From:
http://www.itbusinessedge.com/cm/community/features/articles/blog/hurricanes-wrath-drives-disaster-planning/?cs=23228

Hurricanes' Wrath Drives Disaster Planning

by Susan Hall, IT Business Edge
Jan 9, 2009 12:00:00 AM

On the morning of Aug. 28, 2005, John Chaffe, director of IT for Tidewater Inc., got a call from his CEO. Overnight, Hurricane Katrina, taking aim at New Orleans, had been upgraded to a Category 5 storm. The company operates about 400 supply vessels for the global offshore oil industry. Since the storm had picked up steam, Chaffe's boss wanted to talk about what to do.

Chaffe had been promoted to IT director six months before. He says that after Hurricane Ivan, his predecessor and the IT steering committee decided to run another T1 line to the company's Houston office in case the New Orleans headquarters lost power for more than a day or two. Their plan was to use the company jet to fly the data center to Houston. The major dangers were considered high winds and power outage. Housed on the 19th floor of a downtown New Orleans office building, flooding didn't seem that big a threat, he says.

But when he went to do that, the pilot said he couldn't take any personnel. And he'd just drop the servers off in Sugarland, Texas, before taking the jet to Dallas. Chaffe wasn't about to let the servers go unattended at an airport, so he rounded up two SUVs and packed what he could for the drive to Houston.

"It's not a lot of equipment. We have a pretty lean data center here," he explains. "We have an AS/400, which handles most of our financials, and a hodge-podge of other servers. Our business is very boat-specific. While we have 8,400 employees probably 7,000 of them work on boats and never see a computer except for what might be on the bridge of the vessel. So we provide computing services for about 1,400 back-office

personnel. Most of this is company-specific information, logistics software and all the corporate e-mail. ... Operationally, e-mail is very important to us."

Enacting a New Plan

But the trek to Houston was not to be a short trip.

"We saw that everything went to hell in a hand basket and we realized we were going to be around (Houston) for a while," Chaffe says. Katrina weakened to a Category 3 storm before slamming the Gulf Coast, but created massive damage nevertheless.

"We were in shared office space (in Houston), one of these executive office-suite areas. We commandeered as many of the spare offices as we could and got an electrician to come the next morning and we built a data center. We ran all the cable and put all the wireless up we needed to build a new main office in a matter of eight hours."

"We had a $75,000 line of credit in 15 minutes and went over to CompUSA and started buying everything we needed off the shelf, because I didn't take everything. I took boxes; I didn't take KVM switches, monitors, keyboards, that sort of thing. We spent the day putting it back together and failing the Internet connection over to Houston. ...The only thing we didn't have running was the AS/400 because it generated too much heat, and we had to wait for a spot cooler to put in that room. Once we got that, we were able to fire it up. But e-mail was only down 54 hours. That's not bad when you don't have a plan."

Unexpected Rita

The only problem was a month later another Category 5 storm, Hurricane Rita, had a bead on Houston. Chaffe summed up the mood, probably in understatement: "By this time, we're pretty much out of sorts."

With the New Orleans building still inaccessible, they struck the data center, loaded it into the SUVs, drove to the airport at Sugarland, loaded it into a jet, and flew to Mobile, Ala. There they spent the night on the ground as the storm hit Texas.

"As soon as the storm passed, we came around the south side of the storm and went around it in Victoria, Texas, the next day," Chaffe says. "My people were still on the ground there, they drove to Victoria, piled the whole data center back into the two SUVs and drove it back to Houston and had it all set back up by 7 o'clock that night.

"At that point, my CEO came to me and said, 'We're not doing that again.'" He wanted a replicated system that would allow the company to operate in another city with the flip of a switch.

The problem up to that point had been getting the company to commit the money to such a project, because IT hadn't been considered a priority, Chaffe says.

"Senior management is not aware of what IT provides to them until they don't have it. So they had a little better focus at this point in time," he says.

As far as it went, the idea to run the T1 line to Houston worked well during Katrina, because the company is so dependent on the Internet, he says.

"But every person has his own disaster. Here in southern Louisiana, we're in danger of flooding from hurricanes. That's what is most likely to happen to us. In California, it's earthquake. You've got to plan your disaster for what's most likely to happen to you, then look at where your problems are going to be," he says.

"We learned that you can't make your alternate data center on the coast that's got the same problem. I want to have my alternate data center away from the coast. An F5 tornado might hit Richardson, Texas, [Dallas] but that's a pretty remote possibility compared with the risk of a hurricane hitting Houston.

"So we decided to use Compellent's thin-replication and boot-from-SAN capabilities and we migrated, rebuilt all the servers, put it on HP DL360s. We put a bunch of them in New Orleans and an identical set in Dallas and ran a DS3 line between here and Dallas. We put a Compellent storage network at both sites and replicated all the data to Dallas. Then Gustav came, so that was the test."

Revisiting the Plan

Chaffe says the company now revisits its disaster-recovery plan at least once a year. It's developed more strategies for communications in a crisis.

In fact, twice this year — for Hurricanes Gustav and Ike — it used AT&T's CrisisLink service, which allows the company to predefine certain phone numbers to be failed over to other phone numbers.

It also now has a contract with a satellite provider for a megabit of Internet and 20 IP phones to be dropped and deployed anywhere in south Louisiana, Texas or Mississippi with 48 hours' notice. That could come in handy if there's an outage at its branch offices in Houma or Amelia, La.

