How Monitoring Systems Reduce Human Error in Distributed Server Rooms and Remote Wiring Closets

White Paper 103

Revision 0

by Dennis Bouley

> Executive summary

Surprise incidences of downtime in server rooms and remote wiring closets lead to sleepless nights for many IT managers. Most can recount horror stories about how bad luck, human error, or just simple incompetence brought their server rooms down. This paper analyzes several of these incidents and makes recommendations for how a basic monitoring system can help reduce the occurrence of these unanticipated events.

Contents

Click on a section to jump to it

Introduction	2
Simple or complex?	2
The nature of human error related downtime	4
Tales of the unexpected	5
Monitoring system components	5
More battle tales	10
Conclusion	11
Resources	12

APC white papers are now part of the Schneider Electric white paper library produced by Schneider Electric's **Data Center Science Center**
DCSC@Schneider-Electric.com

Introduction

Many IT managers recount stories of unexpected downtime events that occurred in their distributed server rooms and remote wiring closets. When analyzing these stories, a common thread emerges: lack of information. This lack of information leads to human error which causes the downtime. Stress levels are high because operators and administrators have no real-time data at their disposal and therefore, cannot prevent the human error from occurring.

Consider the following two statistics:

- According to estimates, there are 2.9 million server rooms and wiring closets in the United States alone[1]
- Over 70% of reported data center outages are directly attributed to human error[2]

This paper details common distributed server room and remote wiring closet downtime incidents. Then recommendations are proposed for how monitoring & automation software integrated with video surveillance and sensors can reduce the occurrence of human error-related downtime problems in these small, distributed environments (see **Figure 1**).

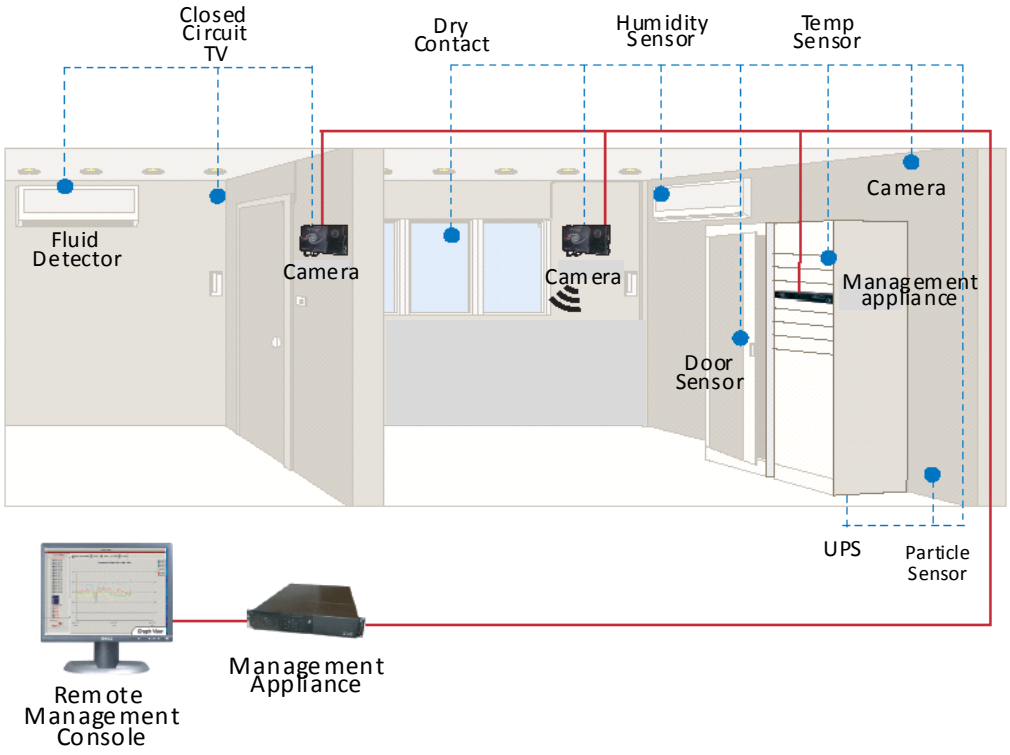

Figure 1

Human error reduction can be accomplished through intelligent monitoring at multiple levels

Simple or complex?

