A bnp media PUBLICATION

January/February 2011

Mission CRITICAL

Data center and emergency backup solutions

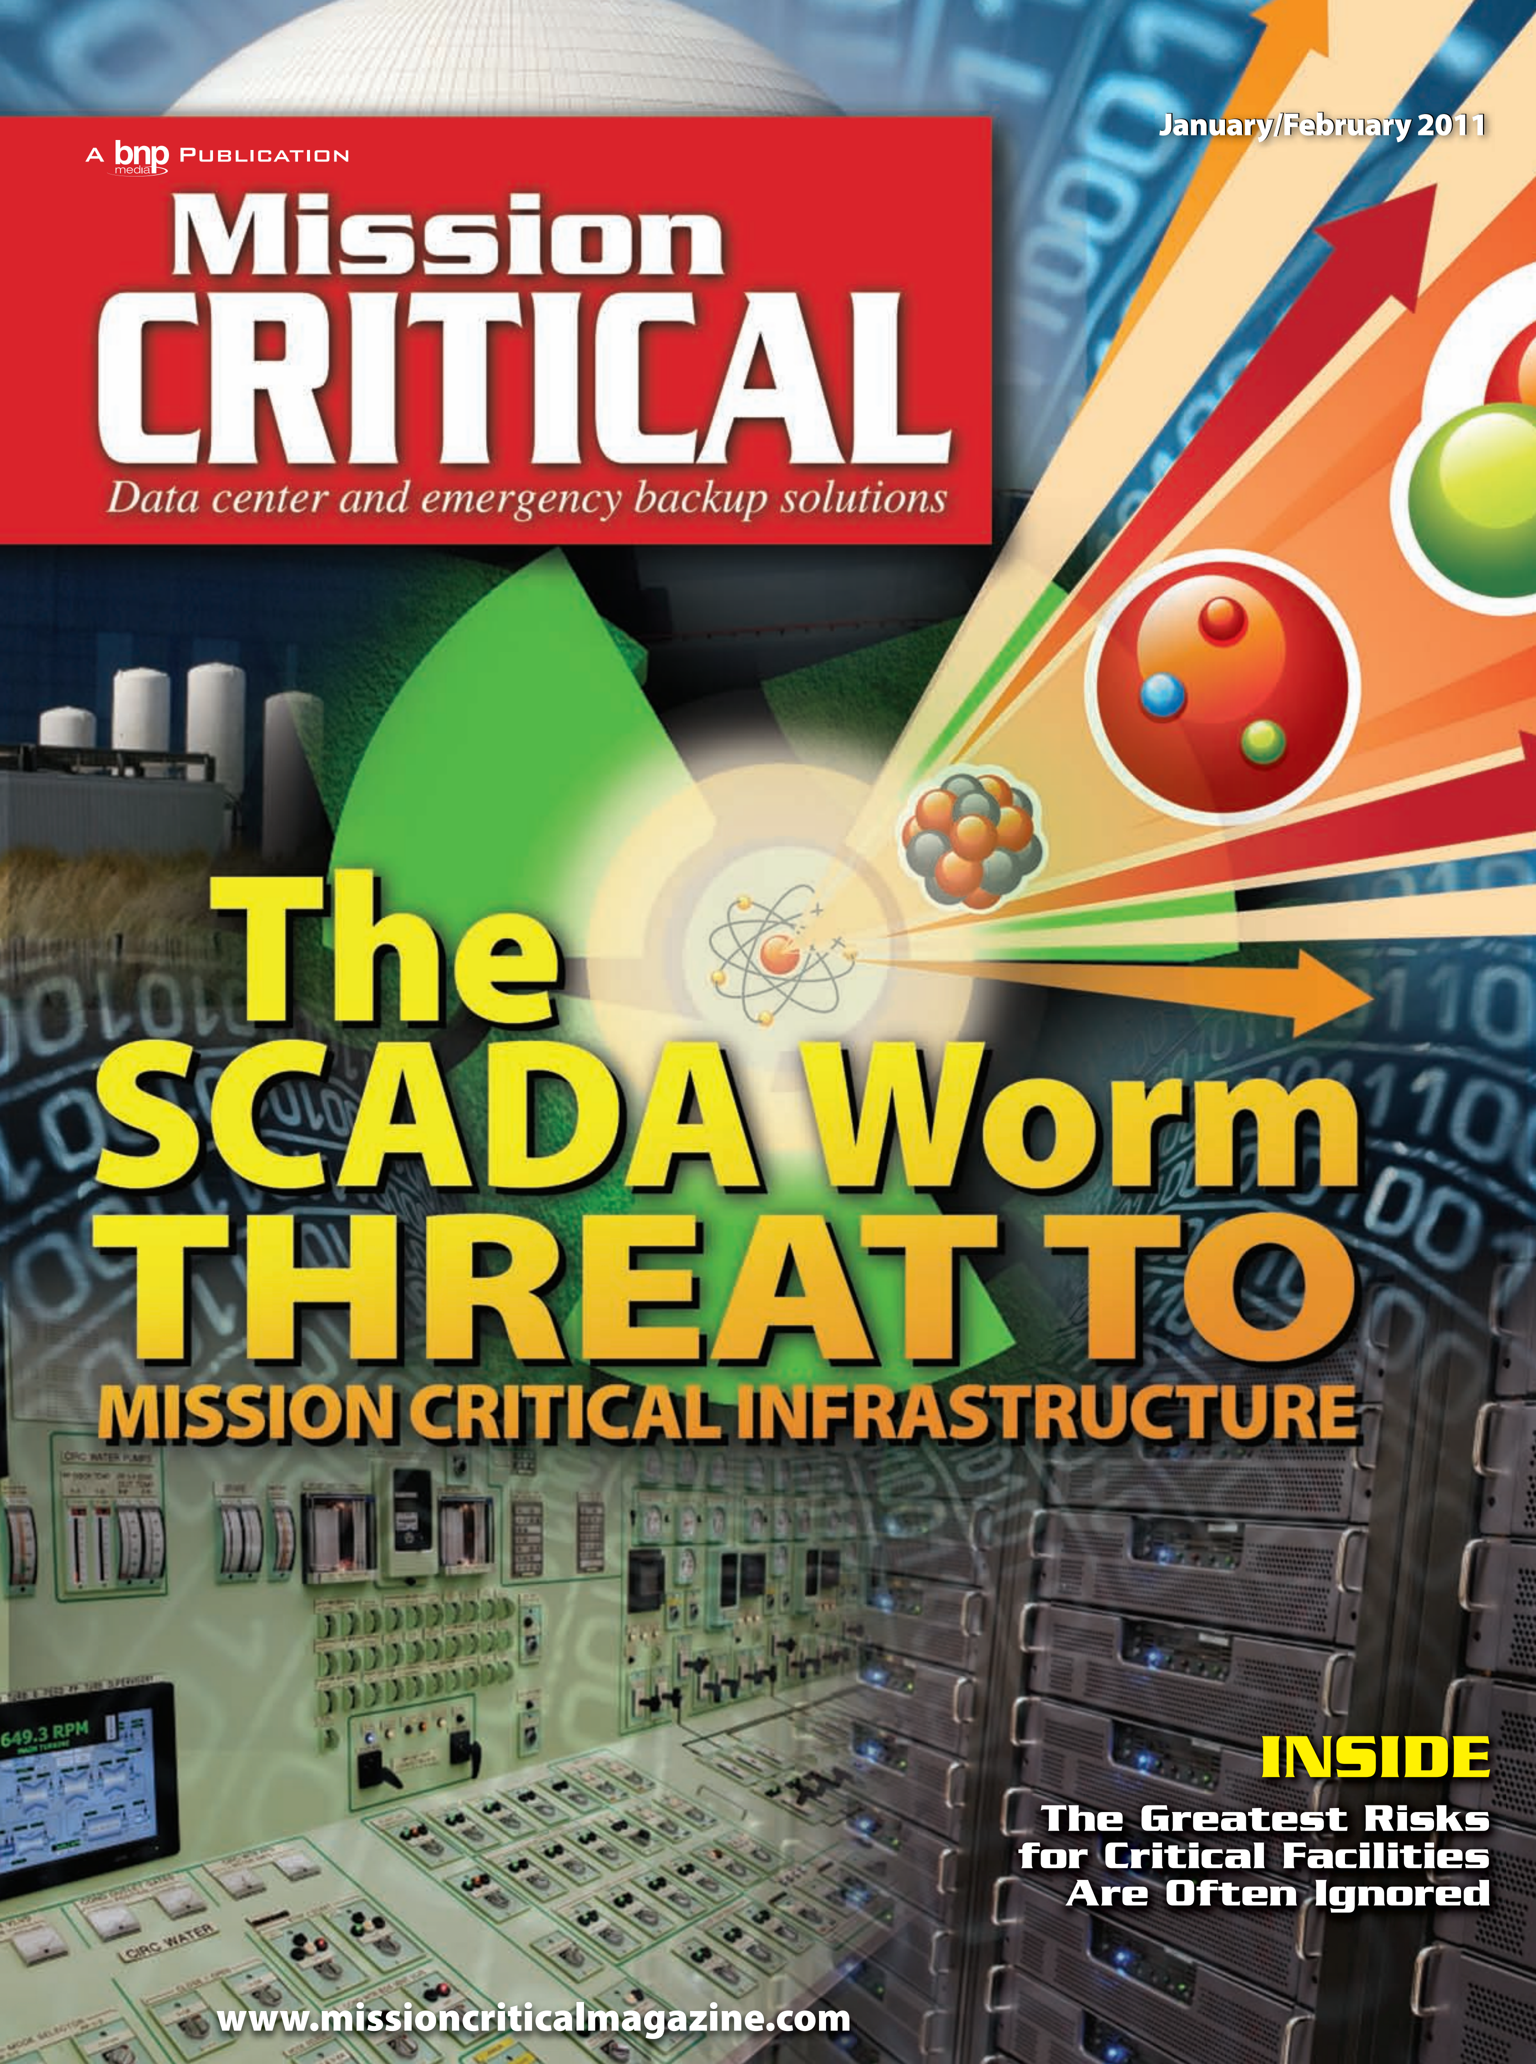

The SCADA Worm THREAT TO MISSION CRITICAL INFRASTRUCTURE

INSIDE

The Greatest Risks for Critical Facilities Are Often Ignored

www.missioncriticalmagazine.com

The Greatest Risks for Critical Facilities Are Often Ignored

Design and equipment configuration can be leveraged to improve availability

BY DAVID BOSTON

Why do data center owners consistently miss their greatest opportunities to enhance the potential for continuous operation of their facilities' systems? The answer may be simple: they do not see examples of comprehensive, yet relatively simple, facilities operations programs when they visit other facilities in their industry. In addition, the issue is rarely discussed at industry conferences. Most conferences and data center tours focus on equipment technology and design configuration, with emphasis on what is new and what level of redundancy is selected.

David Boston is president of David Boston Consulting, a firm specializing in Critical Facilities Operations consulting. As facilities manager for GTE Data Services from 1985-1995, his department achieved multiple years of uninterrupted uptime for a 100,000 square-foot data center facilities operation with a continuous operation objective. While at GTE, Boston conducted the industry's first data center facilities performance and operating cost survey through site visits of 35 data centers across the United States. This survey later became the foundation of the benchmarking programs conducted by the Site Uptime Network, a North American critical facilities consortium of mostly Fortune 100 companies. Serving as primary facilitator and program director for the consortium as it grew from 11 companies to 90 from 1996 through 2009, Boston was responsible for helping member companies optimize infrastructure investment and achieve continuous availability. He may be reached at: (727) 595-3039 or DBoston@tiepoint.com.