"We've revisited how we deal with our vessels in the Gulf of Mexico and the command and control systems. All of that has been revamped," he says. And most of the office staff now have laptops instead of desktops.

"Now if we see a hurricane's coming, we tell people they need to take their laptops with them," he says.

But one of the biggest decisions the company has made is that its data center doesn't have to live at the New Orleans headquarters. Gustav showed there are still some issues with failover, so he wants to eliminate that.

Explains Chaffe: "We've decided we're probably going to turn Dallas into the main site. We're using a Verizon Business Center in Richardson and I'll probably move our data center here to a similar facility in Houston. I need one of these things in a city where I have personnel. It makes much more economic sense for us to use the center in Houston and ride Verizon's fiber optics to Dallas and replicate it that way. It's so much cheaper that way. In Dallas, we don't have to fail anything over."

See also Iowa Hospital's Disaster Plan Helps Keep Its Head Above Water and When All Else Fails, Make Sure You Have Backup.

V&D PLUS

DISASTER RECOVERY: Telecom: Down, when most needed

For telecom service providers, it boiled down to the number of cell sites, and whether proper redundancy planning had been made

Rajneesh De

Wednesday, September 07, 2005

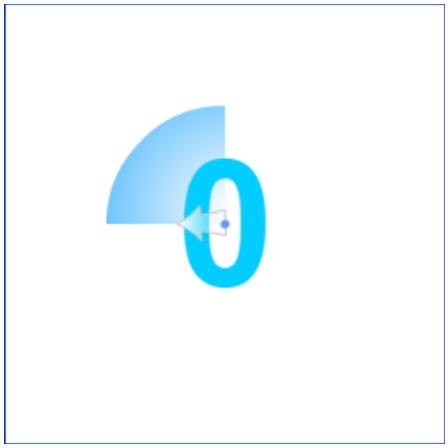

The Mumbai deluge of July 26 exposed the softer underbelly of India's financial capital. True, at 940 mm it was the highest recorded rainfall in a single day anywhere in the world (forcing geography books to replace Cherrapunji with Mumbai) but still, the way most basic services like power, telephone network, and the bank ATMs floundered showed how unprepared India is in the face of a disaster-natural or otherwise.

As the city gradually limps back to normalcy amidst lots of mutual finger pointing, Dataquest decided to have a look at how technology fared on the black Tuesday, and whether it could have really assuaged the problems to a certain extent. Our investigations reveal one interesting result: in most cases the problems were rather man-made, or manual, in nature; technology, more or less, proved to be a savior. However, as all service providers analyze their 7/26 (July 26) learnings, the biggest lesson seems to be: have a proper planning procedure in place to better harness technology during similar disasters in the future.

Drowned Networks

The sudden failure of telephone networks was perhaps the worst thing to have happened. There seemed to be a common problem cutting across all service providers-most of the cell sites, especially in the suburban areas which received more rainfall, had no power supply. Though generators were available, in very few cases were there provisions to stock enough diesel. This lack of foresight came to light since many suburbs like Andheri, Marol, and Vakola had no power for days on end and consequently telephone connections in these areas suffered, even after the 26th.

Tata Teleservices, in particular, seems to have suffered the most. Admits Sandeep Mathur, "Since most of our cell sites run on Reliance Energy power supply-which was down for a week in the Andheri region-TTSL cell sites were down as we had generator capacity for only four hours, after which it was difficult to even procure diesel." What made matters worse was that TTSL's equipment in the Marol site, which delivers services to many clients in the Andheri area, was completely submerged on both 26th and 27th. New equipment, as replacement, could be brought only after the water level receded a bit on the 27th. However, on the 28th it rained again, and the switching equipment in Marol got affected.

While Tata's phone services were badly hit, Reliance Infocomm fared better. The general consensus was that the Reliance network was working the best amongst all. SP Shukla, Reliance Infocomm, however, does not attribute this to a CDMA advantage. Rather, he feels that having more cell sites in more areas helped. He also adds that the extensive network of optic fiber cables also proved to be a boon-rain water did not have any impact on these. Also, Reliance seemed to have a better provision for generators, and adequate amount of diesel to run them.

Where did the station go? High tide at Bombivli station

This proves that it was not a question of a CDMA or a GSM service provider offering better service on that night. It boiled down to the number of cell sites, and whether proper redundancy planning had been made for power and diesel generators. Claims Jayant Khosla, CEO, Bharti Televentures, "The rains had minimum impact on Airtel transformers and base stations, and even in power problem areas generators made the sites fully functional. Apart from that our current battery backup is a minimum of 12 hours and diesel generators were deployed at cell sites at strategic locations." Amongst other GSM providers, Orange customers suffered as they had power backup of only two hours on cell sites. Similarly, there was not enough power back up on BPL and MTNL cell sites.

