Data Center Management

Searching for the Optimum Solution

Management of assets in the data center is incomplete, inaccurate, and difficult to maintain with today's technology

No Limits Software White Paper #1
By David Cole

Table of Contents

Overview

A data center can contain thousands of assets, from servers, storage and network devices to infrastructure support equipment such as Computer Room Air Conditioners, PDUs and UPSs. Keeping track of these assets is an ongoing task faced by data center managers around the world. A Digital Realty Trust survey found that only 26% of data center managers could locate a server that had gone down within minutes. Only 58% could locate the server within 4 hours and 20% required more than a day. The inability to locate equipment in the data center increases the mean time to repair (MTTR) for the equipment and decreases the overall availability.

20% of data center managers can take more than a day to find a server!

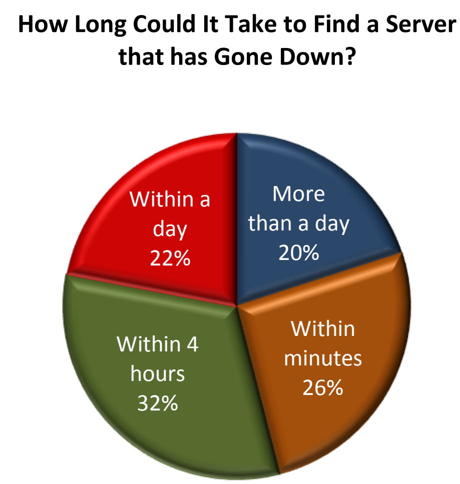

Digital Realty Trust

The ability to manage the IT assets is only part of the solution, however. There are three primary resources which must be considered when trying to determine where to place the next rack asset:

- Space – is there enough contiguous space to house the asset?
- Power – is there sufficient redundant power for the asset?
- Cooling – is there enough cooling to remove the heat generated by the asset?

Comprehensive management of IT equipment at the rack level must consider space, power and cooling. This white paper addresses the state of the systems and related technology currently being used to address rack level management and provides criteria to use in the search for an optimal solution.

Asset Management Software

There are a number of asset management software solutions on the market today. These solutions are intended to address management of the IT assets in the data center, with the primary focus on data center planning. Their purpose is to answer questions such as:

- How much capacity do I have remaining in my data center?
- Where is a particular server located?
- What is the configuration of a particular server?
- What equipment is installed in a particular rack?

Accuracy of Manually Entered Asset Information

Organizations can typically expect a 10% error rate in manual data entry.

There are several drawbacks to the current asset management systems. First, they typically require manual entry of the asset information, an arduous, time-consuming process that is rife with errors. In the Computer Associates technology brief *Striving to Achieve 100% Data Accuracy: The Challenge for Next Generation Asset Management* (Watson & Fulton, 2009), the authors point out the difficulty in maintaining the accuracy of this information. The authors point out that "Manual tracking with pen and clipboard, or even spreadsheets is time consuming and highly error-prone. Organizations can typically expect a 10% error rate in manual data entry due to typing and transcribing errors."

In a survey of the International Association of IT Asset Managers (IAITAM) members, respondents said that 85% accuracy rate for tracking IT assets was above average and that a 90-95% rate was exceptional. Consider for a moment the impact of a 10% error rate on the accuracy of manually-entered asset data. In a data center with 5,000 servers, as many as 500 of the servers will have inaccurate data recorded.

When asked in an Aperture Research Institute survey (ARI, 2009) about the difficulties in gathering asset data, more than 70% of the respondents identified human error as a problem. Another 58% said there were problems in correctly identifying the equipment and 40% responded that there were problems in correctly identifying the equipment location.