Two issues present themselves whenever monitoring systems are proposed for small, remote computing environments like wiring closets and server rooms. The first issue involves installation. How complex is it to install a monitoring system? That is, how much time does it take to capture information about the characteristics of the devices to be monitored and how much time is spent performing data entry of this information (consider hundreds of devices in a multiple site scenario)? How does the system know which devices are cameras, air

[1] IDC, *Building, Planning, and Operating the Next-Generation Data Center,* Michelle Bailey, 2008

[2] Uptime Institute, *Data Center Site Infrastructure Tier Standard: Operational Sustainability,* 2010

conditioners, UPS, heat sensors etc, and how are IP addresses generated so that the devices can communicate back their status? The second issue addresses the unknown amount of work involved whenever changes are made to power, cooling, and environmental monitoring equipment residing in the remote site. For example, how is a new firmware upgrade performed and how is a temperature threshold modified?

Over the last several years, monitoring software packages have evolved to the point where users can choose either to self install or to enlist an outside service for installation assistance. The outside service can usually have the user up and running within 1-2 days.

Monitoring software packages can be delivered either as distributable/downloadable code or as a rack-mounted server with preloaded software. The system can be installed either remotely or in a central data center (if, for example, dozens or hundreds of wiring closets need to be managed). Once the management server is plugged in, the client can be downloaded onto a laptop so that the operator can begin the process of identifying the power, cooling and environmental equipment, and the human activity that needs to be monitored. Most modern UPSs, cooling systems and security cameras come standard with the network interface cards (NIC) that are required for communication. The operator defines an IP address or range of IP addresses to be utilized by the devices to be monitored. An example of this is shown in **Figure 2**. Some systems can then automatically search the network and locate all of the power, cooling, and security devices to be monitored. This ability to "auto-discover" devices, greatly simplifies the challenge of system installation and start up. Once "discovered" the system begins to monitor the remote devices.

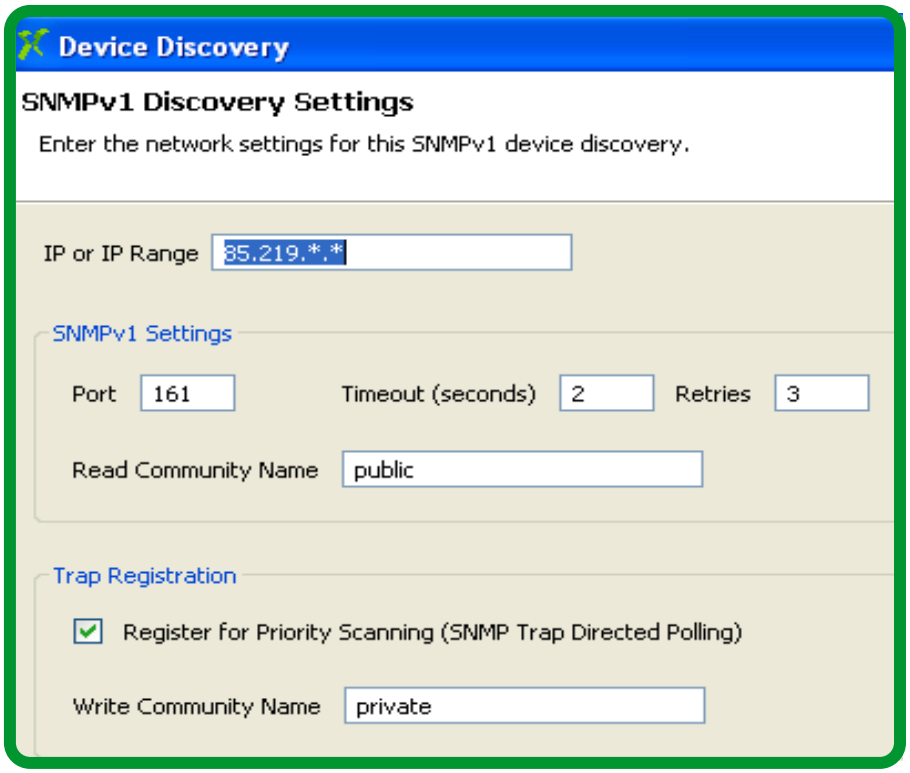

Figure 2

Setting up of IP addresses for multiple devices can be as easy as typing in a range of numbers (sample screen extracted from the Schneider Electric InfraStruxure Central application)

Some monitoring & automation systems also allow for devices to be grouped either by location, by row within a location, or by device type (e.g., grouping all cooling devices, all PDUs, all meters, all cameras etc.). This grouping exercise allows the user to set up policies and thresholds for that group. Common threshold parameters could include temperature, humidity, and designation of an open or closed status (e.g., doors of racks).