Sinking Banks

After phones, another piquant problem Mumbaikars faced was non-functional ATMs. SBI, HDFC, ICICI, UTI, all banks with large number of ATMs were not working. In fact, in case of SBI, even in Kolkata a notice was hung on one of the ATMs, on 27th, that since the ATM switch in Belapur, Navi Mumbai was under water, it could not function properly. Ramanthan of SBI, however, denies this as "sensationalism of media" and claims that their operations were fully in place, thanks to migration to the DR site in Chennai. "On 26th night, when we realized the gravity of the situation, we switched to Chennai DR for the ATM networks, as they are critical, and cater to 5400 ATMs across the country. We also have 10 other bank customers sharing this network and could not afford any down time. We decided that more than any other services our foreign operations and ATM services had to be operational and, therefore, they were made available through our DR site at Chennai. Later, we were able to switch back to our DC at Belapur for our ATM Networks."

The few instances, he claims, where the ATMs in suburban Mumbai were down were due to last premise equipment, where there was physical disruption of the ATM centers like water logging, cable cuts and equipment damage. He also claims there was never any cash crunch and no problem with cash replenishment. "The DR site in Chennai is a mirror site and all the log files were switched; however, we continued to run our core banking apps from the Belapur DC only," he added. MTNL in Navi Mumbai was flooded and so communication was disrupted on the 27th, Ramanathan admitted, perhaps making provisions for why there was disruption in other parts of the country.

Other banks too claim that it was not that their ATM networks were down, rather it was the physical ATM centers in severely water-logged areas that were non-operational. Even VK Ramani of UTI Bank informs that ATMs in the suburbs that were down were so due to cable cuts and water logging. "In our case the ATMs went down for a brief period of four hours because there was no power in our Chembur Datacenter. We considered migrating to our secondary site at Bangalore, but realized that no DR of a bank can give 100% service all the time. And since the ATM could not be down for a brief period also, we decided to continue on the primary site only and that way had 100% service available. Since there was no power from Reliance Energy, we managed our own power supply and when we felt that we might have to restock on fuel, the drums were 'rolled over' to the premises as there was no transportation."

Ramani's argument that DR cannot run 100% services of a bank are echoed by Srinivas of Sanovi too, a DR service provider of Mumbai who has HDFC as its customer. "A remote site cannot offer guarantee as the issue here is to ensure consistent formats between the primary and the secondary sites. There is data loss during failovers.

Only when the failover is of the organized type, which can happen with the help of designed software, it is smooth and 100% data is guaranteed, otherwise a secondary site is prepared to work for durations starting from 7 days to months."

HDFC Bank was a typical example. It had to revoke its DR for its ATM Network. Though the datacenter was operational, but end ATM network centers' functioning was disrupted at various locations owing to the data loss during failover. The bank took a dual approach where it ran its critical apps from the mirror site and continued to run the non-critical operations from the primary site. The learning for most banks seems to be that when there is a communication failure one has to depend on automation, as minimum people involvement is advisable. The system behaves like a split-brain-the software views all the services and starts treating the secondary functions as the primary.

Datacenters-Live Up To Expectations

Mumbai also houses a number of important datacenter hosting sites of crucial financial services and other organizations. A look back at how Reliance, Tatas and the NSE coped with their datacenters could be a useful future reference for service providers encountering similar disasters. The biggest advantage for Reliance, argues Sunil Gupta, is that "the basic design of the datacenter is owned by us. Hence, quality of data control is in our control; it runs on a commercial building, leased out, but owned independently by

us." Another advantage was that they had a totally independent power feeder for the building in 2 separate grids, with power redundancy also from 2 feeders. And even if power from the grid went down, they had enormous power capacity to generate their own power, he added.

Moreover, the Reliance datacenter buildings have Faraday's cage, implying they are guarded against lightening strikes, are electromagnetic radiation safe, and designed to absorb an earthquake measuring 8.3 on the Richter scale. Therefore, all 10 plus banks hosting core banking with Reliance plus other enterprises like HLL, Godrej and Videocon had no operations disrupted. There was also the mirror site in Bangalore that ensured that all data was replicated in Bangalore if a customer opted to go for a secondary site in Bangalore for DR.

Business continuity was not an issue for Satish Naralkar of NSE as they had an extremely robust DR plan in place: " It was so used to mock drill, that when tested in a live scenario it more than lived up to expectations." Power was not a problem as NSE had supply from both Tata and Reliance and also a strong backup in place for fuel capacity to cater for 72 hrs. The main datacenter equipment was housed on the 6th floor but few were in the basement, which had to be protected. Pumps were used to dislodge water, and blocks also put in place. NSE had prepared back up plans in case the adversity continued beyond 72 hours.

Though SEBI had declared only the 27th as a holiday, the market was ready to trade only days after. And as people were stuck in their offices they wanted their systems to be available and be able to trade. The NSE mirror site is in Chennai. The primary site, inspite of being in an area which was in the high alert affected region, Bandra Kurla Complex, was well protected. The satellite connectivity was on. There were instances of only a few LAN connections which were down as they were on leased lines from MTNL-the MTNL exchange in the area was down, Narelkar informs.

For the Tatas, just like their voice services, the disaster provided valuable learnings. The power had to be shut off as there was danger of short circuit in the water, with most of the equipment located in the basement. Emergency power on generators was used, but as there was capacity in place only for 4 hours, they ran out of fuel. Tata Power proved to deliver efficiently by putting in a second ring. The only DC affected was the one at Andheri (E) Technopolis Park. The Vashi and Prabhadevi (VSNL premises) were unaffected as they are located higher up in the building. By 26th late evening, Mathur informs, power at the DC was restored with the help of the second ring.