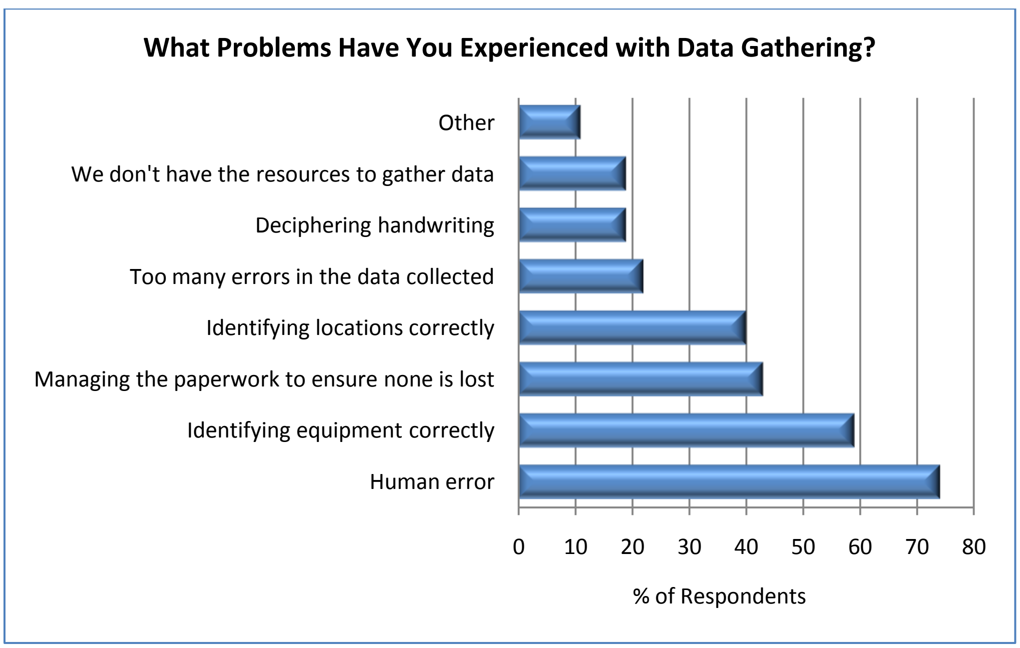

Asset management systems are intended to provide data center management with the information they need to make important decisions about capacity planning, proper placement of equipment in the rack, and so on. If the information is inaccurate, the decisions made based on the information become suspect.

Cost to Manually Collect Asset Information

Another drawback to manually entering data into an asset management system is the time and cost to perform the initial data collection. At a typical cost of $13-15 per asset to manually collect visible data (name, manufacturer, model, serial number and location), the initial data collection for a fully populated rack could cost more than $500. The collection of more detailed hardware information such as processor, memory, disk and network configuration or other information such as virtual machines and installed software would add considerably to this cost.

Another issue with manual data entry is the inability to audit the data without performing yet another manual data collection. While the cost to audit the data may be less than the original data collection, it could still cost hundreds of dollars per rack. With regulations such as Sarbanes-Oxley, HIPA and CFR-11, the need for more frequent audits can quickly increase these costs.

The optimum solution would provide a means to quickly and accurately collect IT asset information, maintain 100% accuracy of the data, and provide an inexpensive and accurate means of auditing the data.

Tracking Authorized and Unauthorized Changes

80% of system outages are due to changes.

80% of MTTR is trying to find out what changed.

Asset management systems generally deal with static information, in other words data that doesn't change unless someone manually changes it. In an ideal world, everyone would follow the specified work flow procedures when changes are to be made to an IT or facilities asset. Unfortunately, this is not always the case. Changes are frequently made without the proper authorization and without following the proper process. An asset management system will be unaware of these changes because they were never entered into the system. In the book The Visible Ops Handbook: Implementing ITIL in 4 Practical and Auditable Steps (Behr, Kim, & Spafford, 2004-2005), the authors examined a number of high performing IT organizations and found that "80% of their outages were due to a change, and that 80% of their MTTR was trying to find out what changed". These organizations learned to first look at "the scheduled and authorized changes for the affected asset, as well as the actual detected changes on the asset. By just looking at this information, problem managers could recommend a fix to the problem over 80% of the time, with a first fix rate of over 90%." The authors found that organizations which implemented automated change auditing were "surprised and alarmed to see how many changes are being made 'under the radar'." The authors recommend scanning systems for changes at least once a day. The ability to track both authorized changes and detected changes – changes made but not necessarily authorized – is key functionality needed to reduce MTTR and increase overall system availability.

The optimum solution would provide a means to automatically detect and track all system changes – authorized or not – and provide a means for problem managers to quickly view these changes as a means to recommend problem resolution.

