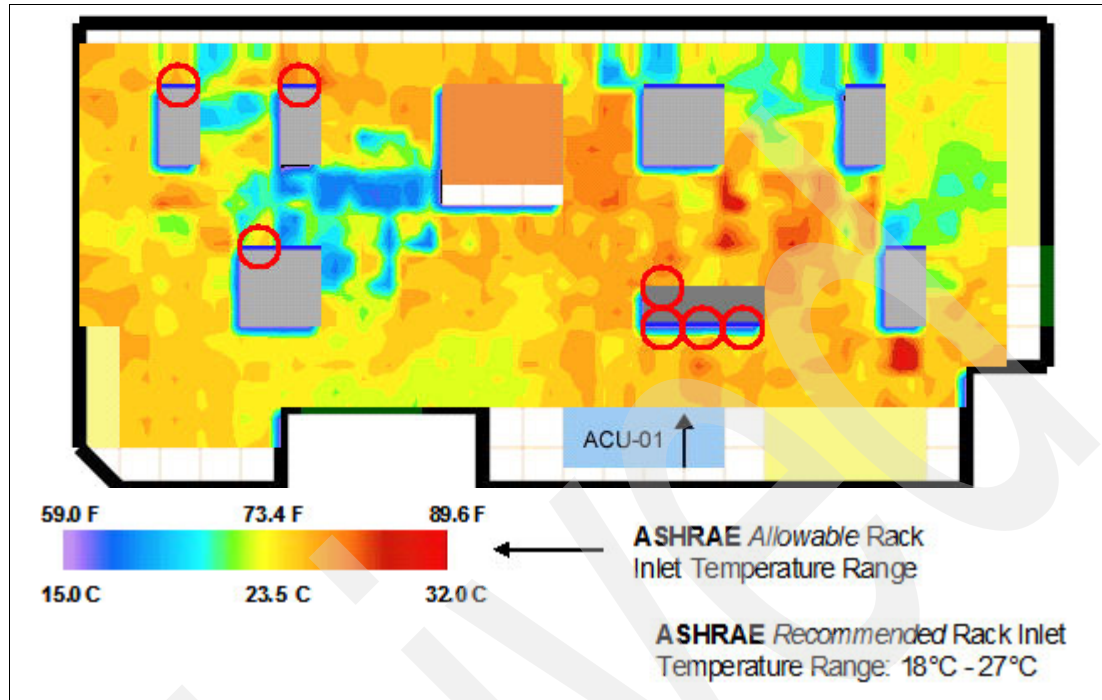


Figure 6 Inlet temperatures showing inlets over 27°C (80.6°F) circled in red

From the MMT survey data, the following conclusions can be drawn:

- ▶ The ACU utilization is 77%.

This metric measures the ratio of the heat removed by an ACU over the nominal ACU removal capacity. This utilization is high and approaching the maximum load for the ACU.

- ▶ The targeted air flow is 83%.

Targeted airflow is the percentage of air from ACUs directed through perforated tiles near the rack inlets.

- ▶ The Inlet Hotspot Factor is 2.9°F (low) with a mean temperature of 76.4°F (high), standard deviation (sd) = 5.7°F.

Inlet hotspots identified IT equipment with the hottest inlet temperatures. The high inlet temperatures on the three servers shown in Figure 6 are caused by significant air recirculation. Exhaust air from the servers is bounced back into the inlet air by the transparent wall in front of the servers. A low value for the Inlet Hotspot Factor is desired. However, the concern is the high mean temperature and standard deviation. This distribution should be tightened up to avoid having any inlets on the high side exceed the ASHRAE maximum recommended value of 81°F.

- ▶ The Horizontal Hotspot Factor is 15.5°F (high value).

The Horizontal Hotspots Factor uses the average rack inlet temperatures to compare the temperature distribution throughout the data center. This metric quantifies the overall temperature spread across the entire data center space. It is used to identify areas that are over-provisioned and under-provisioned with cooling air. A high value means that more of the recommended rack inlet temperature range is being used. If this range is tightened by better air distribution in the space, cooling efficiency is improved. Here, a lower value is also desired.

- ▶ The Vertical Hotspot Factor is 0.6°F (low value).

The Vertical Hotspot Factor analyzes the temperature difference between rack top and bottom inlets throughout the data center. This metric identifies areas affected by recirculation of warm air from rack exhausts to server inlets either coming over the top or around the ends of racks. There is a low value because, on the racks that have been identified as having inlet hotspots, the inlet air temperatures close to the floor are already high, which indicates a lack of sufficient flow coming up from the cool air plenum under the raised floor. This flow results in a low top to bottom of the rack temperature gradient and a low Vertical Hotspot Factor. Here, a higher value is desired.

- ▶ Inlet temperatures are high on the DS8000 storage unit at all measured heights and on three server racks at heights above 5.5 feet.

The high inlet temperatures on the DS8000 storage units are caused by the storage unit requiring inlet air in front and back and not having isolation from exhaust air from other servers in the data center.

Computational Fluid Dynamics analysis

A goal of the Poughkeepsie Green Data Center was to increase IT capacity in the data center and to allow for future IT growth. An inventory of the IT equipment identified servers that could be consolidated into higher density racks in the data center. The target inventory included nearly 100 kW in IT load and included the following configuration:

- ▶ Nine BladeCenter chassis with 111 blade servers
- ▶ One System z10™ Enterprise Server
- ▶ One System z9 Enterprise Server
- ▶ Three DS8000 Enterprise Storage Servers
- ▶ One XIV Enterprise Storage Server
- ▶ Two TS3100 Tape Libraries
- ▶ Five DS4000 servers with 12 TB of storage
- ▶ Two IBM POWER6 p570s
- ▶ Four SAN Volume Controllers
- ▶ Fifty System x rack mounted servers

A number of factors contributed to the design decisions for the data center. For example, the ACU was nearing maximum utilization. The layout of the data center was not conducive to adding an additional ACU to the room. Business needs did not allow for downtime that might be required for construction in the data center.

Another factor is that water through Rear Door Heat Exchangers (RDHx) was used to remove a significant portion of the heat from the distributed high density racks. There are several reasons for using water to support the additional IT load. Water is 3,500 times more efficient than air at removing heat, and the Cooling Distribution Unit (CDU) that supports the RDHx uses less energy than an ACU.

To determine the best possible layout of the equipment, the thermal properties of the data center were modeled by using *Computational Fluid Dynamics* (CFD). CFD is a suite of numerical analysis techniques for fluid mechanics that are used to evaluate fluid flow conditions. Computers are used to perform the millions of calculations required to simulate the interaction of fluids and gases in complex environments. Only approximate solutions can be achieved in many cases, even with the use of supercomputers. Approximations are made in the boundary conditions to provide an estimated solution for a good starting point.

Over forty simulations were run to optimize the floor layout for the IT equipment. CFD modeling of the room was performed to help optimize the layout, locate racks and tiles, determine the effectiveness of the gap filler, try new ideas, and analyze the impact of using

the RDHx. Many simulations were run to ensure that the best possible results were achieved for the IT equipment. This practice was critical because the room only contained one ACU, which already had a calculated utilization of 77%.

Figure 7 shows the recommended layout based on best practices. Figure 8 on page 11 and Figure 9 on page 12 show the CFD analysis and changes made to the proposed layout in Figure 7.

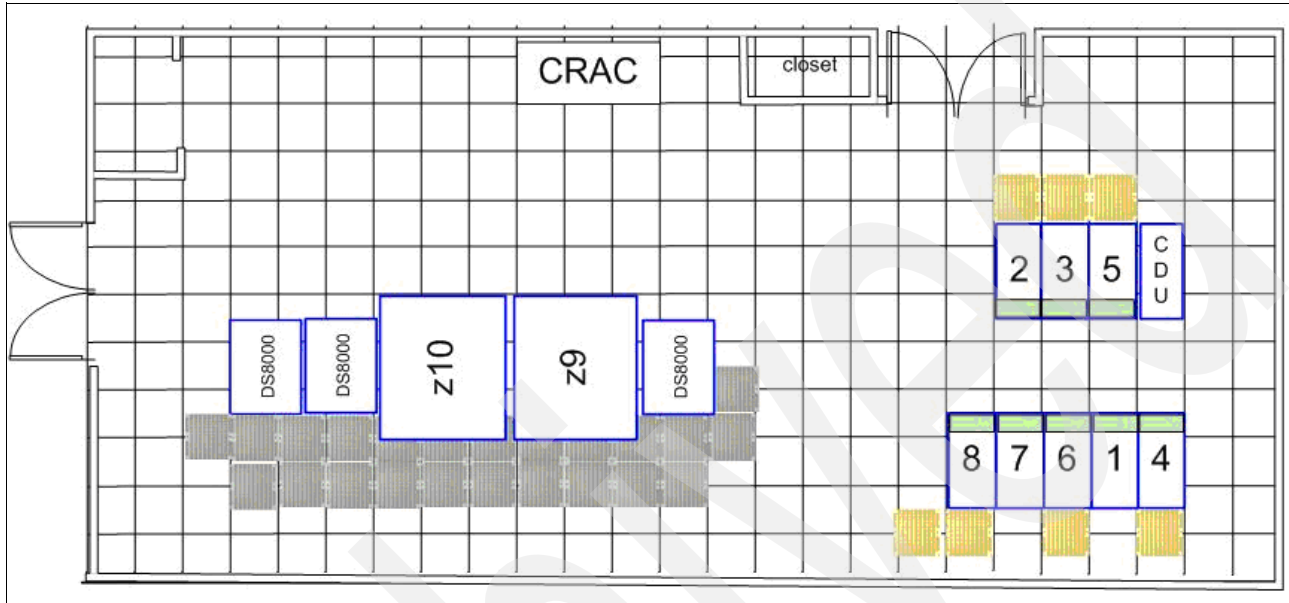


Figure 7 Phase 1 data center

The layout in Figure 7 implements design best practices of the data center. Racks 1 and 3 through 8 use the RDHx to remove the majority of the heat. These racks are situated such that the exhaust heat is in the hot aisle between the racks. The mainframes (IBM System z10 and z9) are situated such that their exhaust is directed back toward the ACU to provide best heat exchanger efficiency. The DS8000 storage racks are situated such that they receive cool air intake (require both front and back) and minimize any hot air recirculation from the z10 and z9.

Using this model, the next step was to perform a CFD analysis to determine if the layout was optimum for the IT configuration, perforated tiles/percent open, and inlet and exhaust temperatures. Figure 8 shows the CFD analysis of the Phase 1 layout.




Figure 8 CFD analysis of the Phase 1 layout

The CFD modeling shows that all IT components receive proper cooling with the exception of the storage units next to the z10 and z9 mainframes. Recirculation of the hot air that affects the inlet temperature of the DS8000 storage units still exists. This analysis shows that the hot outlet air from the mainframes must be isolated so that it does not impact the inlet air temperature of the DS8000 storage racks.