Rajneesh De *and Minu Sirsalewala*

Page(s) 1

PREVIOUS ARTICLES

- Data Services: Sizzlers for you
- Digital Brodcasting: Entertainment goes mobile

WLAN Security: No Entry At Will

VOICE&DATA ZONES

BPOrbit New | Service Provider | Equipment Business | Technology | Column | Regulatory | Enterprise | Interviews

For Voice&Data Print Subscription

[Magazine Subscription] [Contact Info] [Advertise : Online | Magazine | Advertising Print]

Other CyberMedia web sites

[Dataquest] [PCQuest] [CIOL] [Living Digital] [IDC India]
[DQ Channels] [The DQweek] [CyberMedia careers]
[CyberMedia Events] [CyberMedia Digital] [Cyber Astro] [CyberMedia India]
[Global Services] [BioSpectrum] [BioSpectrum Asia]
[Computer Shopper] [College Buying Guide] [Voice&DataConnect]

THE CHRONICLE OF HIGHER EDUCATION
Information Technology

From the issue dated February 28, 2003

Preparing for Computer Disasters

Advance planning and backup systems allow colleges to recover data

By DAN CARNEVALE

Seattle

Marc F. Elvy prepares for disasters, but he wasn't prepared for what he saw when he showed up at work one day in December.

The University of Washington's Educational Outreach office, where Mr. Elvy is the network manager, was on fire. For eight hours the top floor of the two-story building burned while the fire department trained hoses on the flames.

After the fire was extinguished, another seven hours passed before Mr. Elvy could persuade the fire marshals to let him into the building. He found just about everything on the second floor, including the computers, either melted or burned to a crisp. A foot of water flooded the first floor, where the office's 17 main computer servers were housed.

"The whole ceiling came down on the first floor," Mr. Elvy says. "It was just a soggy, pasty mess."

Mr. Elvy fished some of the waterlogged servers out and loaded them into his pickup. He also found some software disks floating in a briefcase in the building. Then he drove home and spent the next several days working with a space heater and his wife's hair dryer.

He was able to save some of the e-mail data the servers contained. But most of the 17 machines were useless, and many of them held student records and registration information for the university's online courses.

Fortunately, the university saves all important computer data every night on tapes that it keeps off the campus. The university also has a supply of spare servers for emergencies. Mr. Elvy and his staff members had the educational-outreach network and its online services back up and running after just one business day, although at a slower-than-normal pace. It took two weeks of work, over the Christmas holidays, to bring everything up to full speed.

But for every story of how an institution was able to save its data, a professor somewhere has a horror story about losing years of research after a fire, natural disaster, or a server crash. And while most large universities religiously back up their main computer servers every night, some only save data on smaller servers once a week. And faculty members may or may not back up information they save on their own machines. Such information, sometimes including grades and research, is often at risk of being lost.

Most colleges have some sort of plan to protect their computer information, although few have faced the kind of disaster that would demonstrate whether those plans actually worked. Experts say staging a mock computer disaster can highlight a preparedness plan's shortcomings, but such tests happen only rarely.

Some smaller institutions, meanwhile, can't afford extensive backup systems, which can cost hundreds of thousands of dollars. Because of the expense, some institutions don't buy backup computer systems at all, or do much to protect the data in their servers.

Backing up computer data for critical functions, like payroll and personnel files, is a common business practice. The attacks on September 11, 2001, and the constant threat of terrorism have helped push the issue of disaster preparedness.

A faulty fluorescent light caused the fire at the University of Washington's Educational Outreach office. The building, which the university rented from a private landlord, was old enough that a grandfather clause exempted it from automatic-sprinkler requirements. Luckily the building was unoccupied at the time of the fire.

The blaze caused about $1-million in damage. Insurance will cover most of the expense, including the $170,000 cost of replacing the computers.

David P. Szatmary, the university's vice provost for educational outreach, says he had never paid much attention to the backup systems until the day of the fire. "The first fear when I saw the flames was, 'I hope that our data was backed up,'" he says.

Even though almost all the information was saved, staff members discovered that "working from backup tapes isn't the easiest thing in the world," Mr. Szatmary says. The staff worked many late nights transferring information to the backup servers, which was a challenge in its own right: A system that had worked on 17 servers now had to be handled by 3, meaning that response times were slower for users.

Mr. Elvy says the cost of having the spare servers ready to use was $5,000 each, and that tapes for the regular backups cost $14 apiece (staff members handled the chore as part of their regular duties). When all was said and done, it was money well spent, he says. "As far as my position goes, this is what I plan for, even though I hope it doesn't happen," Mr. Elvy says.

Fire Claims

The fire at the outreach office wasn't the first disaster to strike the University of Washington. In May 2001, the Center for Urban Horticulture was set afire. No one was arrested for the act, but a radical environmental group called the Earth Liberation Front took credit for the blaze. Members of the group claimed in a statement that they were trying to destroy genetic research being conducted in the center's labs.

In that fire, many faculty members lost their research materials, some of it on hard drives and some on paper or in artifacts. While the professors who work at the horticulture center can back up all their data on the university's servers, many hadn't taken the time to do it. Others had backed up their data on disks but left them in their desks, which burned in the fire.

