



# Global Information Assurance Certification Paper

Copyright SANS Institute  
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

## Interested in learning more?

Check out the list of upcoming events offering  
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"  
at <http://www.giac.org/registration/gsec>

## **Printing the paper and serving the news after a localized disaster**

John Soltys

GSEC assignment 1.4b option 2

February 18, 2004

### **Abstract**

A case study detailing the implementation of a business continuity plan for a regional newspaper. This study covers the requirements-gathering process, testing, and implementation of a series of plans jointly developed by members of the newsroom, IT, online staff, and operations. The plans resulted in redundant systems co-located at an offsite printing facility, procedures for relocating staff, and development of a distributed website presence.

## **Printing the paper and serving the news after a localized disaster**

This publishing company is the largest newspaper in the state and serves a community of more than a million readers with its print and online products. The main office is near the downtown core with several additional buildings located within a few blocks. The printing plant is approximately 20 miles to the northeast. Of all these properties only the printing plant was built in the last 20 years.

The proximity of the majority of the employees led to the development of a centralized network infrastructure based downtown. Internet connectivity was provided by a pair of redundant DS3 connections linking the ISP's data center to the corporate network in the basement of the main building. A different vendor using diverse paths provided each DS3. Fiber connects the rest of the downtown campus buildings, and a DS3 (with backup bonded T1s) links the printing plant to the internal network.

Although this network topology resulted in significant cost savings over an infrastructure with independent networks and Internet links at each facility, it created a single point of failure in the data center where the corporate network interfaces with the public Internet. Should an event occur that disables the main building the network for the entire company would be down.

In addition to providing the gateway to the Internet, the main building hosts several critical systems for publishing the printed newspaper and providing an online news product. These systems, like the network, are built with redundancy in mind, but only at the component level. Without geographic redundancy a building-wide failure would render the clusters and redundant storage devices useless.

Included in this list of critical systems are the content management systems for news and ads, page layout systems, the VoIP phone system, and the publicly accessible HTTP servers and their associated database servers. The loss of

these systems would result in failure to publish, damage to the brand, and would fail the community that depends on the company for its news.

In a time when newspaper readership is falling nationwide<sup>i</sup> the newspaper's relationship with its readers is more important than ever before. The history of continual publication grows more valuable every day. As a result, the executive leadership of the company decided to place a high priority on developing solutions that would mitigate the risks posed by a failure at the company's headquarters.

In order to fulfill management's expectations a cross-departmental group was formed to identify and address the likely scenarios that could threaten the company. The group included representatives from information technology (IT), operations, news, and the online staff.

IT was primarily concerned with the physical infrastructure and servers that provide services to their internal customers. Operations was charged with general disaster preparedness from an employee standpoint. The responsibility of news was to coordinate a newsroom of more than 200 writers, editors, photographers, and designers split between the main building and several community offices. The online group took on new importance in the event the paper could not be printed or could not include all the usual content.

With so many different perspectives involved in the project it was agreed that the only way progress could be made would be if different scenarios were defined and addressed in phases.

The first scenario was one in which the main building was uninhabitable, but otherwise unaffected. This type of disaster might occur as a result of a terrorist act such as an anthrax or bomb threat or a more innocent problem such as a natural gas leak. In this scenario all the systems are still up, but they cannot be accessed locally.

The second scenario was more damaging in nature. It called for the loss of the network and/or hardware in the building. Potential causes for this type of disaster included fire, earthquake, water damage, and power outages. At the lesser end of this spectrum the systems are intact, but cannot be operated. At the opposite end is the complete destruction of the entire building. This scenario was envisioned as a three- to five-day event after which the production systems could be replaced or restored and operations would return to normal.

It is important to note neither of these two scenarios addresses the possibility of a catastrophic disaster affecting the entire region served by the newspaper. Such a disaster would result in the destruction of not only the main building but also the printing plant and the infrastructure needed to provide even the most basic news coverage and distribution. It was determined this type of disaster is so

unlikely as to not have a sufficient return on investment to be worth considering until the more likely scenarios are fully addressed.