The thresholds, when exceeded, should trigger an alarm that is communicated to the systems administrator via email or text message. Care must be taken that only major changes to the remote environment trigger an alarm. If not, the administrator would be faced with the prospect of multiple alarms several times an hour. In this case, the administrator could become "numb" to the alarms and ignore them. Therefore a delicate balance needs to be maintained so that any alarm that comes through to the systems administrator is deemed as meaningful or important.

Updates to the server room or wiring closet, such as an update to firmware, are also simplified when a modern day monitoring system is in place. No longer does the data center manager have to send personnel out to remote locations to install firmware upgrades. Many monitoring systems are capable of performing mass configuration, which allows changes to be sent out over the network from the central location.

The nature of human error - related downtime

Distributed server rooms and remote wiring closets do not garner the same amount of investment and attention as large mission critical data center sites. The large, central sites are staffed with experts and are often equipped with the latest in security technology and an abundance of built-in redundancies. Distributed server rooms and remote wiring closets on the other hand, are staffed by individuals with multiple responsibilities, one of which might be to keep an eye on the wiring closet or server room. These spaces often have few security measures in place and are subject to more unwelcome outages than the larger, more sophisticated spaces. No matter how well a server room or wiring closet is planned, the risk for unanticipated downtime is always present. Some IT managers think they have every angle covered. They are proud of their server room design. Then along comes an innocuous-looking, uninformed technician or custodian that defeats the entire plan in less than five seconds.

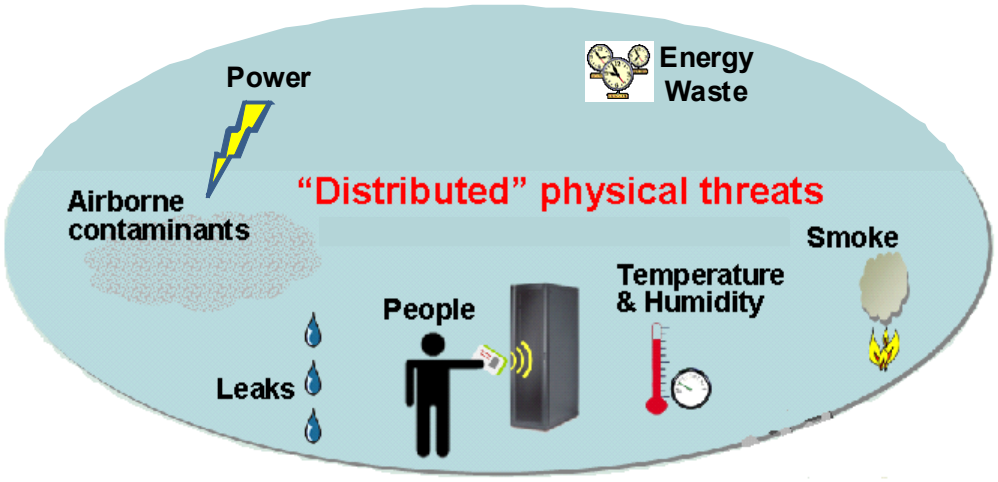

Figure 3

The phrase "an accident waiting to happen" is apt for small and remote server rooms

The list of incidences described below illustrates how the lack of a simple monitoring & automation system can lead to downtime in wiring closets and server rooms. In these environments, either no one is on site, or the person who is on site is not always capable of alerting systems administrators to a failure. An hour delay in the discovery of a cooling failure can make the difference in avoiding a complete outage. Fast, real time alerts enable administrators to supervise a switch over which can avoid loss of service.