Thomas M. Hinckley, director of the Center for Urban Horticulture, had important information on his hard drive, including historical data on Mount Saint Helens and digitized images of other research. The university paid a company $2,500 to recover as much of the information as possible, but about a quarter of it was beyond recovery. It has taken Mr. Hinckley a while to figure out just what data he has lost. "It's one of those things -- as time goes by you look for old files and you can't find them," he says.

Other faculty and staff members' computers required similar attention to save data that were not backed up. "It's lazy -- you're working up to the last minute of the day, and instead of taking 10 minutes to save things, you just assume everything will be OK tomorrow," Mr. Hinckley says.

The one person who didn't lose any data was Toby Bradshaw, a professor in the departments of botany and zoology whose work was the target of the attack. The vandals set fire to a five-gallon drum of gasoline in his office to destroy his research on tree genetics. But he had backup copies of everything off-site.

Learning From Experience

Observers say that most universities are doing at least the minimum necessary to protect computer data, although college officials tend to think more along the lines of server crashes than widespread disasters. And while many institutions have plans that protect data, few have well-thought-out systems for recovery once a disaster has happened.

Jon W. Toigo, chief executive officer of the consulting and analysis company Toigo Partners International, says data recovery often isn't a big priority for an institution unless people there have already witnessed some major catastrophe.

"That's usually the motivator, that they've experienced a disaster, and they're trying to prevent it from happening again in the future," Mr. Toigo says. "You really have to think about recoverability before you build the system."

At the very least, he says, colleges should have some sort of procedure for automatically saving the data. "As long you've got the data backed up, you can buy another system from CompUSA," Mr. Toigo says.

Cole Emerson, chairman of the Disaster Recovery Institute International, which teaches organizations how to protect data, says every business and institution should do at least three things to protect data. First, he says, back up the data. Second, store the backed-up information off-site. Third, have contact information available for people who can help retrieve the information at a moment's notice.

While those are the minimum recommendations, institutions would be wise to also have extra servers on hand so data can be retrieved immediately and services can be restored quickly. And officials should have extra copies of software available to retrieve the stored data.

A Daunting Task

Just 90 miles north of the University of Washington is Western Washington University, which had its own scare in July when the room containing the computer nerve center for the College of Business and Economics was destroyed by fire. The data were backed up and copies were kept off campus. But the officials found the task of putting the pieces back together more daunting than they had expected.

Dennis R. Murphy, dean of the college, says the data were backed up in preparation for a server crash, not necessarily a widespread disaster. "Fire was rather far from our radar screen because this building is brick, so there's not much to burn," he says.

The university had spent over $100,000 on backup servers, and that saved both time and money, Mr. Murphy says. But he says officials had never tested their ability to retrieve backed-up data using other people's equipment. It took a couple of weeks to recover all the information, and the university had to spend a couple of hundred dollars for each computer that needed to be decontaminated. The cleaning bill alone was in the thousands.

Student records and other critical data were saved. But some faculty members lost research data. "There's a lot more individual backing up going on now," Mr. Murphy says.

The computer room in the building has been rebuilt, and now it has a smoke detector installed. And business-college officials have assembled a CD-ROM that contains emergency information, such as a vendor contact list, student-identification information, and a record of the last payroll run. Each administrator gets a copy of the CD, which is updated monthly, so he or she can dig up important information during the first 24 hours after a disaster.

"A lot of places take care of the big things," says Jerry Boles, Western Washington's vice provost for information and telecommunications services. "But it's these little things that can be very frustrating."

The only way to identify those little things is by doing practice drills, Mr. Murphy says. "You need to do dry runs," he says. "It's a pain in the neck, but it's the only way to find out."

Paul Ellis, program director for IBM Tivoli storage management, says drills should be conducted a couple of times a year. Tivoli is a branch of IBM that sells services and equipment for protecting computer data to businesses and

institutions.

During those drills, he says, officials will be able to tell whether the stored data are easily accessible and compatible with other equipment. Drills also make people realize that they need spare copies of the software to run the data they've been storing.

"The scary part is if you're planning for this disaster, and it fails in the calm moment, what's going to happen in a real disaster?" Mr. Ellis says. "Everybody talks about backup, but the really important part is how long it takes to do a restore."

Although larger universities find it is worth their while to develop elaborate backup systems, community colleges don't always have the cash on hand to make that sort of investment.

John R. Moore, associate dean of computing services at Allegany College of Maryland, says the two-year college can't afford to put much money into backup systems. So officials there make do with what they can. "We looked into buying disaster-recovery capability, and it was rather expensive," he says.

Instead of using an automated backup system, staff members manually save information on tapes. Instead of keeping backed-up information off site every night, staff members put the tapes in a fireproof vault in the same building as the computer servers.

Once a week, the tapes are taken to an off-campus location. While the fireproof vault offers some protection, the heat from a fire could make the data on the sensitive computer tapes unreadable. And if the building collapses, it could be weeks before the data could be retrieved. But that's the best that Allegany, a college with 3,000 students, can do, Mr. Moore says.

Spending Varies Widely

Allegany has not had any problems so far with its system. But if a fire did knock out the computer servers, Mr. Moore says it would take two to four weeks to acquire new equipment. Classes could continue, he says, unless they were conducted online.

Not all community colleges have to scrimp when it comes to creating a backup system. The Community College of Baltimore County has off-site backup systems, and the college hasn't lost any data despite two incidents in which networks were damaged by water leaks, says Wally Knapp, director of information technology and technical services there. He says he doesn't know how much the college spends on backing up the data, but it's well worth it if it means a good night's sleep for everyone involved.

