Aperture®

WHITE PAPER

Current Trends in Data Center Monitoring

Executive Summary

Data center monitoring has evolved from a simple alert system, a “nice-to-have” feature, into a critical element required for maximizing uptime and ensuring the increasingly complex environment within the data center is running effectively and efficiently. It has become vital for organizations to develop and deploy a system capable of monitoring the overall support infrastructure—including the power and environmental conditions of each room, equipment row and cabinet within the data center.

Many limitations exist with the traditional monitoring solutions in use today—including cost, operational complexity and limited focus on a select group of devices. What is needed are data center monitoring solutions that go beyond simply alerting data center management of immediate problems and actually enable real-time data to be captured and analyzed for management to make intelligent decisions regarding the data center infrastructure. The most effective solutions are grounded in best practices and integrate both the IT and Facilities groups into a single, comprehensive unit that provides actionable information in the decision-making process.

Evolution of Data Center Monitoring

In 1943, Thomas Watson, chairman of IBM, said "I think there is a world market for maybe five computers." Today, it is estimated that there are over one billion personal computers in use worldwide. Twenty years ago, minimal status information was available for only a small percentage of devices in the data center infrastructure. Today, over hundreds of data points, with metrics relating to everything from status to performance are available for the wide array of devices in the data center. The information now available is staggering—causing challenges in the data center to have evolved from access to insufficient information to an issue of how to comb through the vast amounts of data now available.

The data center has developed into an extremely complex, distributed environment with tremendous computing requirements not even imagined twenty years ago. Just as the data center itself has undergone a dramatic change, data center monitoring has also rapidly progressed. Gone are the days of measuring temperature by whether or not you need to wear a sweater when checking out the mainframe. Simplistic monitoring systems that report a good/bad condition have given way to a multitude of sophisticated protocols for determining the precise state of the critical devices in the data center.

The management structure of the data center has transformed as well. Two disparate groups—IT and Facilities—have arisen, each focused on complementary, yet opposing functions. The IT group encompasses the demand side of the business, requiring increasingly more infrastructure to run its servers, network hardware and storage devices; while the Facilities group, on the other hand, is responsible for supplying the power and cooling to run the data center. As noted by the U.S. Environmental Protection Agency (EPA) while addressing the challenges of energy efficiency within data centers, "In many data centers, those responsible for purchasing and operating the IT equipment are not the same people that are responsible for the power and cooling infrastructure, who in turn typically pay the utility bills. This leads to a split incentive, in which those who are most able to control the energy use of the IT equipment (and therefore the data center) have little incentive to do so,[1]" causing a significant silo effect within the organization.

The revolution and increased complexity of the overall data center environment has also been dramatically impacted by the introduction of new technologies—including blade servers, intelligent devices and virtualization—removing the commonplace practice of walking the raised floor space to "feel" the environment and conduct ad hoc monitoring. Not only is it impractical for data center personnel to manually assess the energy usage and cooling requirements, the current use of high-density devices generates such a massive amount of heat that affects all infrastructure resources; which is impossible to monitor without robust systems. The availability of energy resources is reaching its upper limits and the energy costs are soaring, causing organizations to consider repositioning their philosophical, strategic and technical principles for the data center to include more flexibility and more cohesive monitoring.

The overall evolution of the data center, both in complexity and structure, combined with the myriad of new intelligent devices that comprise the overall data center infrastructure necessitates the need for data center monitoring solutions to be seriously evaluated and upgraded.

[1] "Report to Congress on Server and Data Center Energy Efficiency Public Law 109-431," U.S. Environmental Protection Agency ENERGY STAR Program, August 2, 2007, p. 12.