The first scenario is primarily procedural. Operations had already put in place a series of plans to ensure the safety of employees at the downtown campus. The fire alarm would signal an immediate evacuation of the affected building, instructing employees to exit and meet their floor wardens in the park across the street. An alternate location farther from the building was also designated in the event the park was not considered safe by the authorities.

Once the building was emptied there were three requirements:

1. Administrators must be able to access systems without physical access either to the system or to their usual workstations.
2. A print newsroom must be set up to allow the reporters, editors, photographers, and designers to meet their deadlines and send pages to the printing plant.
3. An online newsroom must be set up to allow the news producers to move content from the print systems to the online systems.

The first requirement was satisfied by identifying desktop systems in other campus buildings that could be repurposed in the event the system administrators needed to access systems in the closed building. The tools the administrators relied on were stored on networked file servers accessible from anywhere on the corporate network so they could be easily installed on the administrator's newly acquired PC.

A training lab in a secondary building was designated the makeshift newsroom. The PCs in it were configured to allow administrators to install the proprietary editorial software needed by the news staff. These systems would enable the newsroom to communicate with the systems in the closed building and transfer pages to the printing plant.

The final requirement was fulfilled by identifying systems in the online department's building from which online staff could import news content from the print systems and prepare it for online publishing.

The documents detailing these procedures were put to the test with much success. There were few problems encountered and pages were successfully sent to the printing plant.

While drafting the plans a great deal of care was given to simplifying the instructions so even the least technical employees could accomplish their goals. Although it was considered somewhat demeaning to some that the instructions were so explicit, their concerns were alleviated when they considered the stress employees would be dealing with during such an event, especially if it occurred close to a deadline when extra minutes equate to thousands of dollars.

At the most fundamental level the second scenario differed from the first in the status of the systems used to print the paper and serve the news online. The events leading to the disaster could be as commonplace as a burst water pipe or a power outage. Though not typically considered disasters in the traditional definition of the word, these events could wreak damage on the infrastructure needed to publish as serious as destruction of the building by fire.

Three main requirements were identified by IT that any disaster plan must satisfy before it could be adopted.

1. Writers and editors must be able to produce content without the use of their usual tools.
2. Designers must be able to build pages.
3. The systems used to prepare the pages for the presses must be available.

Initially, the newsroom pushed hard for complete replication of the systems they depended on. While it was generally agreed this was a laudable goal it was universally acknowledged to be an expensive proposition.

An analysis of the data flow revealed there was a natural break in the publishing process. Before this point the underlying systems served to facilitate management of large quantities of data in a very efficient way. The content management system allowed hundreds of newsroom staff to produce stories, build packages, and follow a complicated workflow ensuring no content went into the paper without the appropriate review and authorization. However, the content management system did not solve any problems that could not be dealt with manually if the number of staff and stories was reduced.

Management accepted that during a disaster the paper would likely be smaller than it would otherwise be. Additionally, it was prepared to see more syndicated than original content as long as the paper was printed and there was a way to expand the content if the disaster lasted more than three to five days. Management was not ready to accept the cost of complete replication.

In order to facilitate some semblance of a workflow, a single fileserver was purchased and was configured to act as a content management system at the printing plant. Rather than the intricate workflow the newsroom was accustomed to, staff would use a system of folders each representing a different state of readiness. Unedited stories and photos would be deposited into the "raw" folder. Editors would pick up this content, edit it and place it in the "ready" folder. Page designers would take items from the "ready" folder and place it on the page.

With this system in place there was no longer a need to replicate the massive print content management system at the printing plant. However, once the pages had gone through composition and were ready to enter the workflow they would need to be sent to the RIPs, which were located only in the data centers in the

main building. The RIPs rasterize the pages producing files with a specific resolution, line screen, and dot shape. Without these systems pages could not be sent to the presses.

As part of an earlier upgrade several older RIPs were made obsolete. Although these systems did not have all the functionality of the primary systems the new workflow could be adapted to include steps to save copies of the pages in a less efficient format compatible with the old RIPs. The old RIPs were moved to the printing plant's data center and their use was restricted to disaster operations only.