This emphasis is understandable but fails to address the most frequent cause of facilities-related downtime: human error. Industry data, such as that published in Symantec's 2008 State of the Data Center Report, the January 6, 2006, edition of *Processor*, and *Mission Critical's* January 2009 webinar survey consistently demonstrate that human error accounts for the vast majority of facilities-related data processing interruptions. Very few data center owners have actively used this information to ensure continuous operation in their facilities. Assessments of over 100 data centers demonstrate that facilities-related downtime is most often caused by inadequate staff coverage, a missing procedure, an incorrect procedure, or inadequate training. Advances in reliable designs and equipment over the years have reduced the number of interruptions to data processing caused by failures of facilities systems and components.

As a result, the amount of investment needed to achieve an optimal operations program is far less than that made in the building and systems installed. Likely, it is only the industry's lack of awareness and focus on facilities operations opportunities that creates this dichotomy.

The astounding rates of success (multiple years between accidents) in industries where life safety is the focus demonstrate that continuous manned coverage, combined with thorough procedures, and rigorous training programs are the key to success. Not only do those who operate nuclear submarines attend multi-month specialized classroom training prior to working on a submarine, they spend 14 to 18 months "qualifying" to operate the various systems once they are onboard. According to Matt Beckman, former Navy nuclear propulsion plant supervisor, submarine crew members are initially provided an orientation to the ship and a systems overview. From that point forward, they participate in two to three hours of training each week. Only after a two-year period do they begin to train others, and only after three years are they normally qualified to train others on all systems. A nuclear submarine team literally operates with over a thousand procedures, most of which are written by the individuals who design the systems.

By contrast, data center facilities team members are typically brought on board just before a new facility is complete, with little knowledge of the construction process. They may have had some other critical facilities experience, but most likely will not. Unless the owner has transferred several team members from another company facility, all will be learning the company culture and objectives while trying to learn to operate the new facility's complex systems.

Most commonly, their only training will occur as part of the commissioning process, meaning it will be informal and rushed (in order to meet the promised start-of-operations date). Design engineering and commissioning consultants repeatedly report the prevalence of this scenario (which the author has also witnessed first-hand as a facilities manager and as a consultant).

Typically few, if any, of the team members will have a chance to operate the equipment as each system is tested. After the hectic commissioning period, some owners will take advantage of manufacturers' offers to provide general training. Ordinarily, this happens when the facilities team is still new, so they do not much retain much of the training. In addition, this training is rarely site-specific. Unlike submarines, ships, or aircraft, which have a limited number of models, each data center facility is unique. The configuration of the systems will vary even when several facilities are simultaneously designed to a "standard." Regional variances in equipment availability as well as building site variances make this an unavoidable reality. As a result, manufacturer provided training is mostly beneficial as a systems overview.

Month	Systems and scenarios	Trainers
January	Fire, severe weather, evacuation	Jim S.
February	Electrical safety, Lockout/Tagout, Arc Flash	Steve G., Art W.
March	Mechanical systems overview	Ann S.
April	Electrical systems overview	Frank L.
May	UPS, UPS swgr, batteries	Paul M.
June	Generators, fuel, controls	
July	Electrical distribution, MCC, EPO, TVSS	Bob R.
August	EPMS & BAS	John T.
September	Cooling towers, water storage, water treatment	Rob M.
October	Chillers, pumps, valves, CRAHs	Rob M.
November	Fire detection and suppression	Ray S.
December	Spill response, hazardous communications	Sarah C.