Limitations to Traditional Monitoring in the Data Center

Multiple limitations exist with the traditional monitoring systems in place today. These limitations include operational complexity, limited support for devices in the infrastructure and an incomplete view of the data center as a whole. Most importantly, information regarding the various devices and applications are housed in so many different systems throughout the organization; making the monitoring and assessment of issues occurring within the data center ineffective and proactive evaluation of potential risks or forecasting the needs in the future impossible. In a survey recently conducted by the Aperture Research Institute™, “60 percent of respondents said that the ability for a monitoring system to work with equipment from multiple manufacturers ranked at least 4 on an importance scale of 1 to 5. Only 25 percent of the respondents had monitoring systems that could do this, however, with 53 percent monitoring less than 75 percent of their critical infrastructure and 30 percent monitoring less than 50 percent.”

Operational Complexity

The limitations due to operational complexity generally stem from the design of the monitoring solution—in particular, how multiple communication methods used in the data center are handled. The use of these widely varied communication methods greatly complicates the development of a solution capable of monitoring different devices with disparate protocols.

The two primary approaches to monitoring include a hardware-based and a software-based solution. The hardware solution requires a protocol to be selected as the primary communications method (Modbus and SNMP are the most widely used options). A hardware device—a protocol converter—then transforms the device protocol into the primary communications method protocol. The drawbacks associated with the use of a hardware solution include:

- Cost of the protocol converters (including the hidden costs of providing power and connecting them to a UPS for backup)
- Added complexity as more devices are added
- Introduction of another potential point of failure

In a software-based solution, the monitoring application communicates with the hardware device directly, using its native protocol; thereby eliminating the need for any protocol conversion.

MOST COMMON METHODS OF COMMUNICATION

Various methods are used to communicate with devices in the data center, with specific types for both the Facilities and IT groups, respectively.

Facilities

* *Modbus*
  The Modbus protocol, developed by Modicon in 1979, is used in master/slave configurations to provide communication with devices, with data and controls defined in a set of registers. Although most often implemented in a hard-wired, daisy-chain configuration using either RS-422 or RS-485, Modbus also supports wireless and networked (Modbus TCP over Ethernet) implementations. Industry analysts estimate more than seven million Modbus devices in use on thousands of different devices from hundreds of manufacturers today.

* *BACnet*
  The BACnet (Building Automation Control network) protocol was designed specifically to manage building applications—including heating, air conditioning, lighting, physical access and fire detection systems. Proposed in 1987 but not widely accepted until 1996, BACnet can be implemented in various ways—including network (BACnet/IP), point-to-point over RS-232 and in master/slave configurations using RS-485. The National Institute of Standards and Technology estimates more than four thousand BACnet installations in the United States.

* *OPC*
  OPC (OLE for Process Control) is a set of standard interfaces used in automation and enterprise systems implemented through the use of OPC servers on an Ethernet network. Originally based on Microsoft's OLE/COM technology, new specifications were proposed in 2004 with Microsoft's de-emphasis of COM in favor of cross-platform Web services.

* *Contact Closure*
  Typically used with legacy equipment, Contact Closure refers to a set of on/off hardware relays that provide information on the current state of a device. A hardware interface reads the state of the relays and passes this information to a monitoring system via other protocols—such as Modbus or SNMP.

IT

* *SNMP*
  SNMP (Simple Network Management Protocol) is the most widely used protocol within the IT domain. Data controls and alarms are defined in a Management Information Base (MIB)—a text file loaded into a Network Management System (NMS)—and transmitted across an Ethernet network.

* *XML*
  XML (Extensible Markup Language) passes structured data across different information systems over an Ethernet network. Primarily used for passing data between applications, it has gained momentum as a means to send and receive device data, as witnessed by the development of XML cards by certain manufacturers.

Limited Support for Devices

Another major limitation with traditional monitoring systems is the inadequate support for devices in the infrastructure. Most monitoring solutions were designed by the manufacturers of the hardware, with the primary intent of selling that particular manufacturer's hardware. Such a vendor-centric approach prevents customers from obtaining support for competitive products, locking organizations into a particular hardware vendor.