Figure 9 shows that, by removing one of the DS8000 storage units and putting a deflector between the mainframe and the DS8000 storage unit, hot air is prevented from recirculating back to the DS8000 storage unit. The DS8000 demands unique isolation because it requires inlet air at both the front and rear of the rack. Therefore, care must be taken to ensure that hot exhaust air is not in front of the DS8000 rack units. The CFD model in Figure 9 is the recommended solution for Phase 1 of the Poughkeepsie Green Data Center. It uses the RDHx liquid cooling solution to remove most of the heat gained by the air exiting the backs of racks 1 through 8. This solution prevents this heat load from increasing the room air temperature and placing even more load on ACU-01. Each RDHx can remove up to 15 kW/rack. The ACU is the primary cooling element for the System z9, z10, and DS8000 storage racks.




Figure 9 CFD analysis of the Phase 1 layout with a baffle

After the recommended Phase 1 layout was completed, a new MMT analysis was performed to verify the results of the CFD modeling. The second MMT verified that the hotspots and high temperatures identified in the first MMT were addressed. We also verified the accuracy of the CFD modeling. The validation is discussed in the following section.

Enhancements to the data center facilities

After the Poughkeepsie Green Data Center was renovated to support the new IT equipment, a final MMT analysis was performed to verify how the changes made to the data center improved the thermal performance of the space. Major renovations included hot and cold aisles, a baffle to prevent recirculation to the System Storage DS8000 racks, LED lights, uncluttered raised floor, and liquid cooling by using seven RDHx units on racks 1 and 3 through 8.

Table 1 summarizes the measurements.

Table 1 Before and after renovation power measurements

Power (kW)	Before	After	After FreeCool
IT	52.73 ^a	91.73 ^a	91.73 ^a
Computer Room Air Conditioner (CRAC)	4.56 ^a	5.1 ^a	5.1 ^a
CDU	0	2.6	2.6
Lighting	1.19 ^a	0.3 ^a	0.3 ^a
Other	0.82 ^a	0.82 ^a	0.82 ^a
Data center total	59.3	100.55	100.55
Facilities	11.79	20.00	4.57
Power Usage Effectiveness (PUE)	1.35	1.31	1.15

a. Measured value

Note that in Table 1, *FreeCool* is the ability to use the environment for heat rejection in the site chilled water plant. FreeCool is available when the temperature of the Hudson River drops below 43°F. This temperature allows data center managers to power off the chillers and only power the pumps (reduces energy required by facilities). This table shows the advantages of using the “FreeCool” option.

Also note in Table 1 that PUE was improved slightly while increasing the IT load by 39 kW in the same space footprint. This improvement was achieved by the use of seven RDHxs and best practices.

Each rack had the following power measurements. Intelligent Power Distribution Units (PDUs) were used on racks 1, 4, 5, 6, and 7.

- ▶ Rack 1 (Storage) had 2801 Watts.
- ▶ Rack 2 (XIV Storage) had 6000 Watts.
- ▶ Rack 3 (5 BladeCenter chassis) had 11830 Watts.
- ▶ Rack 4 (System x storage) had 3453 Watts.
- ▶ Rack 5 (System p) had 3940 Watts.
- ▶ Rack 6 (2U System x racks) had 4094 Watts.
- ▶ Rack 7 (1U System x racks) had 4983 Watts.
- ▶ Rack 8 (4 BladeCenter chassis) had 8540 Watts.
- ▶ System z10 had 23354 Watts.
- ▶ System z9 had 10300 Watts (to be replaced with z10).
- ▶ DS8000 Storage had 5000 Watts.

The goal was to remove most of the heat load from racks 1 through 8 by the use of the seven RDHxs. A reduction was required in the overall inlet air temperature to all racks and to significantly reduce recirculation around the DS8000 racks caused by the z9 and z10 mainframes.

Table 2 shows the before and after results.

Table 2 Before and after results

Measurement	Before	After
Area (sq. ft.)	1040	1040
IT power density (W/sq. ft.)	50.7	88.2
RF power density (W/(sq. ft.))	57.0	96.7
CRAC airflow (CFM)	8101 ^a	8721 ^a
Perforated tile airflow (CFM)	6698 ^a	8225 ^a
Lost airflow (CFM)	1403	496
Average rack inlet temp (°F)	76.4 ^a	69.8 ^a
CRAC water flow (GPM)	43.6 ^a	43.6 ^a
CRAC air heat load (kW)	59.3	73.4
CRAC utilization (%)	76.9	95.2
Water heat load (kW)	0	27.2
PUE	1.35	1.31

a. Measured value

The renovation offered the following benefits:

- ▶ Increased IT power by 39.0 kW (a 74% increase)
- ▶ Increased IT power density by 37.5 Watts/sq. ft.
- ▶ Reduced the average rack inlet temperature by 6.6°F
- ▶ Increased CRAC air heat load by 14.1 kW (23% increase)
- ▶ Removed 27.2 kW out of 39 kW from racks 1 and 3 through 8 using the RDHx
- ▶ Improved the PUE by 0.04

The CDU can handle a load of up to 130 kW. Therefore, there is plenty of bandwidth to add additional IT power to the racks with the RDHxs. In the future, an RDHx is planned to be added to the XIV rack to reduce the cooling load on the ACU, because it is now at 95%. Another enhancement to the layout of the room is to reduce the lost airflow by further blocking leak paths.

Figure 10 highlights how valuable the CFD modeling was for the layout of the Poughkeepsie Green Data Center. It shows an MMT scan at Z=5.5 feet with all of the modifications done so far to the data center.³

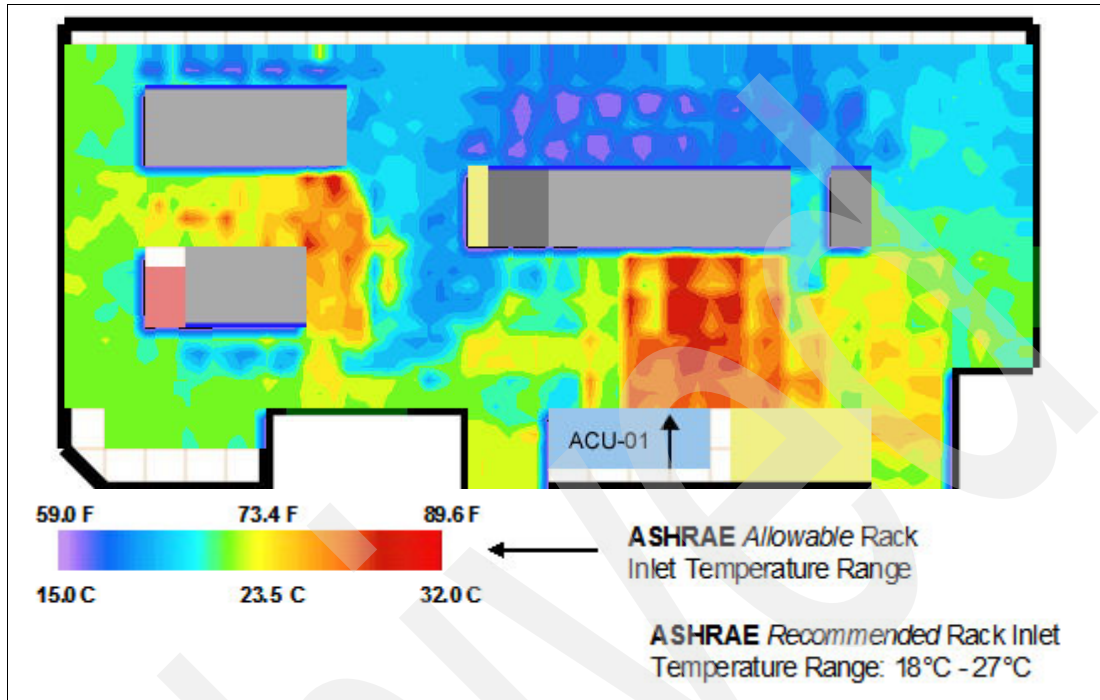


Figure 10 MMT scan after renovations at Z=5.5 feet

By comparing the MMT data in Figure 10 to the CFD modeling results of Figure 9 on page 12, observe that the temperature distributions match closely. The agreement between the MMT measurements and the CFD modeling validates the model. Note that the temperature profile is significantly better than the one that existed before implementing the recommended design changes as shown in Figure 5 on page 7. There is no recirculation in Figure 10, and the temperatures are much lower in the cool aisles. The results show that, even with a significant increase in the total IT heat load in the room, temperatures are maintained within the ASHRAE allowable limits at all rack inlets. These results can be attributed to the careful study of the original layout of the room, modeling of the possible layout changes, and implementation of the recommended layout following data center power and cooling best practices.

In addition to the data center facilities enhancements, energy management techniques are a key contributor to energy efficiency in the data center. The next section discusses these techniques and how they were implemented in the Poughkeepsie Green Data Center.

IT energy management techniques

In today's prevalent data center model, several trends lead to the oversizing of capacity up front. An application is dedicated to a particular server, which remains underutilized for its life cycle. However, typical x86 servers can consume more than 50% of average power when idle. As new models are released, while they might take up less space within the data center, they might use more power. Therefore, the data centers they reside in are actually becoming more outdated because of the changes in hardware. Although these trends have been identified as problems for years, there is a set of new drivers for change. These drivers are basic. Power

³ Z is the height above the raised floor.

and space are becoming limited, and energy costs are consuming a significant portion of IT budgets.

In a distributed computing environment, up to 85% of computing capacity might sit idle. These statistics point to an alluring opportunity inside every data center. Usage of that lost 85% of computational capacity can allow IT operations to capture a portion of the maintenance budget back to allow the investment in new capabilities. The challenge is increasing utilization without increasing operational costs in the same manner.

This challenge can be addressed through a collection of IT Energy Management techniques. Capacity planning is the first and most crucial step. By gaining visibility into current workloads, IT managers can plan for peak utilization across the data center. When the peak utilization and normal utilization are known, virtualization and consolidation can be used across the data center to maximize hardware utilization. Provisioning and cloud computing practices can be used to manage resources. Adopting cloud computing methodologies and integrating systems management solutions have great benefits to the data center. These techniques are discussed and their use in the Poughkeepsie Green Data Center is highlighted in the sections that follow.

Capacity planning

Capacity planning ensures that adequate resources are available just when needed, at the least possible cost. Like many IT functions, understanding the technology is only the beginning. Knowing how the business currently uses the applications and plans to use them in the future is as important as understanding the cost of each transaction.

It is common for IT organizations to manage system performance in a reactionary manner. However, capacity planning is more cost effective and efficient if performed prior to deployment for several reasons. Performance problems resulting from a lack of capacity are more complex and costly to resolve after deployment. Additionally capacity planning provides the information needed to strategically assess future IT requirements for new equipment, additional network capacity, and new, underlying IT architectures. The primary goals of capacity planning are to provide an accurate picture of the return on investment (ROI) for the IT organization and provide satisfactory service level agreements (SLAs) to users.

Virtualization

The goal for platform virtualization is to increase utilization of the distributed environment while increasing ease of management to address resource consumption of workloads. Ideally the capacity planning exercise has given insight into workload and how much consolidation can be achieved while maintaining service level requirements. This consolidation can provide the same workload capacity while reducing the number of physical platforms. In turn, this virtualization can provide more floor space and less power usage.