Table 1. Training calendar used to increase familiarity with essential systems

One owner of a new data center recently provided a two-month window for training the facilities team between the end of commissioning and the start of operations. This critical manufacturing organization also had the foresight to procure a thorough set of site-specific procedures that were utilized during this extended "hands-on" practice period. As part of this effort, the manufacturer gained the additional benefit of testing and editing each of these procedures. All this was achieved before introducing any risk to data center operations. Companies that utilize this approach have found the experience invaluable for their facilities teams. They retain much more knowledge as a result of this extended practice, which enables them to better respond to and resolve unexpected incidents before they result in downtime. And they will have a better chance of seamlessly performing system "changes-of-state" (transferring equipment offline and back) during planned PM activities.

Unfortunately, most owners consider scheduling this amount of time for testing and training an unaffordable luxury, as they need the new computer operations space immediately. Unexpected growth in processing demand and/or a delayed new facility project approval process are the most common causes for this haste.

Without a substantial testing and training window, the average facilities team will remain uncomfortable with the data center's systems and their normal operation until many months after start up. Several owners who have inherited this challenging situation have helped to accelerate the learning curve by ensuring their staff's

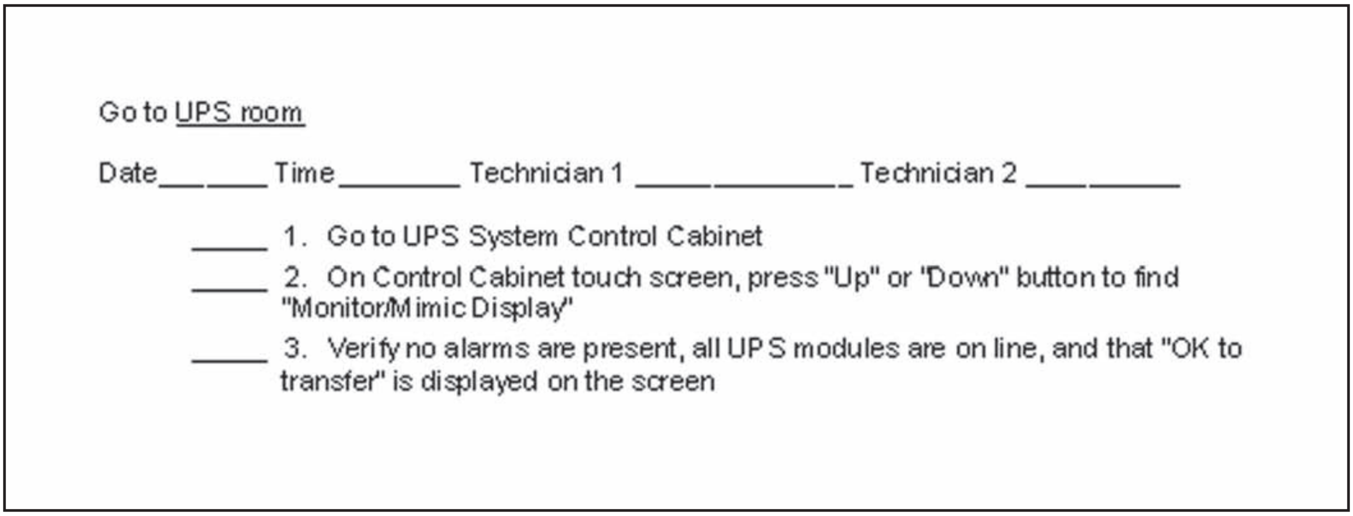

Go to UPS room

Date________ Time________ Technician 1 ______________ Technician 2 __________

_____ 1. Go to UPS System Control Cabinet

_____ 2. On Control Cabinet touch screen, press "Up" or "Down" button to find "Monitor/Mimic Display"

_____ 3. Verify no alarms are present, all UPS modules are on line, and that "OK to transfer" is displayed on the screen

Table 2. Procedure documents should include a clearly marked space or box in front of each written step to permit the individual reading it aloud to check each step as it is completed

participation in all planned preventive maintenance activities, so they may observe and practice the steps necessary to transfer equipment "offline" and to restore a system to "normal."