Accuracy of Asset Power Usage

Comprehensive management of IT equipment at the rack level must consider space, power and cooling. When it comes to power, IT asset management systems have traditionally used either nameplate values – the manufacturer's value to indicate the equipment's maximum power draw – or an arbitrary percentage of the nameplate value (known as the de-rated, budgeted or design value) to measure the asset's power usage. This information is then used by data center managers for capacity planning purposes and to determine optimum placement of new equipment in a rack.

Inaccurate rack power data can result in stranded power or system downtime.

If the power data is inaccurate, there are several undesirable possible results. If the rack power usage is estimated on the high side, there is the potential for stranded power. Under this scenario the data center manager opts to not add any additional equipment to the rack because the rack appears to be at its capacity for power. In actuality, there may be plenty of available power, resulting in a waste of valuable rack space, PDU breaker capability and overall PDU power. The second possible result, estimating rack power usage on the low side, has the even more serious potential consequence of system downtime. Under this scenario the data center manager, believing there to be adequate available power, continues to add equipment to the rack. When the power usage exceeds the available capacity, breakers trip, servers go down, and system availability is impacted.

In the article *The Whys and Hows of Measuring Power in your Data Center* (Chan & More, 2009), the authors cite a study of server power usage at Raritan's own data center. The study illustrates the potential risks of using inaccurate power usage estimates.

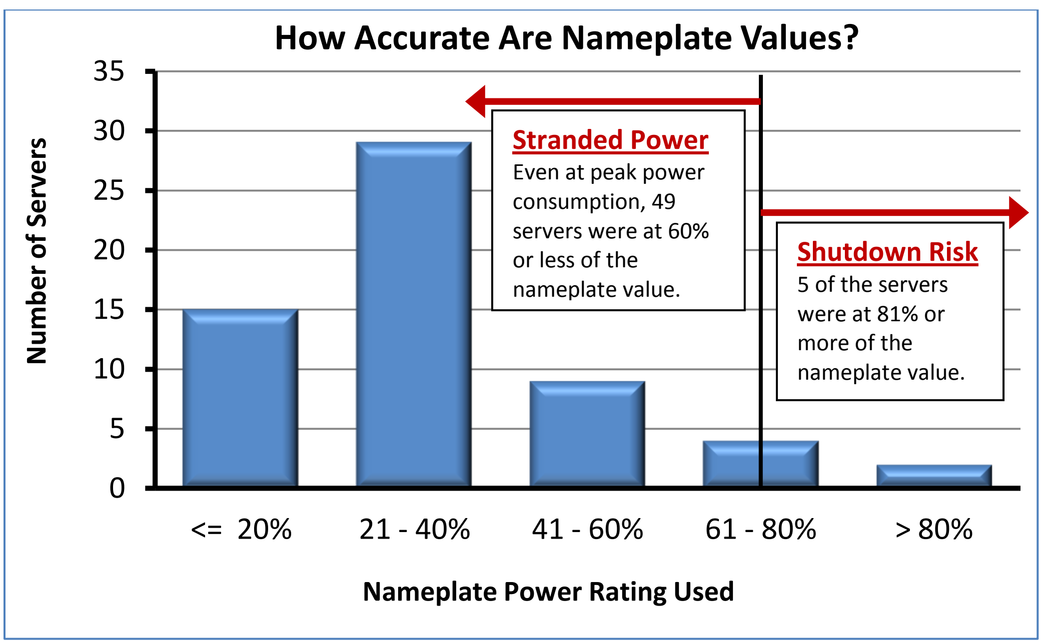

Raritan found that, even at peak power consumption, 49 of the 59 servers used 60% or less of the nameplate value, with 15 using less than 20%. They also found that 5 of the servers were using 81% or more of the nameplate value. Using a de-rated value of 70%, a typical industry value, the study showed that 83% of the servers would report high power usage values (stranded power) while 8% would report low power usage (shutdown risk).

This study, as well as numerous others in the industry, point out the potential problems with the use of nameplate or de-rated values alone for capacity planning or server placement purposes. As the authors of the Raritan white paper "Center Power Distribution and Capacity Planning: Understanding what you know – and don't know – about power usage in your data center" (Raritan & PTS, 2009) point out:

Only through measurement can managers accurately know the numbers they need for energy planning.