Tales of the unexpected

Consider the following summary of human error-related events:

- A systems administrator responsible for a remote branch office server room went in to find out why the servers in the room went down. He discovered that the remodeling contractors during renovations had wrapped racks with shrink wrap plastic to keep dust away from the servers. The contractors neglected to inform IT that they would be doing this, so all the servers were on when they wrapped them up. The servers overheated and shut themselves down.
- A senior business manager decided to take matters into his own hands when he was having trouble accessing the Internet. He went into the server room, took cabling from the router and connected his laptop directly to the Internet thus bypassing all firewall services and encryption and exposing the entire system to outside viruses and other malware.
- As part of a repair, a plumber drilled a hole in the ceiling directly above an Exchange server. He then did a poor job of repairing the pipe joint he was working on. In the middle of the night, water began leaking from the pipe. Nature took its course and the water flowed to the hole in the ceiling and poured onto the Exchange server beneath, causing permanent damage to the server.
- Cleaning people were sent into the server room. They saw pockets of dust not just around the server racks, but inside of them as well, and rack doors were partially opened. The cleaning people did what cleaning people do: they cleaned inside the racks and inside of the servers using window cleaner. No one gave them clear instructions regarding cleaning protocol.
- A vendor was working in a halon protected area. He lit a propane torch without notifying anyone and without shutting down the halon system.
- A vendor shut down a PDU so that he could add a breaker to it. The PDU was feeding a key branch server. Many server room visitors may not know what that can and cannot do in that particular location.

Monitoring system components

When designing a monitoring system whose primary role is to limit the occurrence of human error in remote server rooms, four key components need to be considered: video surveillance, sensors, intelligent rack outlets, and monitoring and automation software. **Table 1** provides a summary of the solutions described in this section.

Video surveillance and sensors

What can help when it comes to these situations? Scalable monitoring & automation systems are available in the marketplace and are capable of collecting, organizing, and distributing critical alerts and surveillance videos. An example of this is shown in **Figure 4**. By monitoring power, cooling, the backs and fronts of racks, and the environment, these systems can generate instant fault notification, enable quick assessment of the situation, and provide resolution of critical infrastructure events that can adversely affect IT system availability.

In the case of the human miscommunication examples listed above, consider how such a monitoring & automation system could have helped:

- A physical infrastructure monitoring & automation system, supplemented by a video camera-based security system which would be panning human activity in the rows, would have recorded whenever the motion detectors were activated. Thus, even though no IT person was on site, the activity of the contractors wrapping up the servers, for example, would have been recorded and an alert sent out to the authorized

administrator. Upon witnessing what was going on, the administrator could have issued an order of "cease and desist" and the downtime could have been avoided.

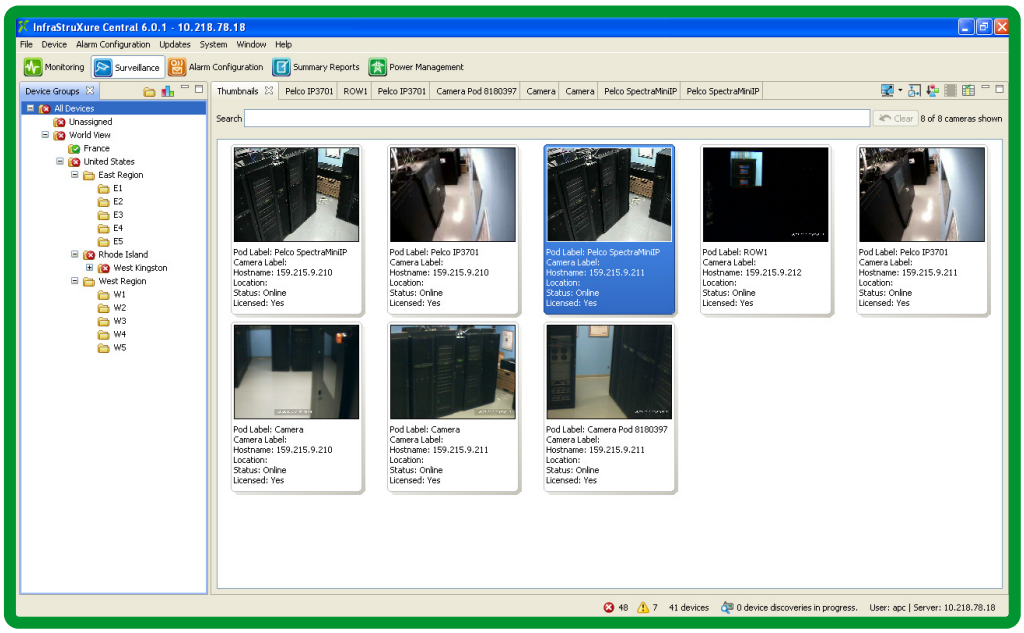

Figure 4

Example of video monitoring triggered by motion detection can help limit instances of human error (sample screen extracted from the APC by Schneider Electric InfraStruxure Central application)

- A monitoring & automation system could also turn devices on or off utilizing low current dry contact output switch closures. Such an approach can be utilized to control locks on racks (see **Figure 5**). Output relay actions can be executed manually or can be configured as automatic alert actions in response to a threshold event or other alarms. In the case of the cleaning people, knowing that they are coming in to clean during the after-hours, the system could be programmed to have all racks locked after 6:00 pm. They could be opened manually or remotely by an authorized individual, but they would be locked to anyone else until the next morning.