Typical server utilization is approximately 10%-15%. With virtualization, utilization rates as high as 80% or greater can be achieved. Server consolidation goes hand in hand with virtualization by reducing physical platforms and more efficiently utilizing them.

This consolidation leads to increased space in the data center. This increased space provides more power due to reduced machine count. Virtualization promises improved continuity and disaster recovery through workload mobility by moving around a failed resources, and virtualized servers can be mirrored in case of failure. Fewer physical servers in the data center leads to reduced power and cooling costs, not to mention savings on server hardware and related maintenance over time.

Consolidation case study

The Poughkeepsie Green Data Center effort replaced an aging infrastructure server farm with new technology to reduce maintenance and depreciation costs, reduce infrastructure complexity, increase availability, and more efficiently use server and data center resources (power, cooling, and floor space).

One consolidation involved replacing 18 servers with two IBM Power 570 systems. Each system is configured with multiple logical partitions (LPARs) by using virtual I/O, micro-partitioning for CPU sharing, and logical host Ethernet adapters. IBM System Storage San Volume Controllers were implemented to more efficiently use and manage storage pools.

The resources hosted on these servers include shared file servers, name servers, installation servers, and Web servers. Additional infrastructure components were added to the new servers as a result of using virtual I/O and micro-partitioning. The benefits of this consolidation are more efficient and smarter utilization of system, facility, and people resources; increased performance and availability; and ease of system management.

Additional consolidation focused on System z mainframes. The Poughkeepsie Green Data Center used the virtualization technology of the System z server to consolidate physical servers onto logical servers to reduce the physical footprint while growing center capabilities.

The Poughkeepsie STG Client Center's missions were hosted on three robust IBM System z9 Enterprise Class (EC) mainframes. The objective was to consolidate two z9 EC systems to a single IBM System z10 EC to reduce the physical footprint while growing business capacity and capabilities.

The System z mainframes offer the following technologies, which enabled the Poughkeepsie Green Data Center to accomplish consolidation objectives:

- ▶ Processor Resource/System Manager (PR/SM™), which enables LPARs to share processor and I/O resources
- ▶ Multiple Image Facility (MIF)
- ▶ Multiple channel subsystem (MCSS)

These technologies enable I/O channel paths and the subsets of I/O devices attached to those channel paths to be shared among LPARs. Because of these technologies, the consolidation was enabled while gaining the following benefits in the data center:

- ▶ Headroom to host more logical servers and projects
- ▶ Physical space
- ▶ IT capacity (network and cooling)
- ▶ Overall LPAR performance
- ▶ Ease of LPAR management and monitoring on fewer physical servers

Provisioning

Provisioning can be used to automate the deployment of operating system images, middleware, and software. Additionally it can be used for patch management, user provisioning, network provisioning, and storage provisioning. Provisioning takes the onus of managing physical resources off of administrators and allows them to focus on managing IT services instead of systems.

For the Poughkeepsie Green Data Center, IBM Tivoli® Provisioning Manager was used in conjunction with a Web portal for users to create a self-service provisioning solution. This portal allowed users to request and provision Linux® for System z resources.

Many organizations face challenges in managing their existing IT resources, evaluating their effectiveness in supporting business needs, and determining how to meet future needs with existing or new resources. To effectively respond to change, financial pressures, and competition, businesses need to maintain an environment that is flexible, responsive, and optimally uses available resources.

Combined with virtualization best practices, provisioning allows IT administrators to focus their time on managing service levels and availability rather than software and machine management. This combination is an important step on the road to cloud computing.

Cloud computing

The cloud computing vision suggests that it is possible to achieve more with less while reducing costs to free operation budgets for new investment. Cloud computing can be roughly divided in three ways:

- ▶ Infrastructure as a Service
- ▶ Platform as a Service
- ▶ Application as a Service

Infrastructure as a Service is the practice of using virtualization, automation, and standardization to effectively manage the physical resources and virtual machines within the environment. The cornerstone of effectively managing Infrastructure as a Service is process and requirements generated during the capacity planning stage.

Managing in this manner is bolstered by homogenization or standardization within the data center. If the hardware platforms are like systems (same architecture, operating systems, and middleware), this practice is greatly eased. The goal is to have a few instances of validated operating system configurations that can be cloned to provide a platform for deployment.

Having a small pool of images allows a finite set of patches to manage. A good practice for image management is to have two versions of an operating system: the current production version and an upcoming development version. All systems in production should be on that level, and all testing and new development should take place on the development version. Given a regular release cycle and rigorous patch management, all systems in the data center can be kept at the same level.

Hardware should be managed as a pool and applications should have resource allocations that can maintain the required SLAs. Then when a new application or expansion of an existing service takes place, the administrator can allocate the required resources from the pool and deploy the new service.

Moving up the stack, the next level of service is *Platform as a Service*. Platform as a Service is similar to Infrastructure as a Service but includes operating systems and required services that focus on a particular application. For example, Platform as a Service, in addition to virtualized servers and storage, provides a particular operating system and application set (typically as a virtual machine file) along with access to necessary services or specialized local resources. That is, Platform as a Service is Infrastructure as a Service with a custom software stack for the given application.

Providing *Application as a Service* capacity within the data center can be addressed in two ways: as a service for developers and testers to request short-term resources or as a line of business offering for customers of IT to request applications. In Application as a Service, a service catalog is offered to the user to specify at least the application and duration they need access. In some configurations, users can also specify underlying details such as CPU,

memory, disk, and software allocations. Application as a Service methodologies allow the administrator to focus on providing services and maintaining SLAs.

The procurement model is also altered in an Application as a Service model. Rather than paying for a system, administrative cost, and hosting cost for an application, the client is instead charged on actual usage of the resource. This model requires sophisticated monitoring and metering. Not only do utilization metrics support the expense model, but they also feed the capacity model and allow the IT systems management to calculate the actual resource expenditure for a given service.

Constructing a robust service level model, workload utilization model, and pooled hardware allows system administrators to manage resources to meet goals based on service level. This management means that low priority workloads (test and development resources) can be migrated to slower hardware or moved out of scope to free resources for workload spikes or high utilization periods.

Management tools, such as Tivoli Service Automation Manager, can help reduce IT operational costs by automating the processes used to deliver and manage a cloud computing environment. With Tivoli Service Automation Manager, IT managers can respond rapidly to business demands for basic IT services such as computing resources and application middleware service deployments. In addition, clients who have provisioning tools and processes in place, but need to automate those processes, can benefit from Tivoli Service Automation Manager. In addition, organizations that seek to provide self-service provisioning abilities for users or request driven provisioning can also benefit from Tivoli Service Automation Manager.

In the Poughkeepsie Green Data Center transformation, a precursor of Tivoli Service Automation Manager, *SOA OnRamp*, was used to implement a rudimentary cloud computing environment. The solution included a user-facing self-service portal built on IBM WebSphere® sMash, which executed provisioning requests through Web services to Tivoli Provisioning Manager. Tivoli Provisioning Manager then executes provisioning workflows that allocate System z resources from a pool and assigns them to the newly provisioned resources. To create the Linux guests within z/VM®, *SOA OnRamp* uses a component of z/VM Center for IBM Systems Director.

In Phase 2 of the Poughkeepsie Green Data Center transformation, provisioning of Linux on System z is performed exclusively through Tivoli Service Automation Manager. The front end of Tivoli Service Automation Manager adds user authentication and authorization as well as improved resource management. In addition to an improved user experience, the IT Systems Management process is greatly enhanced. Pool assets and provisioned resources are tracked through Tivoli Change Control Management Database (CCMDB). Provisioned resources are monitored through IBM Tivoli Monitoring. Many deployments also integrate IBM Tivoli Usage and Accounting Manager to implement chargeback for allocated guests.

Cloud computing provides a panacea of management within a service level requirement model as opposed to managing in a reactionary model. Advanced monitoring and metering of the data center can be managed against business requirements with a focus on the true cost of a dynamic infrastructure data center.

Measuring, monitoring, and managing energy in the data center

Energy measurement and monitoring tools play an extremely important role in the dynamic infrastructure data center. This section discusses the SynapSense tools, IBM Systems Director Active Energy Manager™, and the IBM Tivoli tools used in the Poughkeepsie Green Data Center.

IBM Systems Director Active Energy Manager

Within the Poughkeepsie Green Data Center, the IBM Systems Director Active Energy Manager is used to help implement part of the smarter planet and dynamic infrastructure vision of IBM by monitoring utilization of all IT equipment. IBM Systems Director Active Energy Manager Version 3.1.1 measures, monitors, and manages the energy components built into IBM systems enabling a cross-platform management solution.

Active Energy Manager extends the scope of energy management to include facility providers to enable a more complete view of energy consumption within the data center. It is an IBM Systems Director extension that supports the following endpoints:

- ▶ IBM BladeCenter
- ▶ IBM Power Systems™
- ▶ System x servers
- ▶ System z servers

IBM storage systems and non-IBM platforms can be monitored through PDU+ support. In addition, Active Energy Manager can collect information from select facility providers, including Liebert SiteScan from Emerson Network Power and SynapSense. The Active Energy Manager server can run on the following platforms:

- ▶ Microsoft® Windows® on System x, Linux on System x, Linux on System p
- ▶ Linux on System z

Active Energy Manager uses agentless technology. Therefore, no agents are required on the endpoints.

Monitoring and management functions apply to all IBM systems that are enabled for IBM Systems Director Active Energy Manager V3.1.1. The monitoring functions include power trending, thermal trending, PDU+ support, and support for facility providers. The management functions include power capping and power savings mode.

In the Poughkeepsie Green Data Center, IBM Systems Director 5.20.3 and Active Energy Manager 3.1.1 FP2 were installed on a System x346 running Red Hat® Linux Server 5.3. The tool is currently monitoring nine PDU+ connected to rack mounted IBM System x, Power Systems, and TotalStorage products; nine BladeCenter chassis with over 90 Blade servers; and two IBM System z10 Enterprise Servers.

Figure 11 shows IBM Active Energy Manager displaying the power and thermal trends from an IBM System z10 Enterprise Server.