> *As energy issues come under more scrutiny and as tools become available for accurate measurement, IT administrators and facilities managers should no longer rely on the published nameplate power ratings on their units and factor in accepted industry assumptions. While adequate in the past, "close enough" isn't "good enough," given the potential power crisis that's just over the horizon.*
>
> *Only through individual server measurement can managers accurately know what power their equipment is drawing and acquire precise numbers that will aid their energy efficiency planning efforts.*

It is important to understand all aspects of power usage for a device, including maximum power (nameplate), design state (de-rated) and measured values. Measured values must include steady-state (average), peak and minimum values.

The optimum solution would provide a means to use actual IT equipment power usage values (average, maximum and minimum) for capacity planning or server placement purposes.

Importance of Rack Environmental Conditions

Cooling is a valuable data center resource which must also be considered when managing at the rack level, especially as rack densities continue to rise. The ability to monitor environmental conditions at the rack level becomes an important tool In order to ensure that the inlet temperature to the IT equipment stays in an acceptable range.

As illustrated in the figure below, the occurrence of a cooling failure and the resulting temperature rise in a high density rack can cause a thermal shutdown of the servers in as little as two minutes.

Cooling failure and resultant rack temperature rise can cause server shutdown in less than 2 minutes!

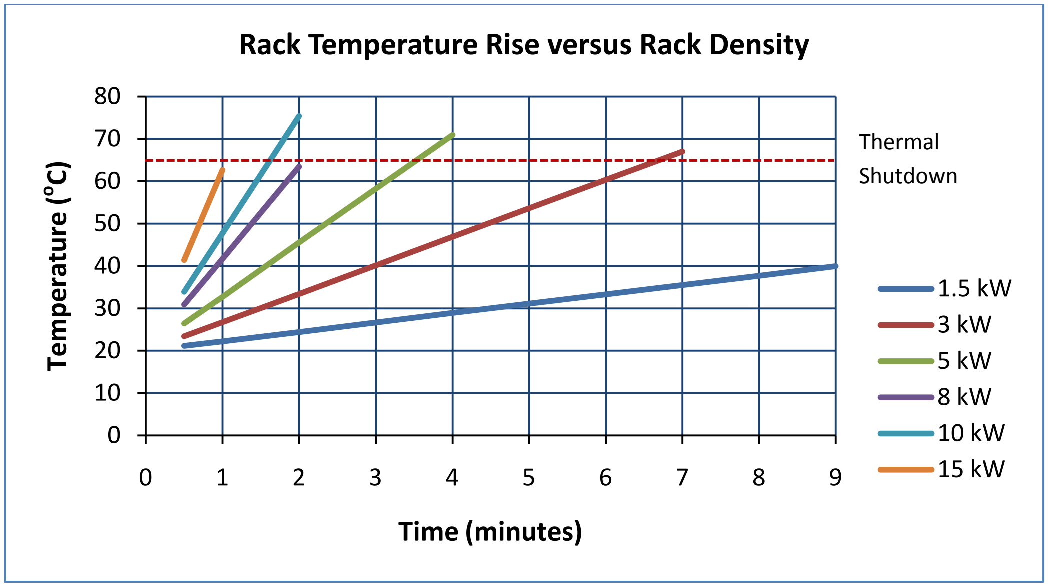

The optimum solution would provide a means to monitor rack environmental conditions for capacity planning purposes and to reduce the risk of shutdown.

Real-Time Monitoring

In order to properly address power and cooling as part of an overall data center management solution, it is important for asset management systems to have access to real-time data. Without accurate energy usage measurements, it is very difficult to make informed decisions regarding capacity and energy efficiency. While energy usage measurement is increasing, there are still many data centers which don't have measurement tools in place. The Green Grid white paper *Energy Measurement Survey Results Analysis* (Bednar, Winiecki, & Winkler, 2009) shows the breakdown of how the respondents measured energy usage.

Only about 1/3 of IT equipment (server, storage and network) is monitored using dedicated meters.