Camera systems make sense especially if the server room is supporting credit card transaction applications. Payment Card Industry (PCI) compliance is becoming an important issue. Some state governments are requiring that businesses notify their customers whenever a data breach occurs. As time goes on, the definition of what data is considered personal information will expand to include credit card numbers. Once credit card information is classified as personal information, punitive measures will be forced upon companies with negligent / non-existent security practices. In the future, direct financial incentives may be granted to companies with evolved levels of security who are recognized as PCI compliant. Video surveillance is one of those requirements under the PCI compliance umbrella.

A camera management system typically allows for tracking of facilities personnel, vendors, security personnel, custodians and other visitors who come into the server room or remote wiring closet. The system can determine who was in the room, at what time, and can detect whether the visitor unplugged an existing piece of equipment or plugged in a new piece of equipment. A camera management system could be programmed to record data when it detects motion. On the other hand, an administrator may want to remotely log into the system, activate the camera that is nearest to the visitor and observe his actions. In fact, some of these systems can be equipped with speakers so that the administrator can project his voice from his laptop microphone to either deliver instructions or to provide warnings to the visitor (e.g., "Whatever you do, don't touch that red button!").

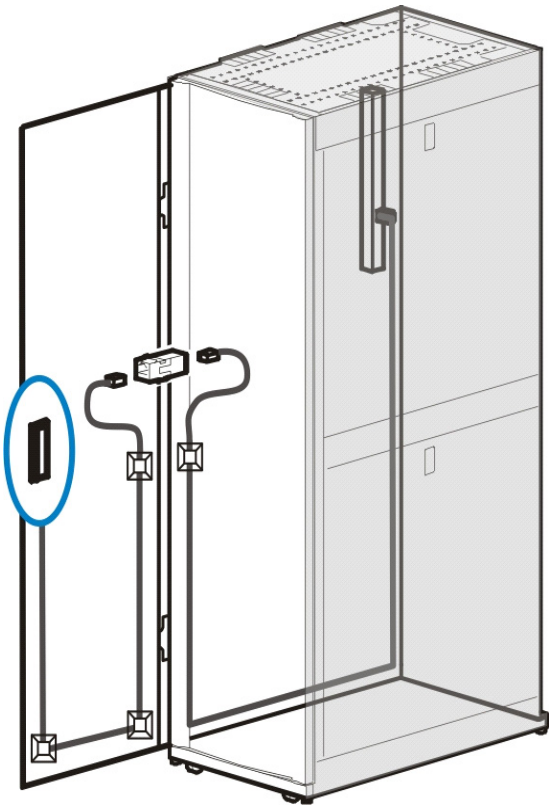

Figure 5

Rack security can be remotely controlled to ward off unwanted entry

Intelligent rack outlets

Intelligent rack outlets are long thin strips of electrical outlets mounted to the inside back of a rack (see **Table 1**). Also known as “rack-mounted PDUs” these devices can be managed by allowing users to remotely recycle power to locked-up equipment. This minimizes downtime by quickly restarting equipment, and any travel time to the remote site for rebooting is avoided.

These devices also allow users to configure the sequence in which power is turned on or off for each outlet. This sequencing allows users to predetermine which piece of equipment is turned on first so other equipment dependent on that unit will function properly. In start-up situations the in rack intelligent power distribution helps to avoid an initial overload of power rushing in which can cause overloaded circuits and further load drops.

In the case of overloaded circuits, the monitoring system prevents overloads by projecting graphical displays of average and peak power usage and by measuring actual consumption through metered rack PDUs (intelligent rack outlets). Therefore the systems administrator has visibility to the power usage of each rack and can make an intelligent decision of where to place the additional equipment that needs to be installed.