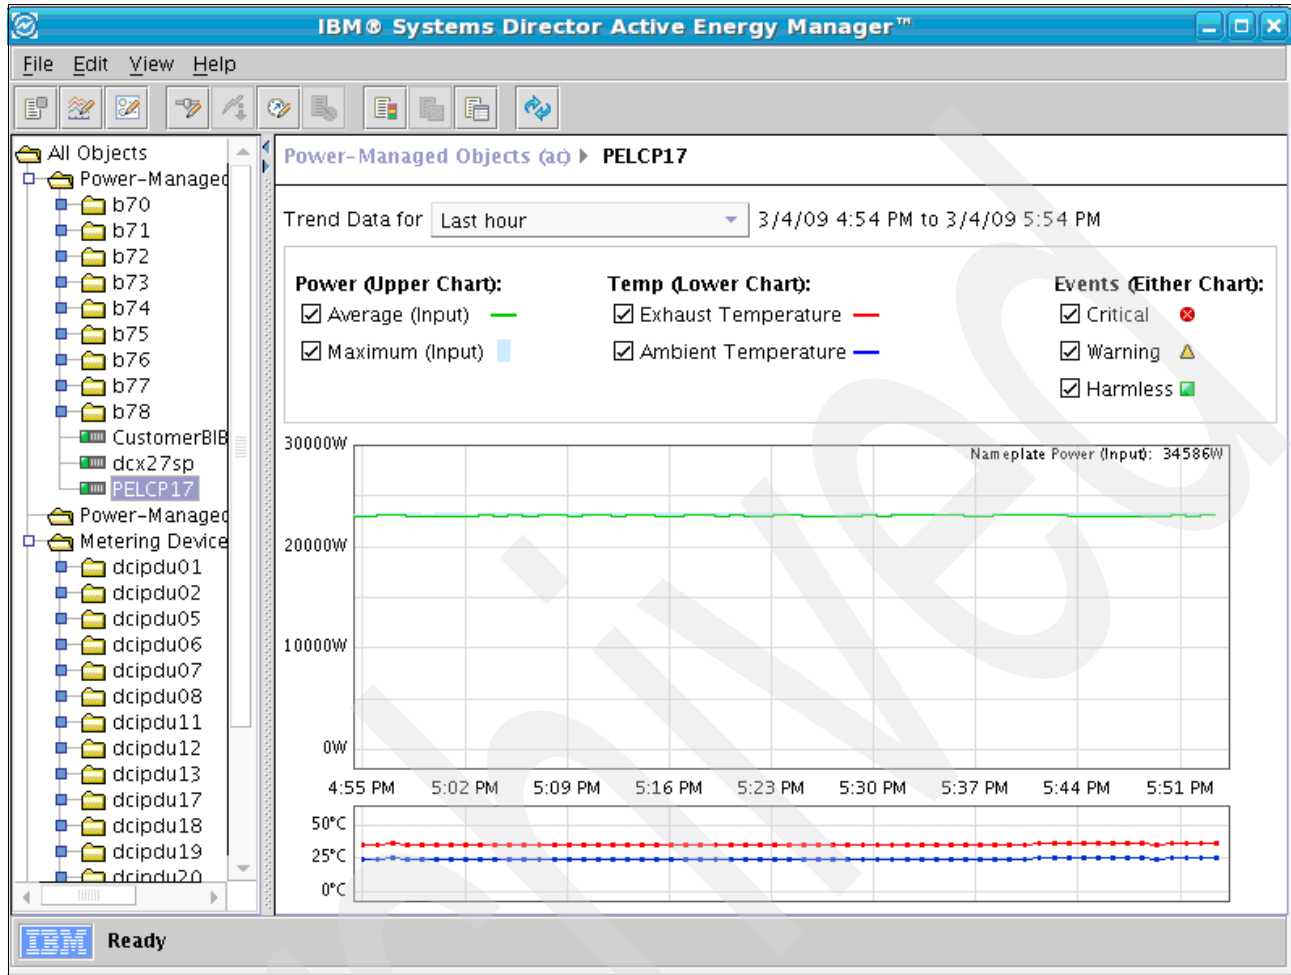


Figure 11 Systems Director Active Energy Manager showing power and thermal trends for System z10 Enterprise Server

Active Energy Manager also provides a source of energy management data that can be exploited by Tivoli enterprise solutions such as IBM Tivoli Monitoring and IBM Tivoli Usage and Accounting Manager. It is a key component of the IBM Cool Blue™ portfolio within Project Big Green.

Enterprise monitoring with IBM Tivoli Monitoring

Within the Poughkeepsie Green Data Center, the IBM Tivoli Monitoring product is used to help implement part of the smarter planet and dynamic infrastructure vision of IBM. The sophisticated agent technology of IBM Tivoli Monitoring has been deployed to instrument the diverse server platforms within the Poughkeepsie Green Data Center infrastructure, allowing critical performance and load statistics to be collected to indicate the overall health of the environment.

Statistics on CPU utilization, disk performance and capacity, process status, and system load are being tracked and aggregated by agents of IBM Tivoli Monitoring across BladeCenter servers and System z, Power Systems, and System x environments. The agents collect key performance and health information across z/OS®, AIX®, Linux, and Microsoft Windows operating systems. From a monitoring perspective, all of these systems are interconnected. Their status, health, and performance are being collected and intelligently displayed through the IBM Tivoli Monitoring Tivoli Enterprise Portal.

In addition to monitoring all of the traditional server based resources within the data center, IBM Tivoli Monitoring for Energy Management also uses its agent technology to provide interconnection to gather critical power and cooling statistics from the facilities-based element management systems. IBM Tivoli Monitoring with IBM Systems Director and the power and thermal instrumentation provided by Active Energy Manager are interconnected within a data center. As result, the Tivoli Enterprise Portal console can integrate both the server statistics and the facilities data into a single intelligent, coherent view that represents the overall health of the environment.

Figure 12 on page 23 shows how IBM Tivoli Monitoring interconnects with all of the power and facilities management components within a data center by using integration with IBM Systems Director Active Energy Manager.

IBM Tivoli Monitoring within the Poughkeepsie Green Data Center

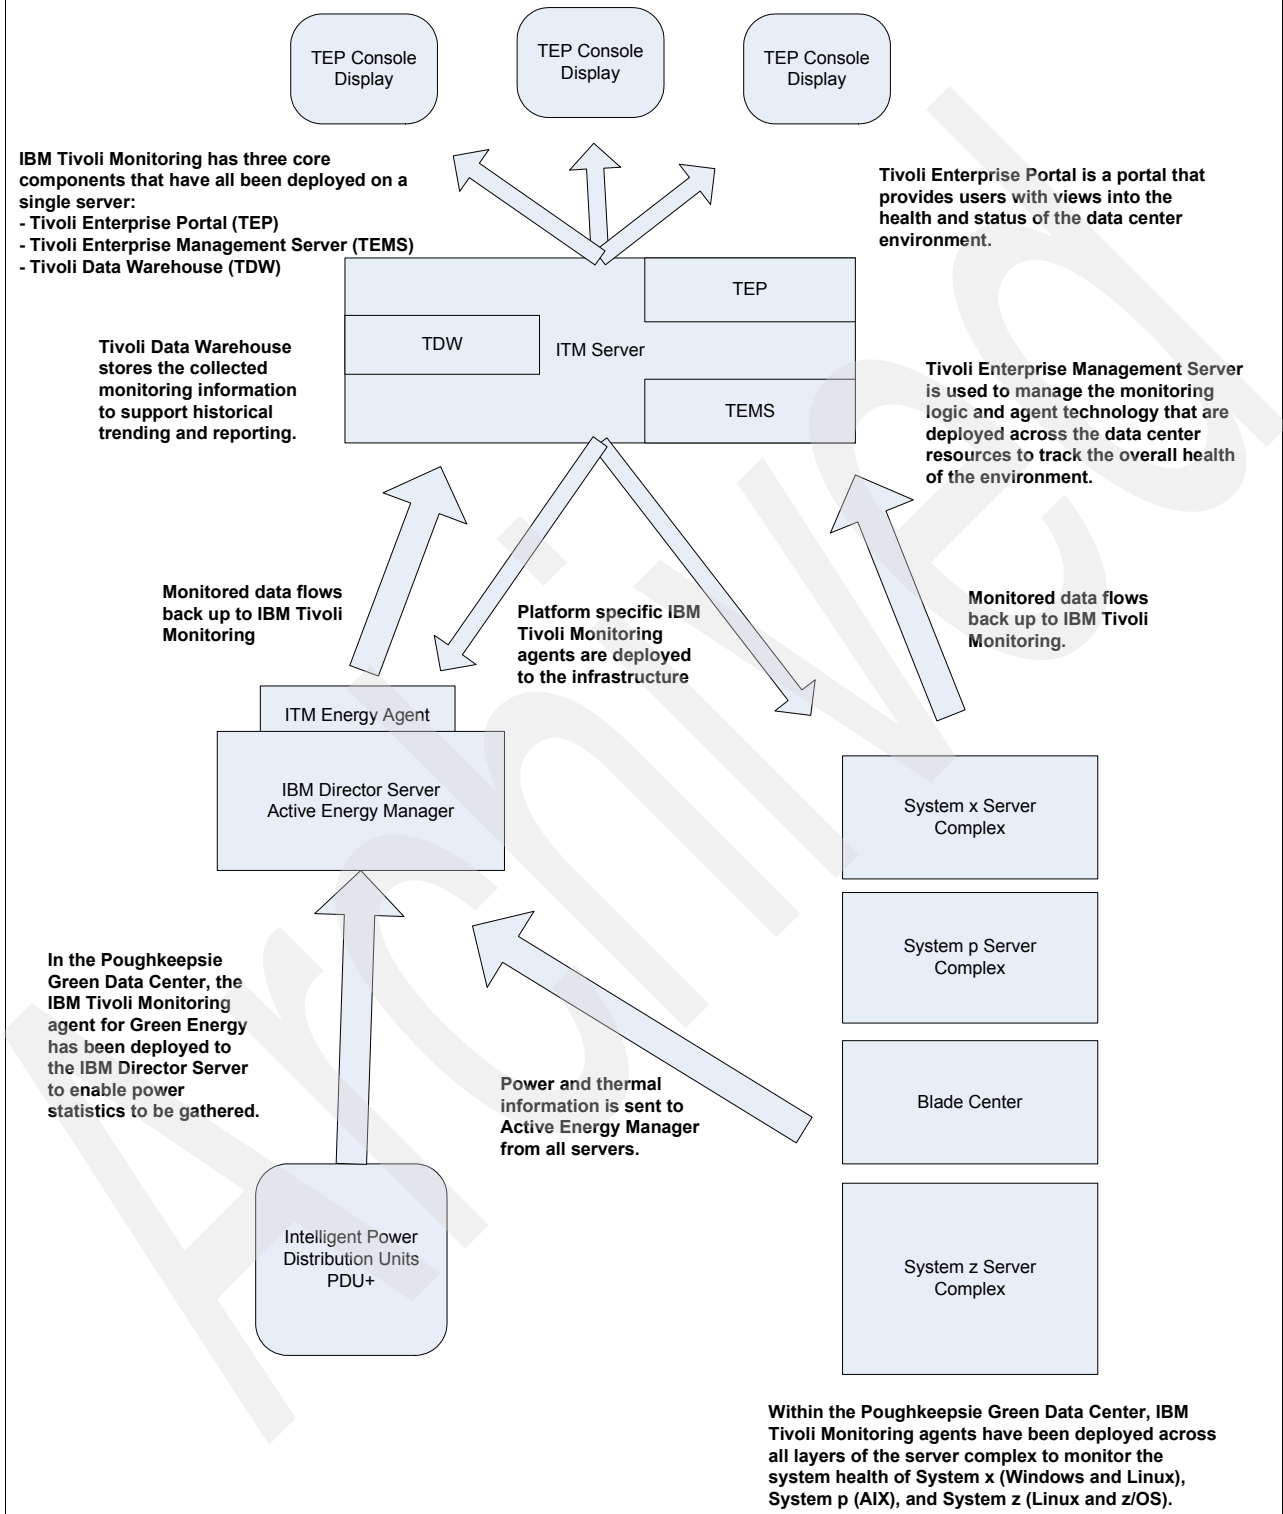


Figure 12 IBM Tivoli Monitoring and Active Energy Manager – Data collection

In the Poughkeepsie Green Data Center, IBM Tivoli Monitoring and IBM Systems Director Active Energy Manager have been implemented in this manner. In addition, the Tivoli Enterprise Portal display has been configured to provide many views into the overall health of the environment. Depending on the role of operations support personnel, different role based views can be set up as the default. For example, a Linux Systems Administrator can be presented with only views that are specific to the Linux systems that are being monitored.