Other institutions have had close calls that have compelled them to take broader precautions. In 1998, the University of Missouri at Columbia was struck by a minor tornado that tore the roofs off some campus buildings. No data were lost in the incident. But Willie Jones, records analyst at the university, realized that its computer records weren't protected well enough.

So the university created a mirror site, complete with computers and duplication of most crucial data, five miles away. The university would not reveal the cost of the operation for security reasons.

Some universities save money by finding ways to make their backup systems productive long before disaster strikes. Instead of having backup servers sitting around doing nothing while they wait for a disaster, Northwestern University uses the equipment for low-priority tasks.

"They may be e-mail servers in a normal situation, but they may end up doing payroll if that's what we need" in a crisis, says Mort Rahimi, vice president and chief technology officer at Northwestern.

The university invested $300,000 in backup servers and spends

about $5,000 a month storing data on backup tapes. Critical information is saved every night while other information, such as research data, is backed up weekly. It's up to faculty members to decide whether to save their data on university servers instead of on their desktop machines.

"The big risk, of course, is if somebody has been working on something for years. If a disaster strikes, they'd lose that data," Mr. Rahimi says.

As Mr. Hinckley at the University of Washington can attest, it's worth the extra effort to save research information on the university servers. "My view is, you back everything up," he says. If there is a fire, you don't lose anything, he says, and "you get a new computer out of it."

http://chronicle.com
Section: Information Technology
Volume 49, Issue 25, Page A33

Local Tech Wire

Posted Aug. 11, 2010 at 2:05 p.m.

'Murphy' goes online: Your disaster recovery plan is the target

Kip Turco
By KIP TURCO, Special to LTW

Editor's note: Kip Turco is chief operating officer at Hosted Solutions, a provider of managed hosting and data center services such as "cloud computing."

RALEIGH, N.C. - I know you're familiar with Murphy's Law. When I was in the? Army, we were always concerned about "ol' Murphy." But please rest assured, the folks at American Eagle Outfitters now know more about Murphy than they ever wanted, and what they know from first-hand experience provides a valuable lesson for the rest of us.

American Eagle's e-commerce website was down for eight days, according to StorefrontBacktalk.com, in what the editor called "complete website death," while the company struggled to recover data from crashed servers.

It wasn't like American Eagle was trying to cut corners, mind you. According to StorefrontBacktalk.com, it had outsourced much of its web operations to IBM. And it was using IBM and Oracle software and hardware. But the company apparently suffered through a "perfect storm." Up to a point.

American Eagle's storage drive went down, followed shortly by the secondary backup drive. The Oracle backup utility software worked, kinda: StorefrontBacktalk.com quotes a source, who says the software was restoring one gigabyte per hour. Which is okay, I guess, but they needed to restore 400GB.

The problem was American Eagle's disaster recovery site wasn't ready to go. The source is quoted as saying, "They apparently could not get the active logs rolling in the disaster recovery site. I know they were supposed to have completed it with Oracle Data Guard, but apparently it must have fallen off the priority list in the past few months and it was not there when needed."

The lesson is painfully obvious. Always assume the worst will happen … and this is exactly why "Murphy" was a popular term used among soldiers in the Army. Work with your team, either in-house or outsourced, to figure out how you'll be able to get back in business when it happens. Above all, test…test…and test again.

If American Eagle had done this, it would have discovered the DR site wasn't up to speed. That lack of preparation apparently cost American Eagle big time.

One other thing; it's a commonly held precept in our business that a DR plan that sits on the shelf is potentially worse than no DR plan at all, because it instills a false sense of security. Things constantly change; servers are added or subtracted, data load grows. Keep testing. If you're outsourcing your DR and storage to a third party, make sure that regular testing is included in any SLA.

So, American Eagle Outfitters is back up and running, and that's a good thing. They have also learned, unfortunately, "If anything can go wrong, it will, at the worst possible moment."

Get the latest news alerts: Follow LTW at Twitter.

Tags: Hosted Solutions, Opinion Analysis

Editor's Blog

The latest blog posts from our Local Tech Wire and WRAL editors.

July 30, 2010

Down For 8 Days: American Eagle's Site Disaster

Retail Realities: One of the Longest Site Outages Ever Hit a Multi-Billion-Dollar Retailer

(CBS / iStock Photo)

(CBS) This column was written by Evan Schuman, the editor of StorefrontBacktalk, a site that tracks retail technology, e-commerce and security issues. Retail Realities appears every Friday. Evan can be reached at E-mail and on Twitter.

In one of the longest site outages ever for a multi-billion-dollar retailer, Tuesday (July 27) saw the apparent end of more than a week of Web problems and days of an outright crashed site for Pittsburgh-based clothing chain American Eagle Outfitters, which outsources much of its Web operations to IBM. The site crashed last Monday (July 19) and stayed dark until Friday (July 23), when it limped along with various parts not functioning until Tuesday afternoon (July 27).

The site's problems, though, shed light on an interesting strategy. During the many days of complete Web site death, the $2.7 billion apparel chain's mobile site was apparently fully functional. (To be precise, we hit the mobile site during the Web outage and it appeared to work fine. But we didn't have the chance to test it by making a purchase. And officials at American Eagle Outfitters, IBM and Usablenet--which handles the chain's mobile site--wouldn't comment on the mobile site's functionality during the crash.)

