

More Never Agains IV

February 2010

It is once again time to reflect on the damage that IT systems can inflict on us mere humans. We have come a long way in ensuring the high availability of our data-processing systems. But as the following stories show, we still have a ways to go. During the last six months, hardware/software and network faults shared responsibility, each causing about one-third of the outages. The rest of the outages were caused by a variety of problems such as power failures, construction mishaps, and hacking.

Rackspace Hit with Another Outage

Techcrunch, June 20, 2009 – On June 20, Rackspace suffered yet another outage¹ due to a power failure. The breaker on the primary utility feed powering one of its nine data centers tripped, causing data center's generators to start up. However, a field excitation failure escalated to the point that the generators became overloaded. An attempt by Rackspace to fail over to its secondary utility feed failed because the transfer switch malfunctioned. When the data center's batteries ran out, the data center went down. Failovers do fail. Have a contingency plan no matter the extent of your redundancy.

NYSE Suffers Several Outages in Less Than a Month

Reuters, July 2, 2009 – On Thursday morning, July 2, brokers on the floor of the New York Stock Exchange found that they could not route orders, causing the NYSE to halt trading in some stocks and to extend the trading day. During the previous month, a software glitch halted trading; and an order-matching problem affected timely order reconciliation. The Exchange has now implemented a new order-processing system that reduces order-execution time from about 100 milliseconds to five milliseconds. However, brokers want microsecond order-processing times.

Comcast Outage Kills Internet for Maryland and Delaware

Ocean City Today, July 3, 2009 – A concrete mixing truck accidentally hit a power pole, cutting Comcast's fiber line. The accident took out Internet, telephone, and video services for a large area of the lower Maryland-Delaware area for several hours. Even after the damage was repaired and after service was restored, multiple temporary outages continued for another several hours. During the outages, merchants could not accept credit cards; and the local newspaper, the Ocean City Today, could not get the emails necessary for its weekly issue.

¹ Rackspace – Another Hosting Service Bites the Dust, *Availability Digest*, December 2007.

Google Apps Down for Hours

Softpedia, July 3, 2009 – Access to the Google Apps Datastore went down for the morning on July 3. Google Apps is Google's answer to cloud-based development and hosting. The outage affected the entire Google Apps cluster. The problem started when the Datastore went into read-only mode with all writes disabled, but the problem quickly escalated to a total failure. All Google Apps applications were down for over four hours. There has been no explanation as to the cause of the problem.

MySQL.com Goes Out Just Before an Upgrade

The Register, July 22, 2009 – A power outage in Uppsala, Sweden, took down the MySQL web site for almost a day. The outage prevented the open community's developers from accessing MySQL source code, documentation, discussion groups, and bug tracking. Interestingly, Sun Microsystems, which owns MySQL, had planned to upgrade the data center that weekend specifically to avoid an outage due to power failures. Bad timing. Many are waiting to see if Oracle is going to continue to support MySQL following its acquisition of Sun.

Undersea Cable Fault Takes Out Internet throughout West Africa

Africa News, July 30, 2009 – There is only one fiber-optic cable linking West Africa to the rest of the world; and on July 30, the landing cable linking Nigeria failed. 70% of Nigerian Internet services went down, and other neighboring countries lost 100%. Though Nigeria had built a redundant land network, there was no backup for the landing cable. The only way to reroute traffic was over limited, expensive satellite channels. Banking services, phone services, and Internet services were all impacted. It took days to repair the cable and to return service to these countries.

Cisco Takes Down its Own Network – Twice

The Register, August 5, 2009 – A human error was responsible for knocking Cisco's web site offline for three hours. Evidently, a maintenance update caused an "electrical overload." System administrators were unable to access Cisco's support resources just after Cisco announced four security vulnerabilities in its network products. This wasn't the first time that Cisco stepped on its own toes. In August, 2007, a preventative maintenance error knocked out the Cisco network for another three hours. In neither case did the backup facilities kick in.

iPhones Down in the UK for the Fourth Time due to O2 Outages

The Register, August 5, 2009 – O2 is the UK network that exclusively provides mobile services for the Apple iPhone. In July and August, the network went down on four different occasions, leaving iPhone users without email and Internet services for hours at a time. O2's problems ranged from a fire in a London substation to a 2G network GPRS (ground packet radio service) satellite issue, a DHCP (dynamic host configuration protocol) problem, and the failure of the system that assigns unique IP addresses to users when they sign on.