Monitoring and automation software

A management & automation system provides the administrator with a wealth of data that will allow for reductions in human error-related downtime. Listed below are a few examples of some of the monitoring & automation management system capabilities available in the market for server rooms and remote data wiring closets:

Table 1

Solution Summary

Solution component	Role	Benefit	Example illustrations
Monitoring & automation	Alarming Equipment Status Reporting Configuration Control	User-set thresholds generate alarms via text message, email, or system postings when conditions such as temperature and humidity rise beyond accepted levels Generates multiple levels of historical data reports in order to spot problematic trends early on Ability to mass configure similar system characteristics (e.g. rack locks, temperature thresholds) to similar devices all at once Ability to reboot hung equipment from remote laptop	 
Video surveillance equipment	Observe human activity	Video storage prompted by motion or alert Detects and records motion, allowing a visual record to be paired with an access or environmental alert, which speeds root cause analysis Storage of error or security breach detection data prevents subsequent occurrences	 
Intelligent rack outlets	Remote startup and shutdown of servers Measured power consumption	Ensures data integrity is maintained during an extended outage Remotely manages outlets so users can turn outlets off that are not in use (prevent overloads) or recycle power to locked-up equipment (minimize costly downtime and avoid travel time to equipment) Allows users to configure the sequence in which power is turned on or off for each outlet – this helps avoid in-rushes at start-up, which can cause overloaded circuits and dropped loads	
Sensors	Door locks, rack locks, fluid detection, temperature monitoring, air quality monitoring	Detects access by unauthorized personnel via door switch Detects the presence of water or high humidity Detects smoke and particles Monitors temperatures in key locations	

Alarming and notification – The alarms set up in a system serve as a trigger. If, for example, a temperature threshold is set at 62° F (16° C) for the bottom of a rack, exceeding that threshold creates an alarm. That alarm then sends out alerts in a number of user-defined ways. An alert could come in the form of an email, a text message, a posting to a web site, or a ring on the phone. These alerts can be as sophisticated as an email to a blackberry containing a graph of the last four hours of temperature in the server room. Or the alert can be as simple as an email indicating that a particular rack door, which shouldn't be open, has been open for more than two minutes.

Equipment status – A basic monitoring system configuration consists of software and a dedicated physical server. The server acts as a central repository that logs information on all configured server room equipment. Information from sensors and cameras is gathered and catalogued as are all the system profiles and thresholds. The level of monitoring can be quite detailed. For example each rack can contain three temperature sensors, one for the bottom, one for the middle and one for the top of the rack, as these temperatures are often quite different from one another.

The status alerts are also useful for monitoring batteries. The failure of a single battery can result in the loss of the critical load. Faulty batteries should be replaced as quickly as possible, but oftentimes no one is keeping track of the age of UPS batteries in remote sites. The cost of replacing one or two batteries is minimal compared to experiencing a failure that causes the closet or server to crash. Basic monitoring can avoid these situations.

Reporting analytics – The data gathered by a monitoring system can be converted into customized reports for the IT administrator to review. In the past, in order to determine temperatures at odd hours in remote server rooms, administrators relied on security personnel or other outsiders to read and manually record information from thermometers on the walls. Now, the administrator can look to the historical data and realize that the temperature has been fluctuating by 10° F (12° C) at night. By reviewing 48 hour reports, 1 week reports or longer time interval reports, the administrator can recognize the problem and then make the case to the building facilities department for the problem to be resolved (if the building comfort system is being used to partially or fully cool the server room). The data gathered by the IT room monitoring system can document that an issue exists and that it could be symptomatic of a bigger problem. From a security perspective, the reports generated by a system also can help the IT administrator quickly determine who has been in which particular rack and for how long.

“The administrator can then immediately identify which UPSs are supporting ‘illegal’ loads and can issue the ‘cease and desist’ order before any of the retail POS systems go down.”

In the case of retail POS stations, for example, a monitoring system can examine the UPSs in the field and produce a report regarding much load is being placed on each individual UPS. If the IT administrator deems that all of the UPSs should be at 50% load, then those that exceed that limit are easily identified. The administrator can then immediately identify which UPSs are supporting “illegal” loads and can issue the “cease and desist” order before any of the retail POS systems go down.

Mass configuration – Upon initial installation, all devices linked into the central monitoring & automation system are logged and profiled into the system. This enables the administrator to configure or initiate mass change (one change affecting multiple devices) later on. Consider the example of door locks on the server room racks. Every single rack door lock does not need to be configured individually. Only one security configuration needs to be pushed to all 50 rack doors (front and back) if that's the decision the administrator chooses to make.