Figure 13 shows the integrated view of the Tivoli Enterprise Portal consoles of all the resources being monitored within the data center. This view includes Power metrics in the two panes in the lower left side, along with views into System z, Power Systems, System x, and the various operating systems. The panels across the top provide message views into the environment.

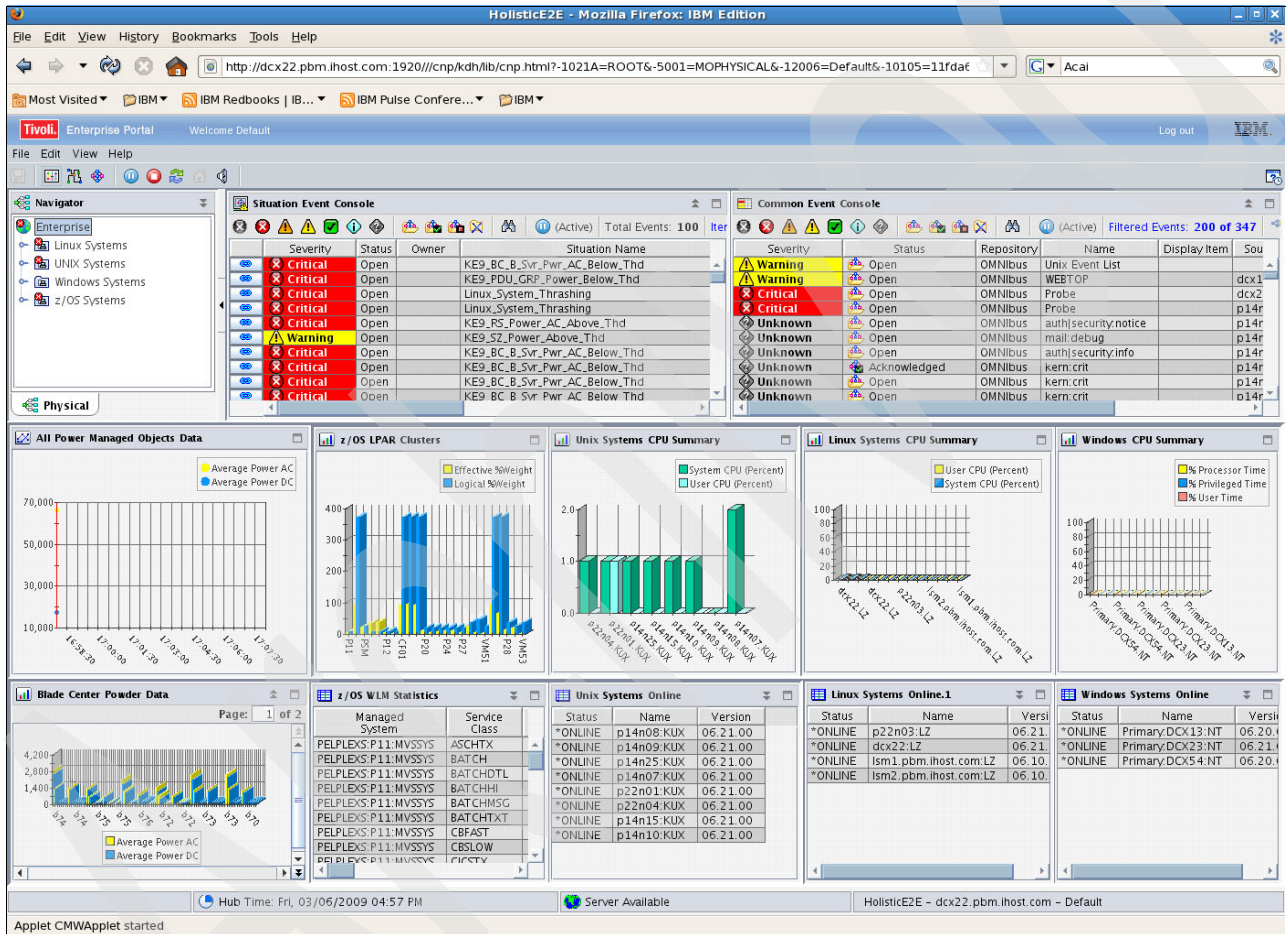


Figure 13 Integrated view of the Tivoli Enterprise Portal consoles

By using the Tivoli Energy Agent and the Active Energy Manager data provider, the monitored power and thermal data are extracted and placed into the repositories of IBM Tivoli Monitoring. Figure 14 shows the IBM Systems Director Active Energy Manager and Tivoli Energy Agent architecture. The Tivoli Active Energy Manager Data provider extracts the requested power and thermal trends and forwards the data stream up to IBM Tivoli Monitoring.

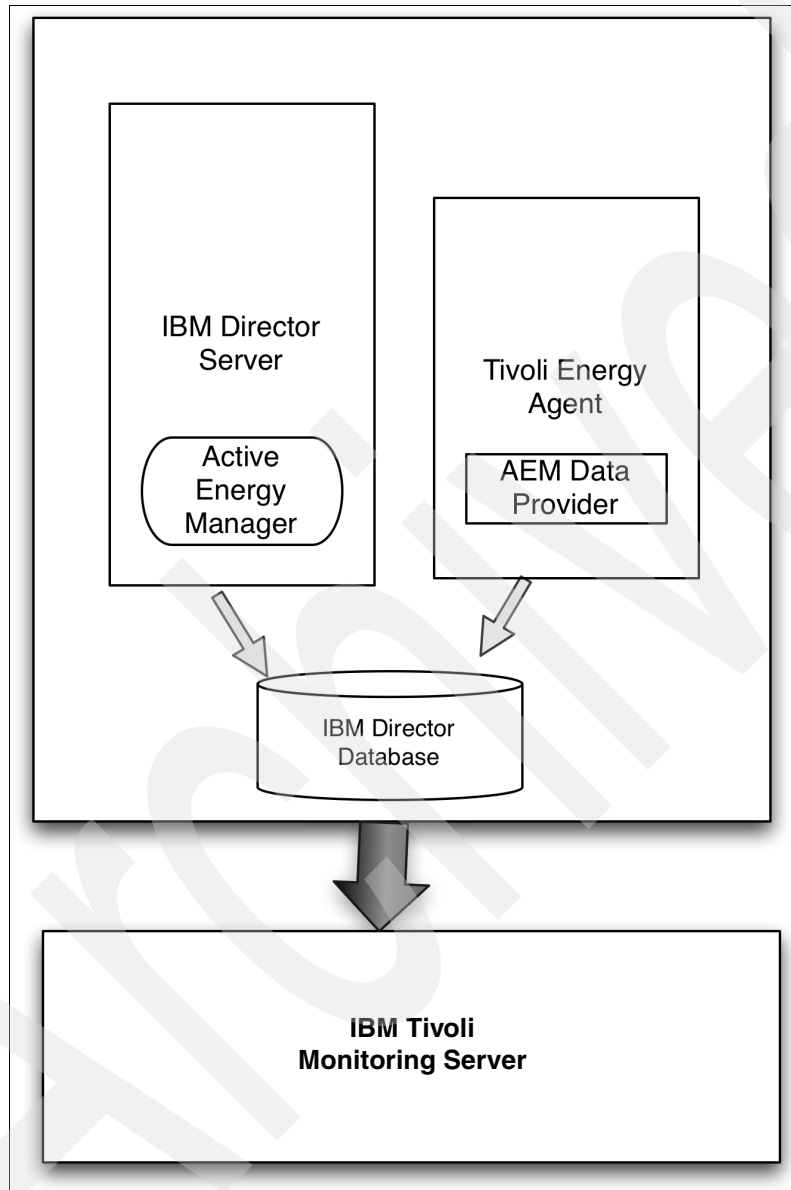


Figure 14 Architecture of IBM Systems Director Active Energy Manager and Tivoli Energy Agent

Active Energy Manager data that is sent to the IBM Tivoli Monitoring server can be displayed by choosing the server that has IBM Systems Director Active Energy Manager installed from the list of IBM Tivoli Monitoring agent systems. For example, the IBM Systems Director Active Energy Manager system installed at the Poughkeepsie Green Data Center is running Linux. Therefore, the Active Energy Manager data is under the Linux servers in Tivoli Enterprise Portal.

Figure 15 shows the BladeCenter Power trends being collected and displayed at the Poughkeepsie Green Data Center from some of the nine BladeCenter chassis.

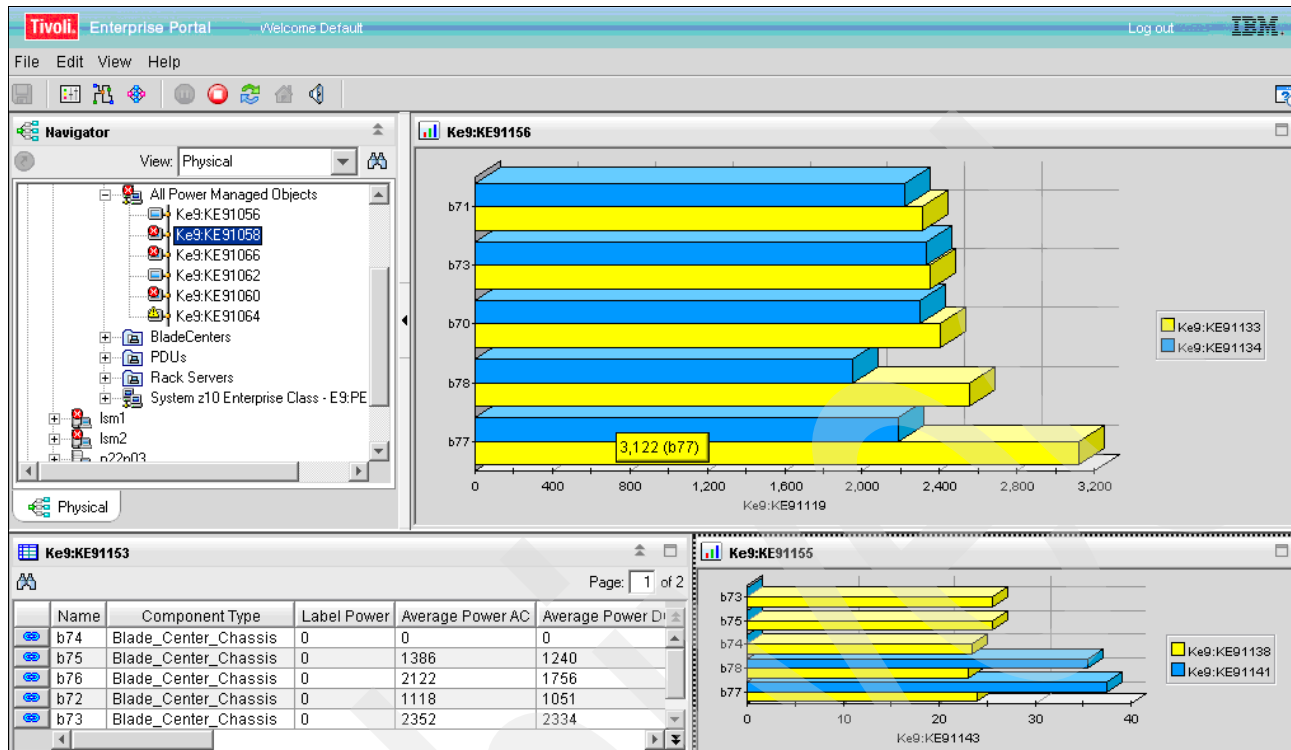


Figure 15 BladeCenter Power workspace in IBM Tivoli Monitoring

Event management using Netcool Omnibus

IBM Tivoli Netcool® Omnibus is being used in the Poughkeepsie Green Data Center to provide intelligent event management. All the servers in the data center produce unsolicited messages about their hardware, operating systems, middleware, network connections, storage subsystems, and applications. The majority of this information is composed of log files that report successful operation. However, error and warning messages are embedded with the informational messages. Netcool Omnibus is used to sort through the messages and present them to the appropriate users or forward them to other tools for integration.