The gap between the IT and Facilities group also impacts the limited support available for devices. Traditional monitoring solutions have generally reinforced the separation between these two groups by designing solutions specifically for IT or Facilities, but not both. IT-based monitoring solutions concentrate on the IT hardware—such as servers, storage and rack-level power devices. On the other hand, Facilities-based monitoring solutions are heavily focused on the hardware managed by the Facilities group—such as UPSs, CRAC units and PDUs. This dichotomy of responsibilities and solutions limits the ability to create a cohesive support system to monitor all devices and infrastructure within the data center.

In addition, most monitoring solutions lack support for multiple protocols. If a solution only supports SNMP, for example, it simply is not capable of monitoring devices that only support Modbus. The lack of support for all devices in the infrastructure has forced organizations to implement multiple monitoring solutions for particular sets of devices. The intrinsic costs of purchasing hardware and software—yearly maintenance charges, training expenses and additional personnel to manage the various systems alone—often becomes a significant investment. Yet, another less obvious cost exists, as well. The consolidation of data from the multiple systems into actionable information for timely reporting and decision making is often extremely difficult, if not impossible—leaving the organization without the ability to comprehensively understand the relationship between the various infrastructure components. As a result, data center managers cannot make the intelligent decisions necessary to effectively manage the overall facility.

Fragmented View of the Data Center

Finally, traditional monitoring systems only provide a partial and often fragmented view of the data center. While in part due to the reliance of information housed in multiple systems, traditional monitoring systems also do not provide a comprehensive, holistic view of the data center, which is paramount in today's organizations. It is insufficient to only identify the load on a particular power distribution unit (PDU); management must also understand the business impact if that PDU is taken off-line; for example, what servers and applications are directly affected, as a result?

Managing the complex environment of today's data centers requires the ***Planning and Management*** of data center resources, the ***Operational Control*** of the People and Processes that run the data center and ***Monitoring and Automation***—the ongoing real-time interaction with the equipment housed in the data center facility. Each of these pillars of Data Center Service Management reinforces the others and magnifies the return on investment.

Data Center Service Management

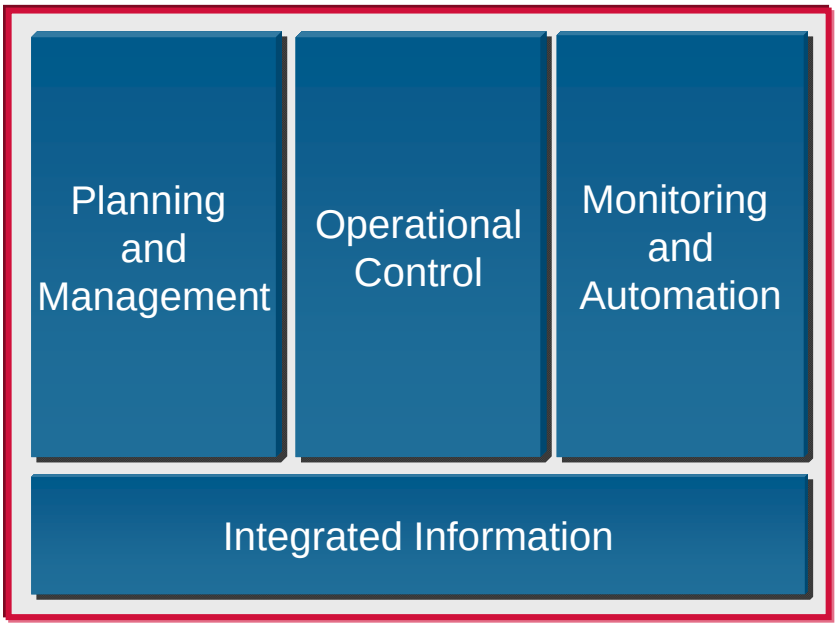

Source: Aperture Technologies, Inc., 2008

Not Just About Alerting Anymore

It is essential for the new generation of data center monitoring solutions to help break down the barriers between the IT and Facilities groups in order to foster collaboration. In addition, these solutions must directly serve the different groups of users with regards to their functions and multi-dimensional perspectives as follows:

1. **IT Operations**
   IT Operations is responsible for maintaining the equipment within the racks. As such, the ideal monitoring solution provides an IT view within the facility infrastructure on the device information within the racks, allows the group to efficiently provision equipment as necessary and generates capacity reports to ensure adequate resources for future provisioning is available. Although IT Operations does not “own” the equipment and does not directly make changes to the equipment, they rely on the equipment for services and require information to respond proactively and with agility to the needs of the business.
2. **Operations Center**
   The Operations Center manages the current state of the overall data center—identifying any potential risks and assessing how to respond to alerts. For this group, data center monitoring requires alarms and alerts to be correlated to the appropriate IT assets and in the context of the IT organization, such that they have the ability to take action. For example, if an issue arises with a particular uninterruptible power supply (UPS), the Operations Center cannot officially address the issue. However, if the impact on the servers

can be assessed, they can respond as if it were an IT issue. In addition, the solution must provide proactive alerting such that the Operations Center is aware of potential problems before they actually occur.

3. **Facilities**
   Facilities is the one group that actually owns the equipment within the data center. Having the ability to access all information regarding issues and the root cause to problems in an integrated, holistic system allows the Facilities group to respond efficiently.

The groups cannot utilize multiple systems with a variety of interfaces to effectively manage the data center. The ideal monitoring solution provides comprehensive, real-time data to integrate information from a variety of different systems into a single pane of glass and encourages holistic data center management processes. Data may be extracted in inconsistent manners, or even pulled from external systems, but monitoring provides accurate, real-time data and reporting to allow all audiences to proactively respond to critical issues—including system administrators, executives, line of business managers, partners and customers—which can save millions of dollars in potential lost revenue. For example, the Data Center Journal reports the cost of downtime continues to climb, with some industries approaching $3 million per hour in lost revenue due to issues related to downtime, as indicated in the following table[2]:

[2] *The Data Center Journal*, Volume 4 | July-August 2007.

Industry Sector	Lost Revenue Potential / Hour
Energy	$ 2,817,846
Telecommunications	$ 2,066,245
Manufacturing	$ 1,610,654
Financial Institutions	$ 1,495,134
Information Technology	$ 1,344,461
Insurance	$ 1,202,444
Retail	$ 1,107,274
Pharmaceuticals	$ 1,082,352
Banking	$ 996,802
Food/Beverage Processing	$ 804,192
Consumer Products	$ 785,719
Chemicals	$ 704,101
Transportation	$ 668,586
Utilities	$ 643,250
Health Care	$ 636,030
Metals / Natural Resources	$ 590,598
Professional Services	$ 532,510
Electronics	$ 477,386
Construction and Engineering	$ 389,601
Media	$ 340,432
Hospitality and Travel	$ 330,654
Average of All Industries	**$ 1,010,536**

Another critical component of the ideal monitoring solution is access to real-time data for capacity planning. The Aperture Research Institute recently concluded "almost 40 percent of [survey] respondents reported they had run out of space, power or cooling capacity without sufficient notice. This can result in delaying provisioning of new initiatives or paying more for additional resources to host those systems. The survey results are indicative of a burgeoning problem in data centers today, where no mechanism is in place to proactively manage and optimize resources.[3]" Implementing a robust monitoring solution can prevent such late notifications from occurring and the significant risks associated with the lack of infrastructure capacity.

The Green Factor

Energy consumption by data centers within the U.S. doubled between 2001 and 2006 and is expected to double again by 2011, unless energy efficiency in data centers becomes a priority. This factor alone has caused many organizations to investigate the benefits of a comprehensive monitoring solution to provide information to measure the performance of the data center in response to the escalating "Green IT" movement, enable cost reductions, increase operational

[3] "Organizations Struggle with Data Center Capacity Management," Aperture Research Institute, November 23, 2006, p. 3, http://www.aperture.com/about/research_institute/ari_cap_mgmt_112306.pdf.

efficiency and ultimately optimize utilization of data center resources. As stated by the Green Grid, a non-profit consortium of information technology companies and professionals, "Given that unlimited access to energy is unlikely and the cost of building or retrofitting data centers is high, it makes sense that customers are provided the tools to increase the efficiencies in today's data centers. Measuring and improving these facilities using the proper performance and energy metrics, along with a set of best practices, will help companies plan a more efficient data center.[4]"

Two fundamental metrics are currently proposed for use in measuring data center effectiveness and efficiency—the Power Usage Effectiveness (PUE) and its inverse, the Data Center Efficiency (DCE).