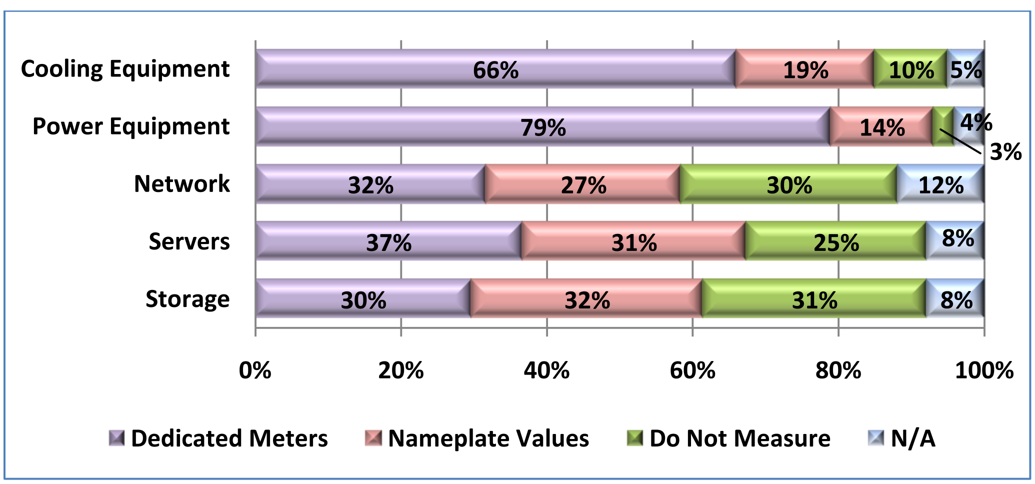

Real-time monitoring systems typically focus on either IT or facilities, but very rarely on both. Monitoring of facilities equipment is usually accomplished using dedicated meters. The meters may be a separate component or integrated into the equipment. Typically, these systems use a variety of communication protocols, including SNMP and Modbus.

There are two primary methods for monitoring the power usage of IT equipment. The first method is to use rack power distribution units (PDU) with outlet-level metering capability. The typical communication protocol is SNMP. There are two drawbacks to this method. First, rack PDUs with outlet level monitoring are more expensive than a rack PDU that monitors only the power draw of the entire PDU. Second, dual-corded devices split their power between two different rack PDUs. A management system must be used to monitor both rack PDUs in order to determine the entire power draw of the device.

The second method for monitoring IT equipment is to poll power readings directly from the device. The typical communication protocol is Intelligent Platform Management Interface (IPMI). There are three potential issues with using this method. First, not all IT equipment supports IPMI. Second, equipment manufacturers may use their own extensions to IPMI, or even their own protocol, to provide the device information. These protocols include Dell Remote Access Controller (DRAC), HP Integrated Lights Out (iLO), IBM Remote Supervisor Adapter (RSA), and Sun Integrated Lights Out Manager (ILOM). Additional information may be obtained using protocols such as SNMP, WMI, and WBEM. This requires a management solution which supports multiple protocols to retrieve the same information from devices from different manufacturers. Finally, while a device may support one of these protocols, it may not provide power usage information.

Because of the issues with polling power usage information from specific devices, it is more typical to see monitoring of power usage at the rack level rather than at the device level. While it doesn't provide device-level detail, monitoring power at the rack level is very useful for capacity planning purposes as it can provide information on whether or not additional equipment can be placed into a rack. Monitoring power usage at the rack level can be done using either metered rack PDUs or using branch circuit monitoring at the PDU.

The optimum solution would monitor both facilities and IT equipment at either the device or rack level. At the device level, the solution would support multiple protocols to retrieve information from the device.

Server Monitoring to Improve Energy Efficiency

Estimates put the number of ghost servers at 8-10% although the actual number may exceed 30% in some data centers.

Sun estimates that 8-10% of servers are "ghost servers", unproductive servers which are running no applications but still consuming valuable space, power and cooling. Other studies show that the actual percentage of ghost servers may exceed 30% in some data centers. Mark Monroe, the director of sustainable computing at Sun, says "the cost of running a server for three years exceeds its original acquisition cost, so keeping the ghost servers around has an easily measured effect in energy savings. Identifying and eliminating wasted resources are key components of green or eco-computing" (Dunn, 2007). The cost to run a server for a year can be $800 or more. If 10% of the servers in a data center with 1,000 servers were ghost servers, the cost of these unproductive systems could exceed $80,000 per year.

In order to determine if a server is actually performing work, it is important to monitor parameters such as CPU utilization, I/O and network traffic. The CPU utilization values are also important in identifying servers which are good candidates for virtualization.