This availability raises the question: Should retailers look to their mobile sites as emergency backups for their Web sites? Should pages indicating that a site is down automatically include a link to the site's mobile version? Mobile sites, of course, work just as well on desktop machines as they do on phones. American Eagle Outfitters, which has the admirably short URL of ae.com, exists as a mobile site at http://mobile.usablenet.com/mt/www.ae.com/web/index.jsp.
Before we dive into that mobile-as-site-backup issue, let's look at exactly what happened with American Eagle's site. None of the players involved would get specific as to what was wrong with the site, other than to say that there was no upgrade going on at the time and that the site experienced "a hardware issue."

A server failure almost certainly would not have caused this problem; redundant servers would likely have kicked in while the defective machine was replaced with a new server and a backup was restored. That process would have taken a few hours, not almost eight days.

This delay suggests some sort of storage problem. Say the storage array begins to fail. OK, no problem, we'll just find the bad drive and replace it. Whoops, looks like something has corrupted multiple drives. (That could happen if power gets flaky inside the array.) Now we have a catastrophic failure of the storage array. No problem, we'll just fix the hardware and restore.

Whoops, new problem: Turns out this problem has been going on for a while. The last set of backups is corrupted. So is the set of backups before that. Sorting through to reconstruct good data is going to take time.

Alternatively: All recent backup sets are toast. Maybe nobody was verifying that the data was actually being written. However, all the transactions are being logged. No problem, then: All it takes is a lot of time and special expertise to essentially rerun all the recent transactions (since the last good backup) into an empty database, merge the new stuff with the old stuff and then load it all back into the replacement hardware.

By the way, it seems that American Eagle was recently searching for a "Manager - Business Continuity

(Given that the site wouldn't be fully back up for eight days, the site must use the same definition for "a few minutes" that my 12-year-old daughter uses when she offers a time estimate for when she'll wash the dishes.)
That sign was up through Wednesday (July 21). On Thursday (July 22), the sign changed to "We're making updates to our sites. Free Shipping on us when we're back, thru July 25." Note: The site didn't fully return until July 27.
Back to the timeline, the site came back up late Thursday morning (July 22): Strand put the time at "about 11:15 AM" and added that the site at that time had "a little bit of limited functionality. Some of the saved data hasn't been restored, but it is shoppable now."

By Friday morning (July 23), the site's message had changed again, offering some details on the limited functionality. "We're still working through some issues, but you're able to shop! Everything should be completely fixed very soon. Thanks for hanging in there. Stuff we're still working on: Order tracking, Registered Information Functionality, Wish List, Order History."
Screens through Tuesday (July 27) thanked consumers for "hanging in there while we work through some site issues." On Tuesday afternoon, the warning messages came down and IBM said the site "is now fully operational."
Although this outage was much worse than retailers typically suffer through, even a several-hour outage poses the potential for losing customers. As such, does it make sense to look at your mobile site as an emergency backup to your Web site?

A few issues need to be considered before you make that leap: Many mobile sites leverage the content and the database of the main site. That would mean a mobile site would only be helpful if the main databases of the site--pricing, inventory, order placement, etc.--are still functioning. Payment is often independent, so there's a fine chance that system may survive. If it's merely a hosting server that has crashed, the mobile site could be a powerful backup.

As such, would it be a good idea to run a mobile site on duplicate mirrored versions of the key databases, explicitly so that that site could more likely survive a main Web site outage?

Even if it does survive, there are other concerns. Despite ongoing signs that mobile is growing with staggering speed, some mobile sites are not given the infrastructure to support enterprise-level Web traffic. Unless the sites are redesigned to handle much more bandwidth, pointing your Web traffic to your existing mobile site may cause a mobile crash.

Another concern, albeit much less significant, is design. Mobile sites are deliberately kept utilitarian. That look may not appeal to some marketers, who could argue that "no site" is better than a "bland site." Perhaps, but we're guessing the CFO may disagree.

The opinions expressed in this commentary are solely those of the author.

By Evan Schuman

More From SciTech

- Intel To Buy TI Cable Modem Chip Business
- Cell Phones: Easy ID Theft Targets
- Dell To Buy 3Par For $1.13 Billion
- Spacewalking Astronauts Plug In New Cooling Pump
-  Play CBS Video

  America's Quest for Alternative Energy

From: www.cio.com

Inside American Eagle Outfitter's 8-Day Website Nightmare

– Thomas Wailgum , CIO

August 09, 2010

HANG IN THERE. That was the first plea American Eagle Outfitters (AEO) offered its nearly 1.5 million Facebook friends three days after its trio of e-commerce sites had crashed.

The online retail outlets—ae.com, aerie.com and 77kids.com, which target teens and kids—suffered what American Eagle Outfitters (AEO) termed a "major incident" on July 19. "This has caused all of our sites to go down," a Facebook update said on July 22, at 1:26 p.m. "We're working together 24/7 to get them up and running as soon as possible. Hang in there....we'll make it up to you. Check back soon."

Those back-to-school shoppers who checked back too soon at the sites received this message: "We're making updates to our sites." This was no planned maintenance, however. Behind the scenes, AEO and its external hosting provider IBM (IBM) were scrambling for several days as the sites stayed dark.