Netcool Omnibus uses a three-tiered architecture to accomplish event management. Figure 16 on page 27 shows the Netcool Omnibus architecture. In the collection tier, *probes* are deployed to gather events from various data sources. The probes can suppress informational messages and forward warnings and errors. In the aggregation tier, the *object server* categorizes events and suppresses duplicates. In the display tier, the *desktop client* is used to view errors and warnings. A browser-based client is also available by using the *Webtop server*. All of these components were deployed in the Poughkeepsie Green Data Center.

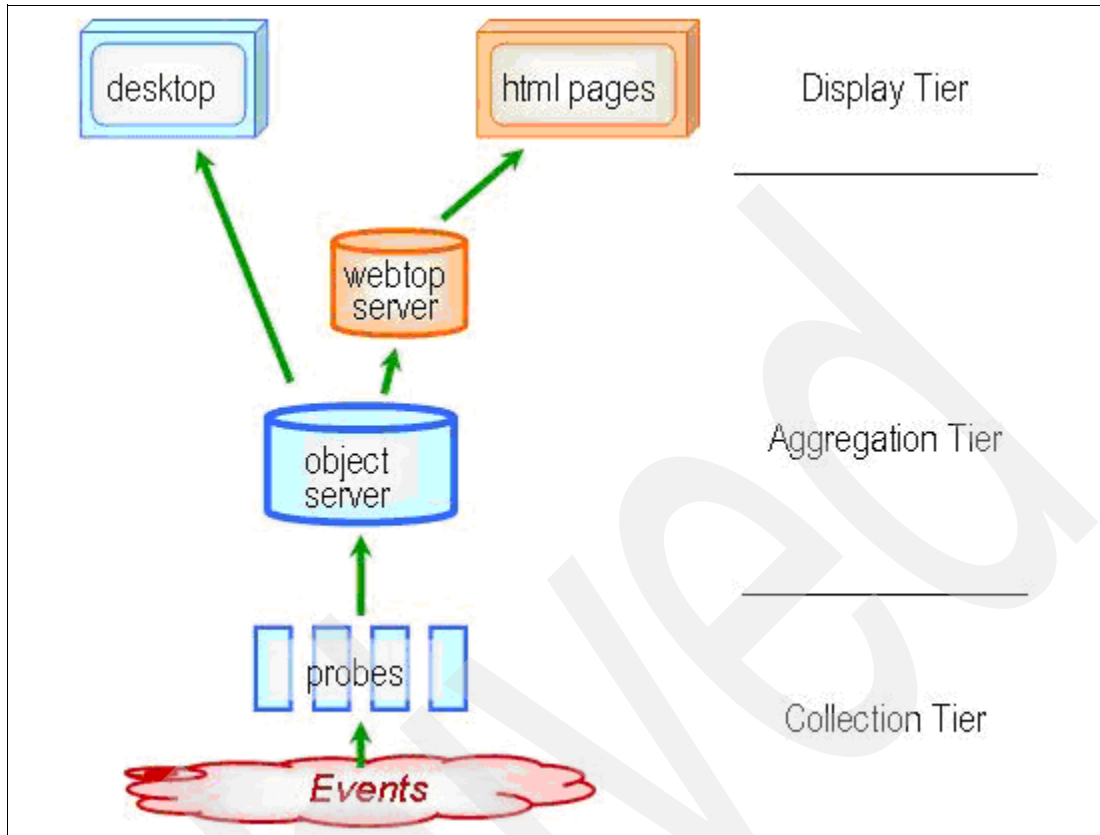


Figure 16 Netcool Omnibus architecture

In the Poughkeepsie Green Data Center, Netcool Omnibus probes have been deployed to AIX, Linux, and Windows servers to collect events. AIX and Linux Syslog probes were deployed to collect operating system events written to the syslog and forward errors and warnings to the Netcool Omnibus object server for processing. Likewise, Windows event log probes were deployed to collect events written to the event log and forward errors and warnings to the Netcool Omnibus object server.

Figure 17 shows the Netcool Omnibus Webtop client. Events collected by probes are displayed. Events are color coded by severity.

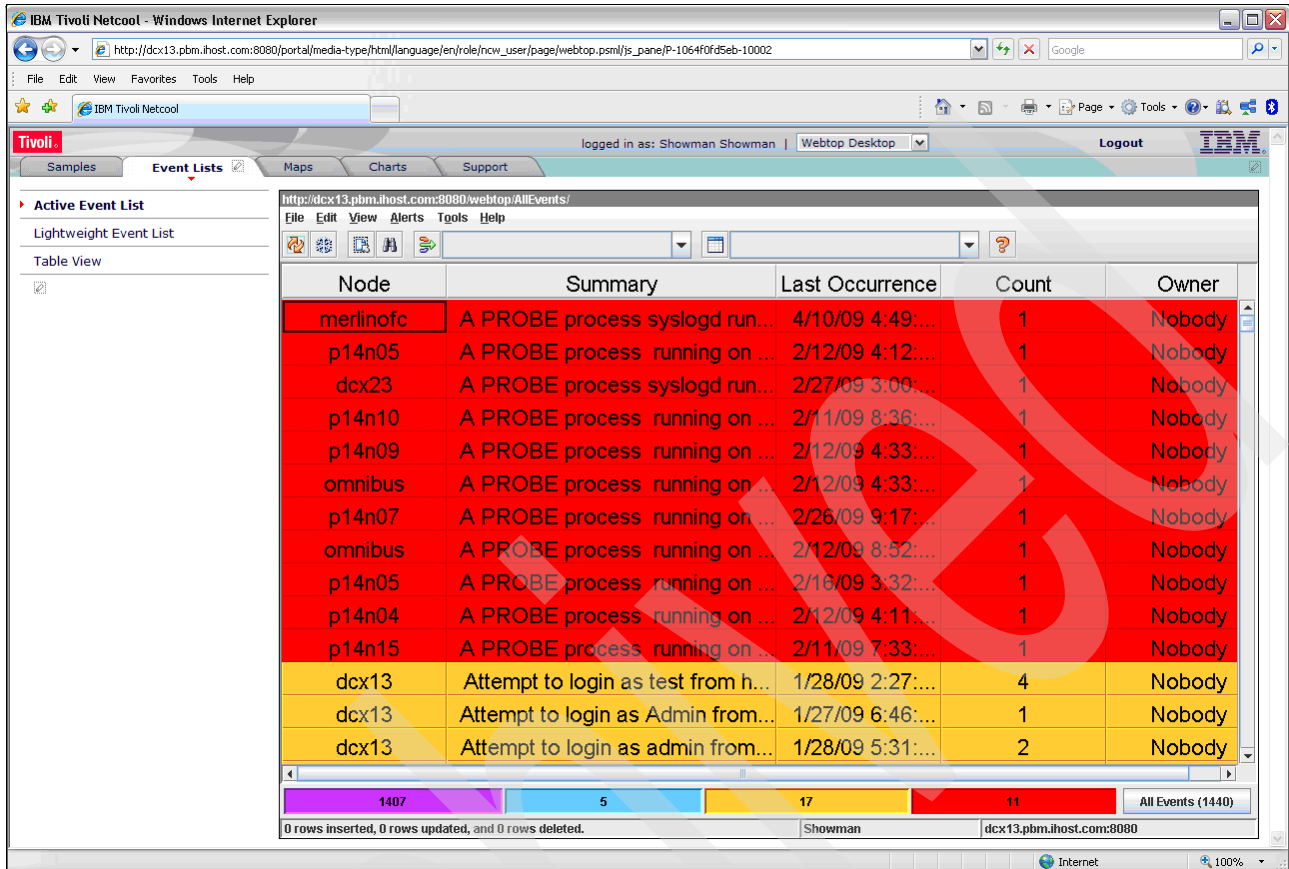


Figure 17 Netcool Omnibus Webtop client

The following event attributes are indicated along the top of the Event List:

- Node** Identifies the name of the source specified in the Probe. Syslog and Event Log Probes display the host name.
- Summary** Provides the text of the error or warning. Cryptic messages can be reformatted to display meaningful information.
- Last Occurrence** Provides a time stamp of the message. If duplicates were received, the time stamp of the most recent is displayed.
- Count** The number of events of the same type from the same source. This feature allows for duplicate message suppression.
- Owner** Users can take ownership of events. Multiple users can work to resolve problems simultaneously.

Each attribute can be used to sort the event list. By selecting the Node attribute, the display re-sorts the event list by Node. Figure 18 shows how sorting by Node allows a user to focus on errors and warnings associated with a single server, dcx13. As problems are resolved, events are closed and leave the event list. Data center staff can work smarter by focusing on discrete resources to resolve problems.

Node	Summary	Last Occurrence	Count	Owner
dcx13	A OBJSERVAUTH process ru...	3/26/09 6:16:...	2	Nobody
dcx13	A WEBTOP process running o...	3/2/09 7:02:5...	1	Nobody
dcx13	A WEBTOP process running o...	4/8/09 1:17:0...	1	Nobody
dcx13	Attempt to login as Admin from...	1/27/09 6:46:...	1	Nobody
dcx13	Attempt to login as Hilon from ...	3/31/09 2:39:...	2	Nobody
dcx13	Attempt to login as Netcool fro...	3/26/09 6:36:...	2	Nobody
dcx13	Attempt to login as Show from ...	2/4/09 9:18:1...	2	Nobody
dcx13	Attempt to login as Showman f...	1/27/09 6:45:...	7	Nobody
dcx13	Attempt to login as admin from...	1/28/09 5:31:...	2	Nobody
dcx13	Attempt to login as greendc fro...	2/11/09 8:10:...	2	Nobody
dcx13	Attempt to login as root from h...	3/31/09 2:12:...	20	Nobody
dcx13	Attempt to login as showman fr...	4/23/09 8:42:...	28	Nobody
dcx13	Attempt to login as test from h...	1/28/09 2:27:...	4	Nobody
dcx13	Attempt to login as test1 from ...	1/28/09 5:31:...	2	Nobody

Summary: 1407 (purple), 5 (blue), 17 (yellow), 11 (red), All Events (1440)

0 rows inserted, 0 rows updated, and 0 rows deleted. Showman dcx13.pbm.ihost.com:8080

Figure 18 Netcool Omnibus event list sorted by Node

Figure 19 shows how Netcool Omnibus is integrated with the IBM Tivoli Monitoring infrastructure. Events from Netcool Omnibus are now visible from the Tivoli Enterprise Portal, integrating event management with enterprise monitoring. The IBM Tivoli Monitoring Tivoli Enterprise Portal console can now display both IBM Tivoli Monitoring and Netcool Omnibus messages, which allows users to have a consolidated view for smarter systems and service management.