Control – Administrators feel under much less pressure having access to detailed monitoring & automation system data. For example, a system can map the power path and physical system relationships and dependencies. This helps to avoid a mad scramble when a problem occurs to find out where the source of the problem lies. Some systems can also recommend the best location for placement of new equipment based on available power and network ports. This avoids the problem of experiencing an unanticipated shortage of power in a

particular rack. A system may also illustrate the consequence of device failure on rack-based equipment for instant identification of critical business application impacts. This allows the administrator to formulate a plan ahead of time in case a problem occurs, so that any downtime experienced is minimized.

More control over the environment, more alerting, and more historical data can help to foster an environment of less stress. If an investment in video surveillance and centralized monitoring & automation is already being made, then the addition of temperature control, humidity control, dew point data, and other environmental alarms represents a minor additional cost. The evaluation of environmental trends and the review of video surveillance data help the administrator to nip problems at the bud, so that human error is kept to a minimum.

More battle tales

Power and cooling systems are particularly vulnerable to human error because of a lack of knowledge concerning these systems. The incidences below detail some of the risks involved.

- In one incident, the UPS overheated because packages of toilet paper were piled high on top of the unit and obstructed the air circulation.
- A small server room for a temporary project had been set up in an upstairs area of an office building. The team who had set up the room made sure that everything installed was cheap but per spec. They used one of their domestic air conditioners for the cooling, as it had the right thermal rating to match the heat dissipation required for the gear in the room. It wasn't long before a service call had to be placed because of hardware failure. An engineer was sent out, and it was discovered that the temperature was about 110° F (43° C) in the computer room. Unfortunately, the installation team installed the air intake and air outflow of the air conditioner in the same tiny room.
- An unused convenience outlet is like a magnet for anyone who wanders into a server room or wiring closet. Many server rooms have been knocked off-line because of problems with unauthorized equipment plugged into convenience outlets. Vacuum cleaners and drills are perfect examples of what should NOT be plugged into a UPS-powered outlet. In one case, there was a short in the drill which caused a breaker to trip on ground fault, which took down a significant portion of the server room.
- A large retailer had no one at the store site knowledgeable about how the server room or wiring closet functioned. The cashiers came to work and found that their cash registers were down. The headquarters advised them to bypass the UPS and run their systems on street power until a battery could be drop shipped. . Once the battery arrived, a knowledgeable individual had to be sent out to install the battery. Thousands of dollars in transactions were lost that day with the potential for much more if a power outage had occurred.
- Another retail operation was running into problems maintaining uptime on point of sale (POS) stations. This became a major issue, because each time the systems went down, the scales used for weighing goods to be shipped would have to be recalibrated, which would significantly extend the downtime. Upon investigation, the IT manager discovered that employees at the retail sites were “illegally” plugging devices like space heaters and fans into the UPS that were supporting the POS. Since the systems were only designed to handle the normal electrical loads of the POS stations, the overloads caused by the unanticipated extra loads were bringing the systems down.
- A rack of servers was lost because an IT administrator unintentionally overloaded an already maxed out power strip.

Anyone who has spent time managing remote server rooms, in all likelihood, can add human error stories to the list presented in this paper. Fortunately, a number of monitoring tools are available that can help ease the concerns of operators who worry about unanticipated downtime in these remote environments.

Conclusion

Server rooms and small, remote closets are prolific and are often subject to downtime caused by human error. Managing these smaller data centers is time consuming and problematic. Many of these facilities are unattended IT rooms that are minimally supervised.

A four-pronged approach consisting of a monitoring & automation software system, video technology, intelligent rack outlets, and sensor technology can greatly reduce the incidences of human error in these small environments. These systems place critical data in the hands of knowledgeable administrators who can manage remotely and identify problems before they result in downtime.

About the author

Dennis Bouley is a Senior Research Analyst at Schneider Electric's Data Center Science Center. He holds bachelor's degrees in journalism and French from the University of Rhode Island and holds the Certificat Annuel from the Sorbonne in Paris, France. He has published multiple articles in global journals focused on data center IT and physical infra-structure environments and has authored several white papers for The Green Grid.

Resources

Click on icon to link to resource

White Paper Library
whitepapers.apc.com

TradeOff Tools™
tools.apc.com

Contact us

For feedback and comments about the content of this white paper:

Data Center Science Center
DCSC@Schneider-Electric.com

If you are a customer and have questions specific to your data center project:

Contact your **Schneider Electric** representative