The output of the RIPs is transferred to the next step in the process. Those servers and the rest of the infrastructure already existed at the printing plant. The need for disaster planning from the print systems perspective was completed. Producing those pages, however, would require an additional 40 employees onsite at the printing plant, each with a networked workstation.

Early drafts of the plan called for PCs already at the printing plant to be repurposed as newsroom systems. However, the printing plant supported a very small population of computer users. Those with desktop systems capable of performing well enough for the newsroom to use would need their systems throughout the disaster.

Redirecting the surplus from a planned mass upgrade to storage at the plant solved this problem. The systems were rebuilt with a stripped down OS and set of applications that would just meet the requirements of the newsroom staff without incurring any additional expense. The network devices needed to support the new nodes on the network were purchased through eBay, which was recognized as a valuable resource for making the most of limited funding<sup>ii</sup>.

Late in the planning a purely logistical question came up. How would an additional 40 employees impact not only the plant, but also the disaster supplies stored there? As part of Year 2000 preparations stores of emergency food and water had been put in place to ensure employee safety during a regional disaster, but these supplies would be quickly exhausted with the influx of additional employees.

The team had accounted for the survival of the systems needed to publish, but had neglected to consider the personnel side of the equation. Sleeping space for employees was identified, but budgets didn't allow for purchase of cots or sleeping bags. Employees would be responsible for their own gear. The plant had a kitchen, but no stocks of food, so food would need to be brought in from outside. A supply of non-spoiling military meals was purchased should the staff be isolated in the plant.

During the initial risk analysis three primary risks were identified:

1. Failure to publish.
2. Damage to the brand.
3. Failing the community that depends on the company for its news.

The plan to print the paper with content generated at the printing plant addressed only the first. In the event the disaster was truly localized to the company's headquarters the general public would likely not care about the company's problems, but they would be concerned that they were receiving a smaller paper with fewer stories.

When the online department was formed in the late 1990s it was considered an experiment. Many of those already employed at the company thought the new endeavor a waste of resources or a threat to their livelihood. Even after the explosion of the Internet the online department was not thought to add much value to the enterprise.

The threat to the brand posed by a company disaster forced even the most print-focused members of the team to reconsider the benefits of working closely with the online department. Posting content online required comparatively little effort as opposed to preparing the printed product. Most importantly, it would be possible to protect the brand and the relationship with the reader online. By adding only a few additional employees to the disaster team located at the printing plant the newsroom would be able to deliver an online experience very close to normal.

There were three technical obstacles to overcome in order to serve the news online during an event of this type. The first was the loss of systems during the disaster that forced relocation to the printing plant. Three principal systems were needed to deliver content online: the content management system, the database storing all the content, and the webserver with which the end user interacted.

Unlike the content management system employed by the newsroom, the online content management system ran on a desktop-class system. The cost of replicating it was low compared to the cost associated with building the site by hand during a disaster. Using the online content management system a single employee could publish the entire paper in an eight-hour day compared to dedicating 10 employees to the effort if they attempted to do it manually.

The database was also a relatively lightweight server that could be used as a development server during normal operations. The only requirement was that production data be sent to the printing plant database once a day. At most, a day of data would be lost. This would allow not only for new content to be posted, but also for existing content to be maintained as usual.

Making the content available to the public posed a more significant challenge. On a typical day bandwidth consumption topped 40Mbps and 2.6 million pages.

Although a vendor's content delivery network reduced the load on the origin webserver to a mere 10Mbps the pool of servers was still integral to the process.

The content delivery network's proxy servers stored a page in cache only after a user had requested the page. As a result the caches could only be populated with the assistance of both a functioning origin webserver and traffic through the proxy server prompting the cache to store the page. Without a live connection between the content delivery network and the origin webserver, content could not be served.

Although the default behavior of serving content from cache without communicating with the origin servers would work well during disasters of limited duration it did not provide the relocated staff the opportunity to update the caches. During a long-term disaster the content would become progressively more stale and less valuable.