Another important component of reducing data center energy usage is the use of server power management features. Newer multi-core processors feature demand-based switching (DBS) to match power usage to computational demands. The following table (Source: Intel, 2008) illustrates the benefits of using DBS (Minas & Ellison, 2009):

	CPU Utilization		
	15%	**30%**	**45%**
System Power	258 W	291 W	316 W
System Power with DBS	201 W	220 W	240 W
Power Savings	57 W	71 W	76 W
Cost Savings	$148/year	$186/year	$200/year
Savings for 500 Servers	$74,000/year	$93,000/year	$100,000/year

Source: Intel, 2008

Server power management can save $100,000 per year for 500 servers!

Unfortunately, many of the servers ship with the power management features disabled by default. Most data center managers don't know if their servers have the power management features enabled because they don't have a tool that will readily provide this information.

Without power usage metrics, data center managers are forced to make capacity planning decisions based on nameplate or de-rated values, numbers that may be dramatically different than the actual power usage of the equipment. They will be unable to make informed decisions on which servers to purchase. Without CPU, I/O and network utilization metrics, data center managers do not know which servers are underutilized or perhaps not being used at all, resulting in significant costs to provide space, power and cooling for servers that are not being used efficiently. Without information on which servers are running power management, data center managers will miss out on significant opportunities to reduce server power usage.

When surveyed on what actions they would likely take if they had more information about their data center energy efficiency, such as CPU utilization and server power management configuration, the Green Grid respondents replied as follows (Bednar, Winiecki, & Winkler, 2009):

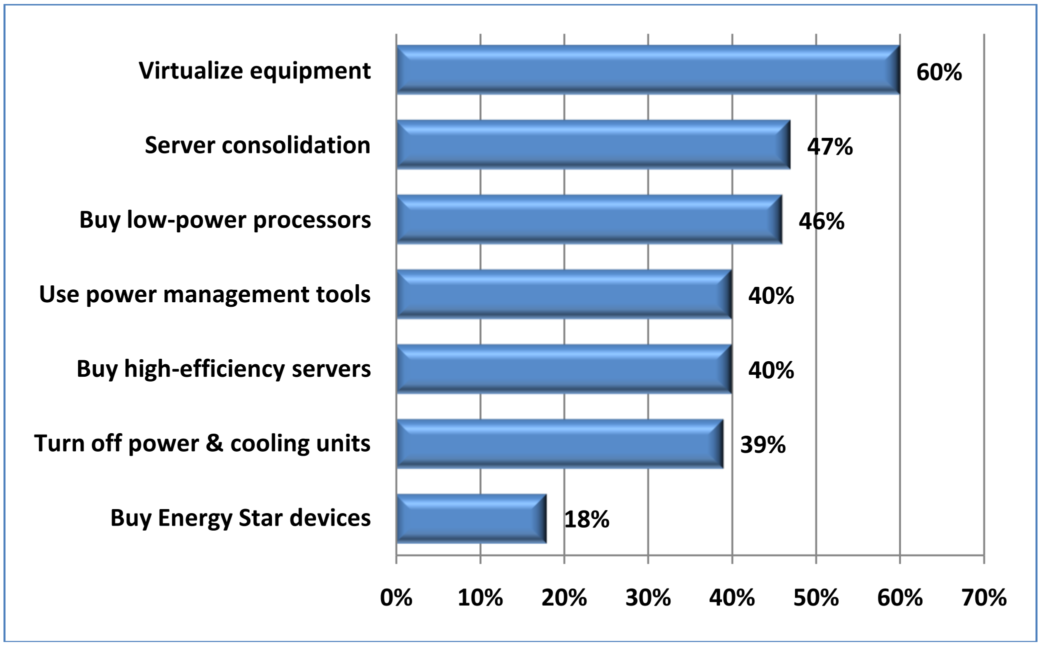

Server and storage utilization information allowed PTS to reduce IT power usage by 24%!

In the PTS Data Center Solutions white paper "Impact of IT Upgrades on Energy Usage and Operational Cost" (Sacco, Davis, Niessen, & Soundararaj, 2010), PTS used the information from server measurements to upgrade their IT infrastructure. The upgrades provided the following outcomes:

- Average CPU utilization was increased from 2.1% to 9.3%
- Peak CPU utilization was increased from 5.5% to 21.2%
- Server virtualization was increased from 14% to 64%
- System storage utilization was increased from 28% to 53%
- Data storage utilization was increased from 27% to 48%
- IT power consumption was reduced by 24%
- Power consumption due to cooling and power protection/distribution was reduced by 29%

The optimum solution would provide monitoring of server parameters such as CPU utilization, I/O, network utilization and server power management configuration in order to make informed decisions on capacity and energy efficiency planning.