This repeated experience, combined with the creation of site-specific procedures, permits facilities teams to acquire ownership of these critical change-of-state activities, rather than relying on the system service vendor to perform or lead these activities. Because of the variance in installed equipment configurations from one customer facility to another, reliance on vendor technicians for system transfers presents an increased risk.

Practicing system change-of-state procedures is just one method successful organizations have implemented to hasten the acquisition of knowledge when the luxury of a dedicated training window is not possible. Additional strategies include:

- Identifying and conducting key system training sessions on a monthly basis (see figure 1)
- Annual repetition on the most critical systems
- Both classroom and "hands-on" applications, with written tests
- Designated "experts" as trainers (staff members, engineers, vendors)
- Development and testing of site-specific emergency and change-of-state procedures
- Contracted assistance, unless staff size permits rapid progress
- Emergency procedures first, then change-of-state procedures in the order of upcoming PM activities
- Clear and consistent format, written at the level of understanding for the least experienced team member
- Testing with each team member separately (or in pairs), by shift

The commercial airline industry has an excellent practice to draw upon when creating procedures and training programs. Procedure documents should include a clearly marked space or box in front of each written step to permit the individual reading it aloud to check each step as it is completed (see figure 2). During both training and actual PM activities, the individual holding the procedure should read the step aloud, the person performing the step should repeat it aloud, and after both nod their heads, the person performing the step should proceed. This "pilot/co-pilot" process should be followed rigorously. Without it, the potential to miss a step is too great. This process causes both individuals to contemplate each step. It will allow a mistake or unclear instruction to be noted, allowing the team to back out and clarify the instruction before proceeding again at a later time or date.

Aside from a comprehensive procedures and training program, experience as a data center facilities manager and as a consultant who has assessed over 100 critical facilities supports the following initial strategies:

- Providing continuous shift staff coverage (system problems occur equally on all shifts)
- Assigning at least two individuals per shift, so that functional work may be accomplished safely. A savings may be achieved over three to five years by in-sourcing selected maintenance activities typically contracted
- Recognizing successes, sometimes with financial rewards
- Shared continuous operation incentives
- Public praise and small monetary rewards are highly effective
- Retain those who innocently cause a downtime event. Often they become advocates for caution, which is contagious
- Clearly defining responsibilities and ownership
- Limiting the number of individuals involved in power

cabling and precisely identifying (physically) demarcation points
- Master planning the computer hardware layout: team one person from each department to create and maintain the plan, to provide maximum cooling potential for hardware
- Controlling processes
- Making annual sign-off on a thorough rules document mandatory
- Substantially limiting unescorted access to the computer room
- Requiring executive endorsement
- Requiring clear endorsement of each department head, which is particularly critical when multiple departments are responsible for a data center's continuous operation, in order for these processes to be successful

There are two common misperceptions regarding data center facilities that should be dispelled. Proactively sharing these with executive management team should enable the facilities staff to avoid (or minimize) scrutiny at a later date.

- A fault tolerant facility will always outperform any facility with less than fault tolerant design. Experience shows that facilities with less than fault tolerant designs will often experience longer periods of continuous operation when they have comprehensive staff coverage, training, and procedures programs in place (if the fault tolerant facility does not).
- My new data center is invulnerable the day it begins operation. Instead, most new facilities teams will not be provided the site-specific procedures and rigorous training experience, prior to operation that would allow this to be true. The first one- to three-year period will be higher risk. This risk period will be longer if a strategic procedures and training program is not implemented.

Those who are able to successfully justify and implement the recommended processes in this article will realize substantial rewards for their extensive time and effort invested. Continuous operation will increase exponentially. Several examples in North America alone have been verified where data center facility systems have operated without an interruption to data processing (single server or entire facility) for periods of five to 10 years. A common application that is deemed critical, values the cost of a downtime event at more than $1 million per hour, and has a normal total system recovery time that matches the industry average of four to six hours will see dramatic payback on the initial investment. ■