Even MIT Can't Keep Its Email Up

The Tech, August 5, 2009 – Over 10,000 students and teachers were without email for over a day when both controllers in the Storage Area Network that manages the email disk drives failed. The aging email system was slowly being migrated to an Exchange Server environment when the failure occurred. About 90% of MIT students and teachers were still using the old system. However, the lucky 10% who had migrated had already felt the pain when the Exchange email server went down twice in the prior weeks.

HP Legacy Sales Commission System Plagued with Glitches

The Deal, August 5, 2009 – HP is facing challenges paying about 2,000 of its 23,000 sales people because of problems with a legacy system known as Omega. Omega was inherited by HP as part of its acquisition of Compaq. Compaq, in turn, inherited it from Digital when Compaq acquired that company. As a result, most of the affected sales people are those selling back-office systems to corporate customers, a business that Digital was in. HP has aggressively consolidated thousands of systems, but they still have a ways to go.

Router Failure Takes Down WestJet Airlines

Calgary Herald, August 7, 2009 – A router in WestJet's Calgary network center took down the airline's computer system that it uses to check in passengers. The router was redundant, and its backup should have taken over immediately. The problem was, the router didn't quite die. Though it failed to forward traffic, it was not sick enough for its backup to take over. About 1,000 passengers were affected for over an hour, and several flights out of Toronto and Montreal were seriously delayed.

Availability at Work – The Mars Orbiter

Tehran Times, August 11, 2009 – NASA's four-year old Mars Orbiter is controlled by a computer system comprising a redundant pair of computers. In early August, the Orbiter's primary computer unexpectedly switched to its backup and went into safe mode – a self-defense mechanism that allows the spacecraft to shut down and wait for instructions from Earth. Ground-based engineers restarted the Orbiter's scientific experiments two days later. So far, the Orbiter has switched to its backup computer three times. Redundancy works.

Texas Loses Access to its Business Records

Marshall News Messenger, August 12, 2009 – A computer hardware fault cut access to the business records kept by the Texas' Secretary of State. The outage lasted for almost two weeks. During this time, business incorporations, real-estate closings, voter registrations, and other activities were put on hold. The hardware fault was in a data center managed by IBM and its partners as part of an \$863 million contract to consolidate data-center operations across state agencies.

Typhoon Disrupts East Asian Internet and Voice Services

Reuters, August 13, 2009 – From August 9 to August 12, users in Taiwan, Singapore, Hong Kong, and the Philippines watched their Internet services deteriorate from slow to none. It turns out that Typhoon Morakat damaged several undersea cables, not directly, but through undersea landslides. The landslides occurred over a several day period, breaking cables one at a time. Though it may take months to repair the cables, service was restored within a week by rerouting traffic over other networks.

ISP Shuts Down Web Site, Goes Down Itself (Sabotage?)

Computerworld, August 25, 2009 – On Friday, August 21, a Swedish court ordered Black Internet, a Swedish ISP, to shut down access to Pirate Bay, which had been accused of providing copyrighted songs, movies, games, and software for free. On Monday, Black Internet did just that. But a few hours later, Black Internet itself went down, apparently due to hacker sabotage. The outage took down hundreds of Internet service resellers. Meanwhile, Pirate Bay returned quickly through another ISP, a new one set up by a convicted Pirate Bay financier.

DNA Results Scrambled by Hardware Incompatibility

Bio-IT World, August 27, 2009 – deCODEme is a genome service that will analyze up to one-million single nucleotide polymorphisms (SNPs) in a client's DNA. One client found that the presentation of his results differed on different browsers. Several dozen SNPs were presented differently depending upon the browser used. It turns out that the problem was an incompatibility with 32-bit processors. The presentations were correct on 64-bit processors. The problem has since been corrected.