The reasons for the four-day outage (and subsequent four more days of technical aftershocks) related to backup and disaster-recovery technologies and processes that failed, according to several articles on retail tech site StorefrontBacktalk.com.

As the articles describe, atypical and concurrent failures with IBM's hosting servers and backup plans as well as with Oracle's (ORCL) Data Guard utility program ultimately proved to be the sources of problems. (Stephanie Oschwald, a spokesperson for Pittsburgh-based AEO, did not return several messages seeking comment. IBM and Oracle declined to comment for this article.)

[Read about e-tailer Zappos.com's $1.6 million pricing mistake]

On July 23, at 5:39 p.m., AEO announced on Facebook that the sites were A-OK. "We had a few glitches but we're back! We missed you while we were gone. Free shipping now thru Monday, July 26 11:59 PST. Automatic on your entire purchase (US and Canada only)."

However, noted StorefrontBacktalk.com, four more days of technological healing lie in wait for the AEO site, "when it limped along with various parts not functioning until Tuesday afternoon, July 27th."

An Unprecedented Outage

The shock of a 100-hour e-commerce outage now over and the damage done to their critical back-to-school sales period, executives of the $3 billion AEO have to be left scratching their heads in frustration—and wondering just how much "teeth" they had written into their service-level agreements (SLAs) with IBM. (For the record, StorefrontBacktalk.com reported that AEO had recently filled the position of "Manager—Business Continuity & Disaster Recovery," and not a moment too soon.)

AEO's websites weren't the first to go down unplanned, nor will they be the last. Dating back to the mid-1990s, history is littered with tales of site outages, complex IT fixes and lost revenue—embarrassing for the companies and career killers for those in charge. Any Amazon.com unplanned downtime, for instance, is still reported with as much zeal as a Jennifer Aniston "new boyfriend" sighting.

But the duration of AEO's site outages could be considered unprecedented for an e-retailing brand.

"In my career, I can't remember a mission-critical application that was down for four days," says Paula Rosenblum, an analyst at RSR Research and former retail CIO. "It's almost impossible to conceive—and I say that not as an analyst but as a former CIO."

Deals, Discounts and Disappointment

The timing couldn't have been worse for AEO, at the outset of its "back to school" summer push. In its stores, AEO's July performance was a "disappointment," according to a Zacks Investment Research report on retailers. The retailer had also lowered its quarterly earnings outlook in July.

The company not only lost online revenue during the outage, but also attempted to make up for the downtime by offering its customers deep promotions—such as this deal promoted on its Facebook page on July 30: "Thanks for hanging in there while our site was down...this Friday & Saturday (7/30 & 7/31) online only take 25% off your entire order & we'll throw in free shipping, too!... Thanks again for hanging in there!"

Brian Walker, a Forrester Research principal analyst who follows e-commerce, doesn't term the AEO outages unprecedented, but unusual. "Sites routinely will become slow and the customer experience and conversion rate may be affected," he says via e-mail, "but to be down for four days has a significant business and customer impact."

AEO's "direct businesses" (those three websites) rung up more than $340 million in sales in fiscal 2009, which accounted for approximately 12 percent of the retailer's overall sales.

Outsourcing E-Commerce Ops: Deal or No Deal?

Outsourcing e-commerce operations in retail is not a new proposition—and IBM and Oracle, for that matter, are not fly-by-night operations. "Using managed services or SaaS solutions for all or key parts of the e-commerce platform is routine today," says Forrester's Walker. "On average, eight hosts will be involved in supporting a typical e-commerce transaction."

Even with the site outages, the outsourcing proposition is still sound, Walker says. "There is no guarantee that doing this in-house is any better," he opines, "and in many ways can be much harder to do as resources, budgets and skillsets may be harder to have at the right levels."

However, it's important that a company outsourcing its operations not abdicate every decision and responsibility over to the hosting company. "Often in these situations there will be some shared responsibility between the business and IT-support staffs," Walker says. Trouble happens, he says, when something as simple as a failure to communicate an upcoming big promotion or event that will likely result in a spike in online traffic big enough to take down a website.

Interestingly, during the sites' outages, AEO's mobile site was still functional, according to StorefrontBacktalk.com. RSR Research's Rosenblum says that's a "good news / bad news" situation. The good news: The mobile site was up, and AEO could have done a customer redirect from the online site that was down to the mobile site, if they chose. The bad news: It would appear that AEO is running separate, siloed systems (for customer databases and queries, for instance) for what are, essentially, similar e-commerce channels, Rosenblum says.

"Our research shows us that too many retailers are in a similar situation," she says. "They spend an inordinate amount of scarce resources—IT and human—forcing multiple channels into synch. The opportunity costs associated with these activities

are incalculable."

Forrester's Walker advises other retailers to use AEO's outage as a wake-up call. "Get your responsibilities, processes and systems together now, test them, and correct as needed," he says. "Having a clear plan in and of itself will pay huge dividends in the event something happens, and with the technology and architectures we have available today this can be done at a reasonable cost."

"Consider the downside of having the business offline for a few days during the holidays," he adds, "and the ROI is pretty clear."

Do you Tweet? Follow me on Twitter @twailgum. Follow everything from CIO.com on Twitter @CIOonline.