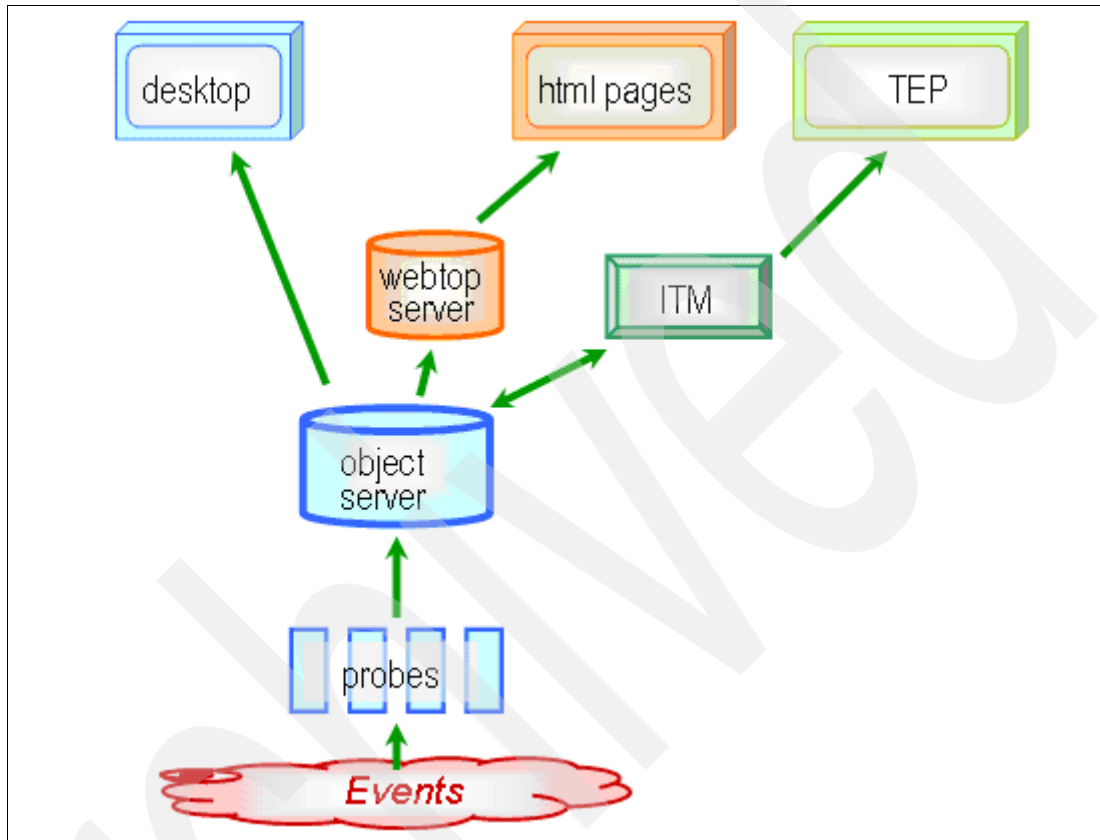


Figure 19 Netcool Omnibus integration with IBM Tivoli Monitoring and Tivoli Enterprise Portal

Phase 1 deployment of Netcool Omnibus in the Poughkeepsie Green Data Center allows operating system log messages from AIX, Linux, and Windows systems to be collected, aggregated, and displayed. Future phases will address middleware, application, and network level monitoring through the deployment of additional Netcool Omnibus probes. Additionally, in future phases, Netcool Omnibus will be integrated with Netcool Impact, allowing for complex event correlation and automation.

The SynapSense solution

The SynapSense SynapSoft solution provides real-time monitoring of thermal conditions within a data center and was an integral component in the Poughkeepsie Green Data Center. This monitoring is accomplished through a series of wireless sensors, which communicate with the SynapSoft server. This data provides thermal maps of the data center, thus allowing data center managers and facilities operators to make changes to their floor plan and energy use, while reducing their risk of downtime and reducing energy costs.

The SynapSoft environment provides the following functions for the SynapSense infrastructure:

- ▶ Sensor Instrumentation
- ▶ Efficiency Consulting
- ▶ Containment Services
- ▶ Non-disruptive Installation
- ▶ SynapSoft Data Center Application Suite
- ▶ SynapSense Wireless Network
- ▶ Alarm and Alert System
- ▶ SynapSense Gateways
- ▶ SynapSense Server
- ▶ SynapSense Wireless Nodes

To collect thermal and pressure data for the SynapSoft server, thermal and pressure sensors were installed at the Poughkeepsie Green Data Center. Figure 20 shows the SynapSense floor plan indicating where pressure, humidity, and temperature sensors were installed. A total of eight rack sensors, one CRAC sensor, and two floor pressure sensors were installed.

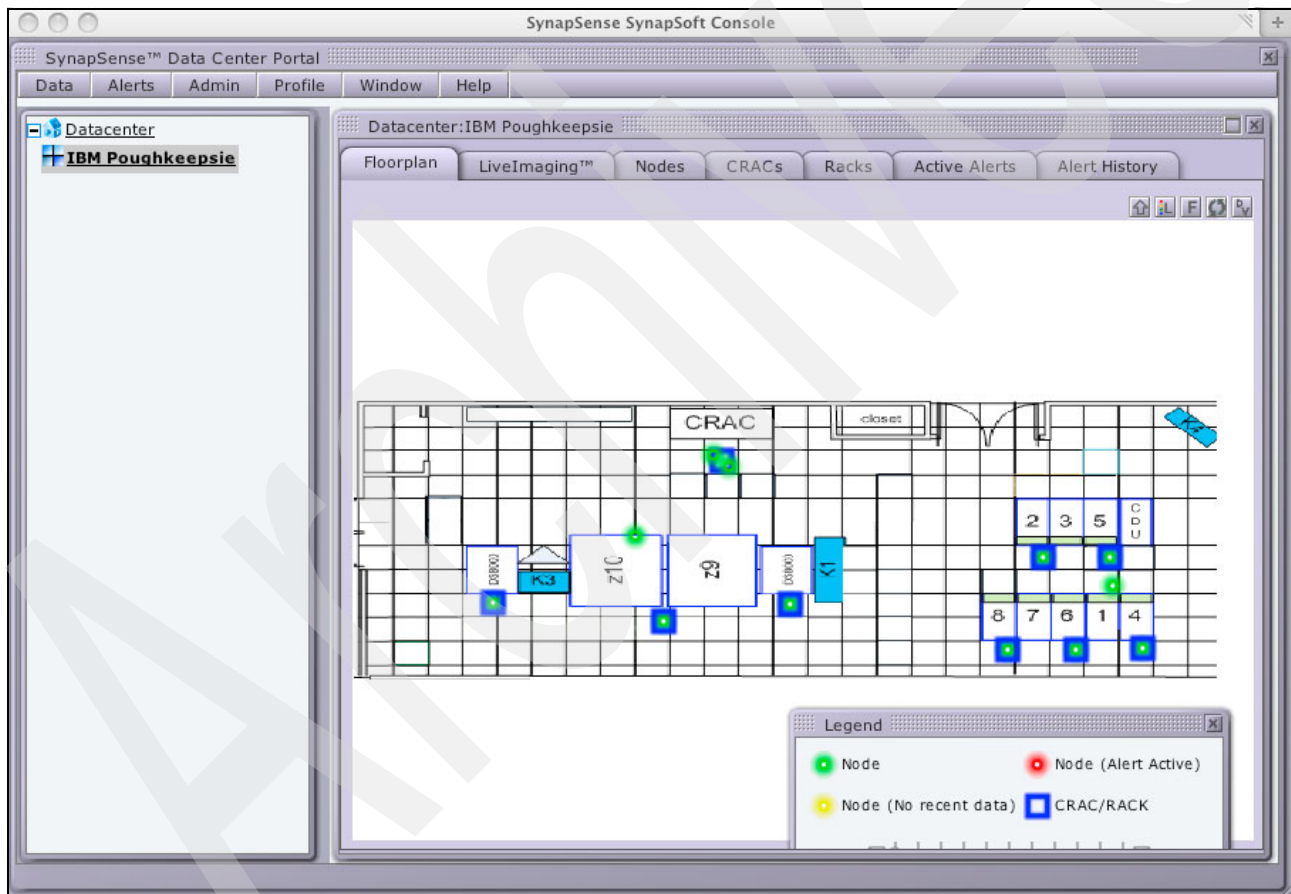


Figure 20 SynapSense Sensor Floor plan from the Poughkeepsie Green Data Center

SynapSoft 4.0 also features *LiveImaging*, the first real-time energy and operational efficiency solution that enables data center operators to literally see temperature, humidity, and pressure trends in their data center, and take corrective action with confidence. LiveImaging maps highlight specific environmental conditions by using 3-D real-time data at multiple levels of the data center. It also provides a visual display of actual examples of thermal air mixing, pressure differential, and humidity.

Figure 21 shows an actual LiveImaging thermal map that was obtained from the Poughkeepsie Green Data Center by using the SynapSense infrastructure installed at the center.