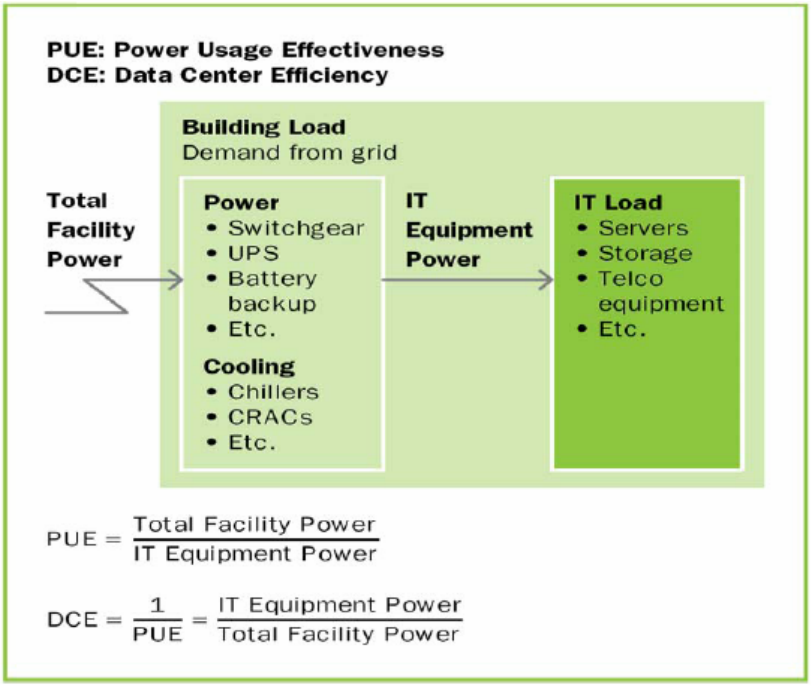

Source: The Green Grid, 2007.

[4] "Decreasing Data Center and Other IT Energy Usage Patterns," The Green Grid White Papers.

MONITORING: A CRUCIAL COMPONENT TO DATA CENTER MANAGEMENT

The following excerpts indicate how data center monitoring has become a crucial component to achieve operational efficiency:

- "Of all the techniques available to users, right-sizing the physical infrastructure system to the load has the most impact on physical infrastructure electrical consumption. Right-sizing has the potential to eliminate up to 50 percent of the electrical bill in real-world installations.[5]" For a typical 25,000 square feet data center, a 50 percent reduction in power can result in saving of approximately $1.3 million per year.
- "Based on a review of a range of incentives and voluntary programs that have been used in other sectors, and considering the unique aspects of the server and data center market, a number of recommendations can be made to pursue improved energy efficiency in the near term." One recommendation includes:
  - **"Standardized performance measurement for data centers**
    Data center operators need standard metrics to assess and report the energy performance of their facilities. The Federal government and industry should work together to develop an objective, credible energy performance rating system for data centers, initially addressing the infrastructure portion but extending, when possible, to include a companion metric for the productivity and work output of IT equipment.[6]"
- "The Green Grid recognizes the importance of establishing metrics for data center efficiency…these metrics and processes will help determine if the existing data center can be optimized before a new data center is needed.[7]" With a cost of a new Tier IV data center now exceeding $1,000 per square foot, optimizing your current data center to delay the building of a new data center can result in considerable savings.

Looking Through a Single Pane of Glass

The ability to sense and proactively respond to a data center event in real-time is critical to the effectiveness of the data center. Holistic data center management—including real-time monitoring—allows organizations to better understand the relationship between space, power and cooling resources and react quickly to changes in the environment.