Bandit Game Servers Take Down Much of China's Internet Services

PC World, August 28, 2009 – China is rife with private servers used to illicitly run copies of popular games for profit. The operators of these servers have a history of launching distributed denial-of-service (DDOS) attacks on each other. However, in one case in May of 2009, an operator went too far. He launched a DDOS on a DNS (Domain Name Server) that provided URL-to-IP address conversions for several competitors. Unfortunately, this was a major DNS server in China's network. As it became overloaded, it sent excess DNS requests to other servers, which overloaded. Much of China's Internet service was down for hours.

Maintenance Disrupts California Airspace

PR Newswire, August 28, 2009 – A maintenance subcontractor's mistake shut down the FAA Telecommunications Infrastructure (FTI) in the Oakland Air Traffic Control center, affecting flights in Northern California, Nevada, and the Pacific Ocean. Controllers had to rely on cell phones to coordinate flights with neighboring FAA facilities. It seems that a maintenance error created problems with the communication system. The backup system was put into service but failed the next day, causing the outage. The controllers were never notified that they were running with a single point of failure. Normally, such a notification would have put them on alert status.

Customs System Failure Piles Up Cargo in Malaysia

Business Asia One, September 8, 2009 – The Malaysian Custom Department's mainframe electronic document-processing system, located at the Kuala Lumpur International Airport, failed, cutting Custom's communications between airports, seaports, and border checkpoints. The custom clearing process, which normally takes two hours, became a manual process taking 24 hours. Cargo, including perishable goods, stacked up for days. It took well over a week to return the system to service. What? No backup?

Wild Street Saw Disrupts New York Trading

Reuters, September 15, 2009 – Just before the end of the trading day, when hedge funds depend upon millisecond trade execution times, a construction saw in one stroke severed 144 strands of fiber carrying 60,000,000 connections in New York City. The fiber cables fed just one building, a critical hub for global Internet connectivity in the Wall Street area. Though the fiber network throughout New York City is highly redundant with many alternate paths, one engineer said that 99% of New York buildings have single points of entry for all of their communication facilities.

Sydney Experiences Internet Outage When Contractor Drills Through Cables

Bloomberg, September 16, 2009 – The very next day after the Wall Street disaster, and halfway around the world, Internet users in Sydney, Australia, shared the pain of their Wall Street brethren. A contractor drilling test holes for a power-cable upgrade in Sydney's city center severed cables belonging to Australia's largest telecommunications provider, Telstra. The incident damaged eleven cables containing 10,000 communication lines. It required a week for Telstra to return full service to all of its subscribers in the area.

Déjà Vu! The Malaysia Customs Bug Hits Delhi

Express India, September 26, 2009 – Within a month, the customs problems that hit Malaysia (see above) traveled to Delhi's Indira Gandhi International Airport. The Custom Service's sixteen-year old Electronic Data Interchange (EDI) system was taken out of service when a communication cable snapped. The airport's 2,000 shipments per day had to be documented manually, resulting in nearly empty cargo aircraft leaving the airport. It was estimated that exporters were losing about one-million dollars per day. It took over a week to restore the system to service.

“America's Most Convenient Bank” Inconveniences Customers

CNN Money, October 1, 2009 – A year ago, Toronto-based TD Bank acquired New-Jersey based Commerce Bank and advertised the new TD Bank as “America's Most Convenient Bank.” But when it tried to integrate the back-office processing of the two banks, “convenient” became “inconvenient.” On three successive nights at the end of September, the integrated batch processing failed, leaving customers with access only to day-old data and no online services. End-of-month automatic bill paying was late, leading to penalties. The bank absorbed the fees.

Single Point of Failure Causes Verizon Outage

Computerworld, October 5, 2009 – A stalled Verizon router interrupted DSL and fiber Internet service to much of New York and Massachusetts for an hour or more on October 5. Though the router was part of a redundant network, it failed in such a way that it appeared to be operable to neighboring routers. Therefore, other routers continued to send traffic to it, which the stalled router simply dropped. Likewise, the backup router thought that the stalled router was still operational and did not take over. Ultimately, the offending router was rebooted; and service was restored.