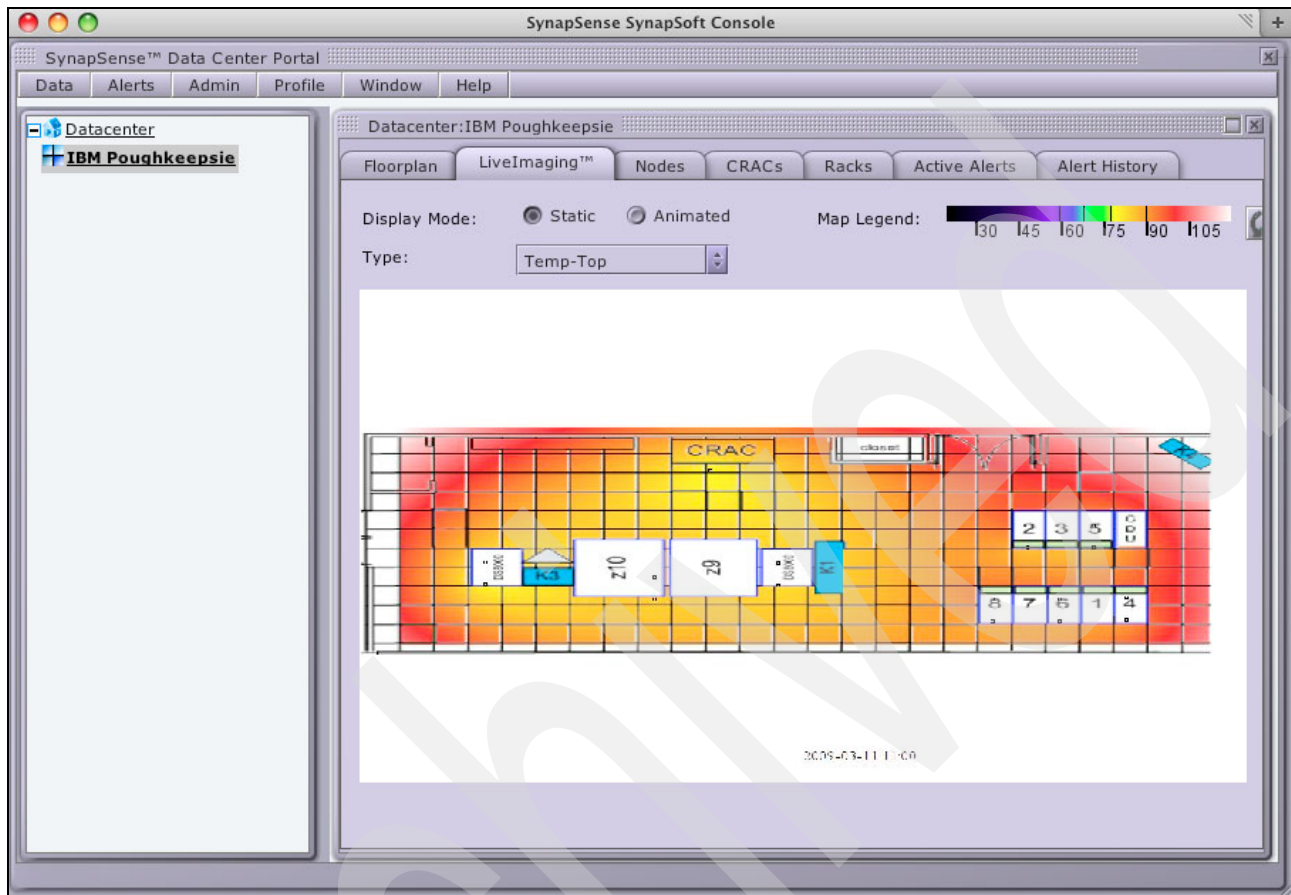


Figure 21 LiveImaging Thermal Map from the Poughkeepsie Green Data Center

This server also communicates with the IBM Systems Director and Active Energy Management Server to provide data center environmental data to Active Energy Manager by using Simple Network Management Protocol (SNMP). The SynapSense data can be displayed through the SynapSoft user interface or through Active Energy Manager.

Conclusion

By implementing the Poughkeepsie Green Data Center described in this paper, IBM was able to prepare for continued growth while adding cooling capacity. Through the use of high efficiency water cooling technologies, data center best practices, and energy monitoring and management techniques, the 100 kW load required is now being satisfied in the existing data center area with more headroom for future growth. Additionally, this upgrade was accomplished without any interruption to critical workloads. The results are significant. The ROI is less than one year, and the environmental impact is reduced by 7.5 kW per year of electricity annually.

IBM has led the technology industry in energy-smart innovation for over 40 years and is committed to climate protection. It is IBM's goal to sustain leadership in energy conservation and management by continuing to deliver power management and cooling technologies. With

these technologies, systems use less power, generate less heat, and use less energy to cool the IT environment.

In the dynamic infrastructure data center of the future, we should expect to see even more efficient system and data center designs. We should expect to see increased server energy efficiency techniques available such as power capping for all equipment and specialized power napping of system components. Performance benchmarks will incorporate even more sophisticated energy efficiency metrics. In addition, we will see an extension to energy efficiency tools that brings automation to the management and reporting of energy consumption by non-IT assets, such as an office building air conditioning system or streetlights in a city. With this software, organizations can visualize energy dynamics and then take appropriate action while extrapolating how changes will yield different business outcomes by using sophisticated “what if” calculations.

Through innovative technologies, energy management strategies, such as consolidation and virtualization, and energy management tools and techniques, it is apparent from the Poughkeepsie Green Data Center that a dynamic infrastructure data center is a truly outstanding energy efficient solution and will continue to be so in the future.

References

Consult the following resources for further information:

- ▶ Smarter Planet
<http://www.ibm.com/ibm/ideasfromibm/us/smarterplanet/20081106/index.shtml>
- ▶ Steve Lohr, “I.B.M. Has Tech Answer for Woes of Economy” New York Times, November 6, 2008
<http://www.nytimes.com/2008/11/06/technology/business-computing/06blue.html?ref=technology>
- ▶ IBM press release: “IBM Unveils Building Blocks for 21st Century Infrastructure”
<http://www.ibm.com/press/us/en/pressrelease/26550.wss>
- ▶ IBM press release: “IBM Global Financing Chief Advises C-Suite on Importance of Investing in 21st Century Technology Infrastructure”
<http://www.ibm.com/press/us/en/pressrelease/26610.wss>
- ▶ IBM Tivoli Monitoring
<http://www.ibm.com/software/tivoli/products/monitor/>
- ▶ IBM Tivoli Monitoring for Energy Management
<http://www.ibm.com/software/tivoli/products/monitor-energy-management/>
- ▶ Active Energy Manager
<http://www.ibm.com/systems/management/director/about/director52/extensions/actengmrg.html>
- ▶ SynapSense
<http://www.synapsense.com>
- ▶ Jerald Murphy, *Increasing Energy Efficiency with x86 Servers*, Robert Frances Group, 2009
<http://www-07.ibm.com/systems/includes/content/info/x/newgeneration/downloads/files.pdf>

- ▶ Lauri Minas and Ellison, Brad, "The Problem of Power Consumption in Servers," Intel®, April 01, 2009
<http://www.infoq.com/articles/power-consumption-servers>
- ▶ Operating system-level virtualization and platform virtualization
<http://www.ibm.com/systems/virtualization/>

In addition, consult the following IBM Redbooks® publications at:

<http://www.ibm.com/redbooks>

- ▶ *AS/400 Server Capacity Planning*, SG24-2159
- ▶ *Server Consolidation with VMware ESX Server*, REDP-3939

The team that wrote this IBM Redpaper

This paper was produced by a team of specialists from around the world working at the International Technical Support Organization, Poughkeepsie Center.

John Brady is a Certified Consulting IT Specialist with the Worldwide Client Centers of IBM Systems and Technology Group (STG) in Poughkeepsie, New York. He currently works at the Poughkeepsie-based IBM Design Center as a Complex Solution Architect and leads customer engagements to assist IBM clients in creating high level solution designs which incorporate the latest leading technologies. John has over 20 years of experience within the IT industry and has specialized in systems and service management. For the past 13 years, John has worked extensively with the Tivoli product suite and has achieved numerous Tivoli certifications including Enterprise Consultant and Advanced Deployment Professional.

Ira Chavis is a Certified Consulting IT Specialist with the Worldwide Client Centers. Working in the STG Center for Solution Integration (CSI), Ira currently specializes in infrastructure architecture and solutions involving IBM server and storage technologies. He has over 29 years of diversified software engineering and IT experience. Prior to working at IBM, Ira worked at Digital Equipment Corporation in varying assignments. Ira holds certifications as an IBM eServer™ Certified Expert in System x, IBM Grid Technical Sales, IBM Certified Expert in Storage Network Solutions, and Red Hat Certified Technician (RHCT™).

Matthew Finlayson is an IT Specialist working in the Worldwide Client Centers specializing in virtualization, consolidation, and cloud computing. He has had 11 years of experience in IT working as a system administrator, Programmer, Tester, and IT Specialist. He graduated from Clarkson University with a Master in Computer Science degree.

Michael Harkins is a Certified Consulting IT Specialist with the Worldwide Client Centers. Michael currently works at the Poughkeepsie-based IBM Design Center as an IT Architect and assists clients with adoption of new technologies. Michael has over 32 years of experience within the IT industry and has specialized in systems and service management for the past 22 years. Michael has worked extensively with the Tivoli product suite since 1995 and is certified in many products. Prior to working with systems management tools, Michael spent 10 years in mainframe and distributed data center operations environments.

John P. Mullin is an Advisory Software Engineer currently working with the IBM System z Benchmark Center in Poughkeepsie NY. He has been with IBM for 10 years. John has technical expertise in data center fundamentals and design; System z, System Storage, Power Systems, System x performance tuning, benchmarking, solution design and implementation, networking and interconnectivity; Tivoli monitoring and implementation, and load tool and workload sizing and implementation.

Julie Peet is a Senior Software Engineer in the Worldwide Client Centers. She has worked at IBM for 20 years, with 17 years in IBM Power technology related areas. She is currently the Infrastructure Team Leader for the Poughkeepsie Benchmark Centers.

Sheryl Qualters is a Project Manager in Worldwide Client Centers. She has 20 years of experience with IBM's server technologies and ten year of experience as a project manager. She has led projects covering STG server early programs, client benchmark centers, and industry solutions. She is an IBM and PMP certified Project Manager.

Rodrigo Samper is Senior Technical Staff Member in the Worldwide Client Centers with subject matter expertise in energy efficiency and green technologies. He leads client workshops to address client business and IT needs for data centers by improving service, reducing costs, and managing risks. His background includes over 25 years of development experience in the power, packaging, and cooling of IBM Personal Computers and System x servers.

Elisabeth Stahl is Chief Technical Strategist, in Performance Marketing, for the IBM STG and has been working in systems performance for over 25 years. Elisabeth is a Master Certified IT Specialist of The Open Group and an IEEE Senior Member. She holds a Bachelor of Arts (BA) degree in Mathematics from the University of Pennsylvania and a Master of Business Administration (MBA) from New York University (NYU).

Thanks to the following people for their contributions to this project:

Mike Ebbers
ITSO, Poughkeepsie, New York

Pat Coico
IBM Data Center Services, Poughkeepsie, New York

Archived

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.

This document REDP-4534-00 was created or updated on June 1, 2009.



Send us your comments in one of the following ways:


- ▶ Use the online **Contact us** review Redbooks form found at: ibm.com/redbooks
- ▶ Send your comments in an email 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 U.S.A.



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®	Netcool®	System x®
AS/400®	Netfinity®	System z10™
BladeCenter®	Power Systems™	System z9®
Cool Blue™	POWER6™	System z®
DS4000®	PR/SM™	Tivoli®
DS8000®	Redbooks®	TotalStorage®
Dynamic Infrastructure®	Redbooks (logo)  ®	WebSphere®
Enterprise Storage Server®	RS/6000®	XIV®
eServer™	S/390®	z/OS®
IBM Systems Director Active Energy Manager™	System i®	z/VM®
IBM®	System p®	z9®
	System Storage™	

The following terms are trademarks of other companies:

Red Hat, RHCT, and the Shadowman logo are trademarks or registered trademarks of Red Hat, Inc. in the U.S. and other countries.

VMware, the VMware "boxes" logo and design are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions.

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, other countries, or both.

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.