Understanding the future of the data center is a priority for organizations to reach business goals. In order to accomplish this, a robust system with a view into past trends, a holistic snapshot of current resources and the ability to forecast future requirements must be implemented. Aperture VISTA Infrastructure Monitoring™ is the active component of Aperture's comprehensive data center

[5] "Guidelines for Energy-Efficient Data Centers," The Green Grid, February 16, 2007, p. 5, http://www.thegreengrid.org/gg_content/Green_Grid_Guidelines_WP.pdf.
[6] "Report to Congress on Server and Data Center Energy Efficiency Public Law 109-431," U.S. Environmental Protection Agency ENERGY STAR Program, August 2, 2007, p. 16.
[7] "Green Grid Metrics: Describing Data Center Power Efficiency," The Green Grid, February 20, 2007, p. 3, http://www.thegreengrid.org/gg_content/Green_Grid_Metrics_WP.pdf.

management system to help organizations create a central system to monitor and assess the vital signs of the data center. It provides data center management with valuable, real-time information to improve efficiency, effectively manage capacity and right-size the data center infrastructure to meet business demands.

VISTA Infrastructure Monitoring is a Web-based, vendor-neutral solution to monitor the status and impact of the critical infrastructure through a “single pane of glass.” It delivers integrated information, including accurate data, completeness of information and a multi-dimensional perspective such that executives can leverage actionable information, create agility, deliver high quality of service and easily manage operational costs in an efficient and effective manner.

The tools and best practice processes inherent within VISTA Infrastructure Monitoring provide an ideal level of monitoring and control of the data center. With VISTA Infrastructure Monitoring, organizations are able to:

- Balance supply of infrastructure resources to meet business demands
- Track critical infrastructure usage and metrics over time
- Forecast future application pipelines to predict the consumption of data center resources
- Determine when data center resources will run out across all data centers
- Provide lead-time for future data center planning
- Understand future trends based on past use and report on key capacity metrics
- Minimize the risks associated with outages or service disruption
- Bridge the gap between IT and Facilities

Conclusion

Many organizations are faced with the task of meeting operational demands in a facility that pushes capacity to the limit. It is essential for organizations to leverage data center monitoring to properly plan and allocate resources, as well as provide a sufficient infrastructure while keeping costs low. Many companies currently rely on multiple systems with unique methods of extracting and storing information, which lacks the ability to make appropriate comparisons. What is needed is a system that provides an IT view into the data center facility to obtain an accurate, real-time picture of the current state of the critical infrastructure and provides the ability to forecast future needs.

Aperture VISTA Infrastructure Monitoring is the ideal holistic data center monitoring solution for data center management to monitor the status and impact of the critical infrastructure through a “single pane of glass”. With VISTA Infrastructure Monitoring, organizations can benefit from an increase in data accuracy, improved capacity planning forecasts and greater availability. Through the use of proactive monitoring, organizations can identify potential problems before they occur.

This white paper was sponsored by Aperture Technologies, Inc.

About Aperture

Aperture is the leading global provider of software for managing the physical infrastructure of data centers. Aperture's solutions reduce operational risk and improve efficiency through the planning and management of data center resources. Aperture delivers the best practice processes that enable organizations to take control of an increasingly complex physical infrastructure including equipment, space, power, cooling, network and storage. With over 20 years of experience, Aperture provides organizations with the information required to optimize their data center operations, delivering better services at the lowest cost. Aperture's customers include the world's largest companies, half of which are Fortune 1000 and Global 500 organizations.

Aperture®

Changing the way companies run their data centers

www.aperture.com

CORPORATE HEADQUARTERS

Aperture Technologies
9 Riverbend Drive South
Stamford, CT 06907

tel 203.357.0800
800.346.6828
fax 203.357.0809

EUROPEAN HEADQUARTERS

Aperture Technologies
288 Bishopsgate
London EC2M 4QP
United Kingdom

tel +44 (0)20 7959 3024
fax +44 (0)20 7959 3030

02-08