The SwissDisk Storage Cloud Loses Customer Data

The Register, October 19, 2009 – SwissDisk of Las Vegas, Nevada, provides online storage services. In mid-October, it suffered a “catastrophic hardware failure” that prevented users from accessing their data. It appears that some of this data was lost, as users had to reregister once the system was restored to service. SwissDisk announced that it was taking steps to provide 99.95% availability. That means that users will be unable to access their data four hours per year. If you have critical data that you want to put in the cloud, make sure you have a backup and an ability to access the backup data when the cloud goes down.

Casino Gambler Denied His Multimillion Dollar Payoff

Miami New Times, November 3, 2009 – A gambler was playing a \$4 slot machine at the Seminole Hard Rock Casino in Tampa, Florida, when the bells started ringing. \$166,000,000 was flashing on the payoff display, and he thought he was set for life. Not so, said the casino. The computer in the slot machine had malfunctioned. Its maximum payoff was \$90,000, and the maximum payoff on the \$1.50 that he had wagered was a paltry \$2,500. The casino said it would not even pay this because of the computer malfunction but finally relented and settled for an undisclosed amount.

T-Mobile Drops 5% of its U.S. Customers

Channel Web, November 3, 2009 – T-Mobile dropped mobile services to 5% of its U.S. customer base late Monday afternoon, November 2, into early Tuesday morning. Affected customers were without service for more than eight hours. T-Mobile attributed the problem to a software error that

caused abnormal congestion in a back-end system. This didn't help T-Mobile's image after the disastrous loss of customer data by its Sidekick smart-phone service the previous month.² However, T-Mobile's image was unfairly tarnished by the Sidekick failure because Sidekick is run by Microsoft, T-Mobile's partner in the offering.

BBC's Web Site Off the Air for Half a Day

Telegraph, November 5, 2009 – The BBC web site is the site that everyone in the U.K. turns to when they think the Internet is down because its web site never fails. But the U.K. 'go to' web site did fail on November 5. In a statement, BBC acknowledged that the failure was due to a network problem but elaborated no further. BBC apologized for the outage and indicated that BBC engineers were monitoring the network to ensure that no further problems occurred.

Washington, DC, Transit System Impacted by Non-Redundant Power Failure

Washington Post, November 5, 2009 – The Metro is the transit system for Washington, DC. Its operational headquarters in downtown Washington depends upon two 27-year old power distribution systems, neither of which has a backup. One distribution system powers computers that operate the rail system. The other powers about a dozen vital communications systems, including two-way communication with bus drivers, the subway public address system, the bus vehicle locator, and fare card machines. On Thursday, November 5, the latter power distribution system reached its end-of-life. It took a half day to restore service with a temporary fix, but the transit agency has no money to replace the units. What's a poor city to do?

London Stock Exchange Gives Up on Windows Trading Platform

Computerworld, November 11, 2009 – Following a major day-long crash of its distributed Windows-based TradElect trading system a year ago,³ the London Stock Exchange once again halted trading on almost 300 securities due to problems with TradElect. After having moved to it from a fault-tolerant NonStop system, the Exchange has now given up on TradElect, which was developed for it by Accenture. The Exchange has changed CEOs and is moving to a Linux-based trading system from Sri Lanka-based MilleniumIT.

Storms Black Out Half of Brazil

Wall Street Journal, November 12, 2009 – Intense lightening, wind, and rain collapsed three high-power transmission lines, blacking out half of Brazil for several hours. A domino effect caused a 14,000 megawatt hydroelectric plant to go offline, depriving 60,000,000 Brazilians and all of Paraguay of power. People were stranded in elevators and on commuter trains. To avoid a repeat of this during the 2016 Olympic games, Brazil plans to isolate Rio de Janeiro from the national power system and to provide Rio with its own independent energy feed.

Vodafone Maintenance Causes Network Outage

Telegraph, November 12, 2009 – U.K. mobile-carrier Vodafone services almost 19 million customers. On Monday morning, November 12, engineers working on its network evidently caused a problem, which dropped service to some of its customers for several hours. Customers could not make calls, and smart-phone users were unable to access the Internet. The timing was an embarrassment for Vodafone since it will start selling the iPhone next year when Apple's exclusive arrangement with O2 ends.

² Sidekick – Your Data is in 'Danger', *Availability Digest*, November 2009.

³ London Stock Exchange PC-Trading System Down for a Day, *Availability Digest*, October, 2008.