The company had several holdings outside the immediate region that had relatively high bandwidth connections to the Internet, which made the possibility of geographically diverse data centers a possibility. However, a cost analysis quickly showed that such a project would require substantial investment in upgrading the existing load-balancing infrastructure. Although the existence of multiple data centers would mitigate the risk of a regional disaster affecting both the downtown campus and the printing plant, the cost was deemed too high given the likelihood of such an event.

A solution was developed using two additional products offered by the content delivery network vendor and customization to the online content management system. The first offering was traditionally used to store objects that changed very rarely. It was essentially a large storage facility accessible by rsync and ftp and easily configured to simulate the directory structure of the news site.

The second part of the configuration involved the vendor's own failover service, pointing, however, to the online storage instead of a second data center. When the origin webserver failed to respond the user was sent a generic failure page indicating technical problems. After 10 minutes of continuous failure the user would see content served from the storage server as though it came from the origin webserver.

Although the origin webserver served primarily static content that could be replicated on the vendor's storage site there were features on each page that were generated dynamically. These portions of the page needed to be stripped from the code before the pages were sent to the storage site, or the user's experience would be plagued by a series of failures.



The content management system used by the online department was written in-house which allowed it to be substantially modified in order to publish multiple versions of the site. By stripping out features known to fail during a disaster and rewriting URLs to redirect users to explanation pages instead of server failure pages the experience of the user was salvaged.

By using the content delivery network's failover and online storage there was no longer a need for direct access to web servers at the printing plant, but Internet connectivity for the staff's workstations remained a requirement. With much of the news gathering infrastructure in disarray the Internet would be even more important as a writer's tool than usual. The printing plant could sustain only an additional 50 employees so most of the writers and photographers would be filing their content online.

Since the printing plant's normal Internet connectivity ran through the downtown headquarters the network at the plant would be isolated during a disaster. Several ISPs were contacted about providing connectivity, but the monthly cost was excessive given that the link would be idle except during tests and actual disasters. The IT networking group lacked the BGP routing experience to use the additional high-speed connection as another route to the Internet, so a simple DSL connection with a low monthly fee was selected as the primary Internet connection for the disaster network.

The ISP providing the DSL connection also hosted several email addresses and FTP accounts in a backup domain. During a disaster these services would be the primary means by which staff in the field would file their content. With bandwidth at a premium over a relatively slow connection, access to the Internet would be limited strictly to those with a demonstrable need to get outside the printing plant's network.

To provide some measure of redundancy for the consumer-grade primary connection a satellite connection was purchased. Although the bandwidth was limited in terms of throughput per second as well as throughput per hour the satellite would provide a minimum level of service if all other connections failed.

With plans in the late stages of development for both scenarios a series of tests were planned. The first addressed only the situation during which all the systems were up but the employees were evacuated from the facility. The test required little more than selecting a handful of employees and relocating them in the other campus buildings. They were tasked with accessing all their normal systems and producing a page. The test ran smoothly, resulting in a film of the page they worked on.

The second scenario was significantly more disruptive to test. It required that the link between the printing plant and the rest of the company be severed in order to simulate the loss of network services. Doing so would cut off not only those

testing the disaster plans but also the employees located at the printing plant. More importantly, the window for the test was limited by the production schedule for the printed newspaper.

When the link was taken down, technical, but otherwise unprepared, IT, online, and news staff were given the disaster plans, which included detailed, step-by-step instructions on bringing the systems back online, readdressing them to use the DSL link, and producing content.

The concept of giving the plans to employees who had not seen them before was a controversial one. The window during which the network could be isolated was small enough that, should any serious problems arise that could not be addressed immediately, the test would have to be repeated at a later date. However, the unrehearsed test identified problems with the plans that would not have been found by those overly familiar with them.

Full-scale tests are planned twice a year ensuring that the plans and equipment are ready for use in the event of a disaster. With each test a new selection of staff is introduced to the plans so the pool of experienced employees increases. At the same time, access to the plans is strictly controlled because they contain all the information needed to administer the systems, including passwords and other sensitive information.