RFID Solutions

RFID provides advantages over barcodes, including more storage, ability to reprogram and remote reads.

Radio-frequency identification (RFID) is becoming more prevalent as a means of tracking IT assets in the data center. There are some obvious advantages to this technology. The primary advantage is the ability to track an asset throughout its lifecycle, from the loading dock to staging to the rack to decommissioning and disposal. As the asset is moved, it is tracked using RFID readers placed throughout the building, particularly in room entrances.

There are several major advantages of an RFID tag over a barcode. First, an RFID tag can contain much more information than a barcode, which is typically limited to 10-12 numbers or letters. Second, an RFID tag is reprogrammable while a barcode must be replaced if the information changes. Finally, an RFID tag can be read remotely while a barcode requires someone to manually scan it with a handheld device.

There are two types of RFID tags, passive and active. The following chart shows the primary differences between the two types of tags:

	Active	Passive
Signal Strength	Stronger	Weaker
Signal Availability	Always on	Responds when read
Size	Larger	Smaller
Initial Cost	Higher	Lower
Maintenance	Replace every 2-5 years	Indefinite lifetime

RFID tags are read using RFID readers which are distributed in areas in which tags are to be read. Typically, these readers are placed in doorways, enabling the readers and associated management system to determine the zone in which the tagged asset resides. As the tagged asset is moved through a doorway, the reader will detect its presence and record where the asset is currently located.

There are some drawbacks to the use of RFID in a data center. First, each tag must be programmed and installed on each asset to be tracked. Each tag is programmed with a unique identifier that is used to tie the identifier to the asset in the asset management database. This programming requires someone to gather the information about the asset and to then program this information into the tag. As discussed earlier, this manual entry may be inaccurate, resulting in incorrect data being stored on the tag. The information the tag sends to the reader is only as good as the information that was put on the tag.

Unless the tag is programmed by the vendor, there is a cost to tagging an asset. The information must be gathered and programmed into the tag and then the tag must be attached to the asset. There is also the cost of the tag themselves, which can be as high as $50 or more for an active tag. There is an added cost of maintenance for the active tags, as well, as these tags must have the battery replaced every two to five years.

One other drawback of an RFID solution is that an RFID tag, much like a bar code, contains static data. This means the RFID solution must be integrated with other systems to track asset configuration changes and power and environmental sensors.

There are many factors that affect the RFID tag read rate, including the distance from the tag to the reader and the substance on which the tag is mounted. Metal and water are not very good conductors of ultra-high frequency signals, requiring special tags for data center equipment such as servers. The tag orientation and design can also affect how well a tag is read. It is important to know that it is possible that a tag will not be read every time.

Deployment of RFID systems in data centers can be challenging even for the most experienced professional. The entire environment must be analyzed, typically requiring a wireless site survey in order to identify all possible sources of interference and signal attenuation.

RFID is designed as a proximity or portal system, not a location system.

One thing that is important to understand is that RFID was designed as a proximity or portal system, not a location system. In other words, the systems were designed to know when a tag passes a reader or passes through a portal containing a reader. This is a perfect fit for supply chain management or toll road applications. To determine an asset location, RFID must rely on triangulation of multiple readers which measure the return signal strength of a tag. This is a difficult task under any circumstances, but is particularly difficult in a data center environment with hundreds of racks and thousands of servers creating a noise-filled environment. Add to this the need to track the location of a server to the accuracy of 1.75 inches (the height of rack unit) and an RFID-only solution may not be possible without the use of additional equipment installed in each rack.

There are several methods that have been used to track assets to the rack unit using RFID and additional equipment in the rack. One method uses Near Field Communication, a short-range high frequency wireless communication technology which enables the exchange of data between devices over a very short distance (typically 4 inches or less). This method requires an RFID reader for each rack unit. A typical 42U rack would require 42 RFID readers, for example. There are a couple of drawbacks to this method. First, the number of readers required may make the system cost-prohibitive. Second, because the readers have a very short range (as low as 20 mm) there are restrictions on where the tags can be placed on the asset. Proper tag mounting within the acceptable range may be difficult due to the presence of ports, power connectors and other physical limitations. Finally, there is no way to track zero-U equipment (equipment such as a rack PDU that is not placed in a rack unit).

The second method to track assets to the rack unit uses a combination of RFID and infrared technologies. RFID is used to track asset movement through the building using readers at portals or in data center zones while infrared technology is used to determine the rack unit. Again, there are drawbacks to this method. First, there is equipment which must be installed in each rack. Second, there is no way to track zero-U equipment (equipment such as a rack PDU that is not placed in a rack unit).

RFID use in the data center has value, but it must be used correctly to be cost-effective.

RFID has value in the data center, but it must be used correctly.

- **Use passive tags instead of active tags** – the two biggest advantages of the passive tags are much lower initial cost and the lack of maintenance
- **Have the manufacturer program and attach the tags** – this saves the cost of the programming and allows immediate tracking when the device arrives on the dock
- **Use RFID to track room location** – using RFID to track assets through portals requires fewer readers and results in a less expensive system which is easier to design and implement
- **Use other technology to track to the rack unit** – don't try to force RFID to do something it was not intended to do

The optimum solution would use RFID for tracking assets at a room level and other technology designed to track the asset the "final mile".

Summary

Comprehensive management of IT equipment at the rack level must consider space, power and cooling. The optimum solution would provide the following functionality:

- Quickly and accurately collect IT asset information
- Maintain 100% accuracy of the data
- Inexpensive and accurate means of auditing the data
- Automatically track all authorized and detected system changes
- Ability to quickly view system changes to recommend problem resolution
- Monitor rack environmental conditions to reduce shutdown risk
- Monitor both facilities and IT equipment at either the device or rack level
- Provide monitoring of server parameters such as CPU utilization, I/O, network utilization and server power management configuration
- Use RFID for asset tracking at the room level and other technology designed to track the asset to the rack unit

About No Limits Software

No Limits Software was founded in 2009 by industry experts in the field of data center monitoring and management solutions. For more information, contact info@nolimitssoftware.com or visit www.nolimitssoftware.com.

Bibliography

ARI. (2009, November 10). *Organizations Overlooking Significant Opportunities to Improve Efficiency and Reduce Expenses, Despite Adoption of Green Initiatives.* Retrieved from http://www.aperture.com/about/research_institute/ARI-ResearchNote-GreenInitiatives-OverlookingOpportunitiestoReduceCosts_11_10_09.pdf

Bednar, R., Winiecki, J., & Winkler, K. (2009). *Energy Measurement Survey Results Analysis.* Retrieved from the Green Grid: http://www.thegreengrid.org/~/media/WhitePapers/White Paper 26 - Energy Membership Survey_111709.ashx?lang=en

Behr, K., Kim, G., & Spafford, G. (2004-2005). *The Visible Ops Handbook: Implementing ITIL in 4 Practical and Auditable Steps.* IT Process Institute (ITPI).

Chan, H., & More, G. (2009, April 14). *The Whys and Hows of Measuring Power in your Data Center.* Retrieved from http://www.greenbiz.com/news/2009/04/14/whys-and-hows-measuring-power-your-data-center

Dunn, D. (2007, September 26). *The horror: 'Ghost servers' that haunt your bottom line.* Retrieved from http://www.computerworld.com/s/article/9038778/The_horror_Ghost_servers_that_haunt_your_bottom_line?pageNumber=1

Minas, L., & Ellison, B. (2009). *Energy Efficiency for Information Technology.* Intel Press.

Raritan, & PTS. (2009). *Center Power Distribution and Capacity Planning: Understanding what you know – and don't know – about power usage in your data center.*

Sacco, P., Davis, L., Niessen, T., & Soundararaj, S. (2010). *Impact of IT Upgrades on Energy Usage and Operational Cost.*

Watson, T., & Fulton, J. (2009, March). *Striving to Achieve 100% Data Accuracy: The Challenge for Next Generation Asset Management.* Retrieved from http://www.ca.com/files/TechnologyBriefs/33603-itam-tech-brief-qxd-final_201662.pdf