Before the planning began, an event that downed the systems in a single cabinet in the data center would have had a crippling effect on the entire enterprise. Servers, workstations, and phones would all be isolated, and the ability to publish the newspaper would be threatened.

With the plans in place all but the most catastrophic of disasters would not prevent the paper from publishing. Safety concerns forcing the evacuation of the headquarters will result in a paper printed using the same systems, but with all interaction from a remote location.

A short-term systems outage will trigger a 10-minute process resulting in a feature-reduced Web site, but one still serving the news and still available. Ten minutes after the origin web servers come back online the automated monitors will detect that the systems are back up and initiate failback to restore all the site's features.

A medium- to long-term systems outage will trigger the relocation of print and online staff to the printing plant 20 miles to the northeast where the news will be collected and formatted for the presses. The systems will be different and the process less efficient, but the paper will still be printed. A backup version of the online content management system will be used to publish and update what news is available online to maintain the freshness that results in repeat visits.<sup>iii</sup>

Most importantly, the company's brand will be protected. Failing to publish in print or online damages the brand because it betrays the public's trust. Although brand loyalty remains strong<sup>iv</sup> the risk to the brand was paramount during risk analysis. Should readers' impressions of the company be damaged during a disaster, resulting in a switch to a competitor, it would be difficult and expensive to reacquire as loyal readers.

Although not the fully redundant, geographically diverse infrastructure that would guarantee an interruption-free publishing environment, the plans put in place would be sufficient to mitigate the risk of a disaster. By ensuring the public could get the news in either printed or electronic form the business continuity project fulfilled its mission. In spite of the tight budgets imposed by difficult economic times, the paper would be printed, the news would be delivered, and the brand would be protected.

## References

Newspaper Association of America. "Facts About Newspapers 2003."  
[http://www.naa.org/info/facts03/4\\_facts2003.html](http://www.naa.org/info/facts03/4_facts2003.html) (2 Jan. 2004).

Hoffman, Thomas. "Surveys Show No Big Boost Likely in IT Budgets This Year." Computer World.  
<http://www.computerworld.com/managementtopics/management/story/0,10801,82188,00.html> (5 Jan. 2004)

Frauenheim, Ed. "Report: IT spending unlikely to climb." News.com.  
<http://news.com.com/2100-7341-5089199.html> (7 Jan. 2004)

Ulfelder, Steve. "Online auctions offer IT bargains, risks." Computer World.  
<http://www.computerworld.com/hardwaretopics/hardware/story/0,10801,76944,00.html> (7 Jan. 2004)

Czar, James. "Sticky Content..." Web and Electronic Marketing Ideas. CincinnatiArts.com.  
<http://www.cincinnatiarts.com/old/ca/presentations/webpromo/sticky.html> (10 Jan. 2004).

Cummings, Joanne. "Five ways to boost your Web traffic." Streaming Media IQ.  
<http://www.streamingmediaiq.com/resources/tips/686-eMedia%20Tips.html> (10 Jan. 2004).

Sharp, Byron, Dr. "Hook 'em while they're young?" B&T.  
<http://www.bandt.com.au/news/92/0c003192.asp> (10 Jan. 2004).

- 
- i [http://www.naa.org/info/facts03/4\\_facts2003.html](http://www.naa.org/info/facts03/4_facts2003.html)
  - ii <http://www.computerworld.com/managementtopics/management/story/0,10801,82188,00.html>,  
<http://news.com.com/2100-7341-5089199.html>,  
<http://www.computerworld.com/hardwaretopics/hardware/story/0,10801,76944,00.html>
  - iii <http://www.cincinnatiarts.com/old/ca/presentations/webpromo/sticky.html>,  
<http://www.streamingmediaiq.com/resources/tips/686-eMedia%20Tips.html>
  - iv <http://www.bandt.com.au/news/92/0c003192.asp>

© SANS Institute 2004, Author retains full rights.