the Availability Digest

www.availabilitydigest.com

More Never Agains VIII

February 2013

Since our last summary of multiple data center failures, published in our September, 2012, issue of the *Availability Digest*, we have reported on several major outages. They have included the DDoS (Distributed Denial of Service) attacks on several major U.S. banks, attacks that took down their online websites for several days in retaliation for the YouTube video entitled “The Innocence of Islam;” a Republican-party secret weapon whose failure may have cost Mitt Romney the U.S. Presidential Election; a memory leak that took down an Amazon Web Services Availability Zone for several hours; several data centers that were out for days after Hurricane Sandy flooded generators located in basements in lower Manhattan; and a violent storm (not Sandy) that left millions of residents of Northern Virginia without 911 service for four days.

Several other outages made headlines during this time and are summarized below.

Chinese Hackers Attack the New York Times

In October, 2012, the New York Times was preparing to expose a multibillion dollar fortune obtained through perhaps nefarious business dealings by relatives of Wen Jiabao, China’s prime minister. It was then that suspected Chinese hackers infiltrated the Times’ computer systems and stole the passwords of its reporters and other employees. Though the attackers could have wreaked havoc on the Times’ computers, they did not. It seems their purpose was to identify Chinese citizens who may have provided information to the New York Times concerning Mr. Jiabao’s business dealings.

Security expert firm Mandiant was hired by the Times and identified the attacks as similar to earlier attacks perpetrated by the Chinese military. The attacks were not made directly but rather through several computers at U.S. universities in an attempt to hide the source of the attacks.

The security experts found evidence that the hackers stole the passwords of every employee of the Times and used these passwords to gain access to the computers of 53 of the employees. The hackers evidently initiated their attack by snooping around for two weeks through back doors that they had opened in the systems until they found the domain controller that contained the user names and their hashed passwords. Even though the passwords were protected by hashing, hashed passwords can be easily cracked via the use of published “rainbow tables” that give hash values for nearly every alphanumeric character combination.

It is still not known how the hackers broke into the Times’ computers, but it is strongly suspected that it was through a spear-phishing attack that requires just one employee to download just one malicious email that will install malware. This malware can then siphon off oceans of data, including keystrokes, screen images, documents, microphones, and cameras.

Over the course of three months, the attackers installed 45 pieces of malware. The Times' antivirus protection from Symantic found only one of these and quarantined it.

Once detected, the Times' allowed the hackers to spin a digital web for four months to identify every back door that was being used to gain access to the system. The Times then replaced every compromised computer and set up defenses to prevent further infections.

Mandiant reports that it has traced several other corporate cyber attacks to the Chinese military.

The Gozi Trojan is Finally Saddled

Starting in 2007, when computer users around the world would log into their banking accounts, they were greeted with a message that said, "In order to provide you with extra security, we occasionally need to ask for additional information when you access your account online."

This was not a message from the bank. It was a message from the Gozi Trojan that had infected the users' computers. Many were duped into entering their social security numbers, their mother's maiden names, and other personal information.

As soon as the information was entered, the Gozi Trojan sent it to a server in California, and then to a central command and control server in the Netherlands. The personal data was sold to other cybercriminals, who used it to access accounts and to steal deposits. Tens of millions of dollars were stolen from over a million accounts worldwide, including 40,000 in the U.S.

In early 2013, three Eastern European men were arrested by U.S. law enforcement and charged with operating the Gozi scheme. The mastermind was a Russian. A Latvian was charged with writing some of the code, and a Romanian was charged with operating a hosting service that allowed the Gozi Trojan to be distributed. Still at large is an unknown programmer who wrote the original code.

The three men face criminal charges ranging from conspiracy to computer intrusion and bank fraud. They each face 60 to 90 years in prison.

Amazon Takes Out Netflix Over Christmas

Shortly after Noon, Pacific Standard Time, on Christmas Eve Day 2012, as families gathered to watch movies, Netflix stopped streaming. The outage lasted until almost Noon the next day, Christmas Day. A developer error took down a major portion of Amazon Web Services' East Coast Availability Zone. Netflix depends upon the Amazon cloud for its streaming services.

According to Amazon, the outage began when a developer accidentally deleted data while performing maintenance on the East Coast Elastic Load Balancing system (ELB). The developer's action impacted almost 10% of the ELBs. The ELB distributes traffic across several compute instances in different Availability Zones. As a result, some traffic could not be distributed and came to a halt.

Because of this experience, Amazon has instituted a new procedure that requires developers to obtain specific change approval from AWS each time they access the ELB system.

Failures of AWS can impact a major portion of Internet services around the world. AWS now handles more than 835,000 requests per second for hundreds of thousands of customers in 190 countries.

Microsoft's Azure Cloud Down – Again

Just a week after Amazon lost its cloud over Christmas, Microsoft's Azure Cloud followed suit.

A year ago, Amazon's Azure Cloud was knocked out by a 2012 leap-year software bug for over a day.[1] Almost a year later, on December 28, 2012, Microsoft reported that it had lost the Azure Storage service for its South Central U.S. Region. It took over two days for Microsoft to restore service.

According to Microsoft, the outage affected about 2% of Windows Azure Storage accounts that were located in one cluster. Making matters worse was that the Azure health dashboard wasn't working because the outage affected the Azure worldwide Management Portal that depended upon this cluster.

According to a detailed blog posted by Microsoft, three issues led to the outage. First, some of the storage nodes did not have protection turned on. Second, the monitoring system that is supposed to detect this condition had a defect, resulting in no alarms or escalation. Third, a transition to a new primary node triggered a reaction that led to an incorrect formatting of other nodes.

Normally, Azure should have survived the failure of two nodes of storage, as it keeps three copies of the data spread across three separate fault domains. However, the reformatted nodes were spread across all three fault domains, leading to the total unavailability of the data.

Going forward, Microsoft will offer optional geographical redundancy to improve availability. A copy of the data will be replicated to a remote location and will be available for read-only access by the customer. Should the primary copy fail, the customer will have a choice as to whether to fail over (maximize data availability) or not to fail over (maximize data durability).

South Carolina Computers Hacked for Three Million Social Security Numbers

In early October, 2012, the South Carolina Division of Information Technology informed the S.C. Department of Revenue of a potential cyber attack involving the personal information of state taxpayers. After verifying that this may have happened, the Dept. of Revenue retained security firm Mandiant to analyze the break, plug the holes, and determine the damage.

Mandiant's assessment was that a Dept. of Revenue employee unwittingly executed malware by clicking on a malicious link in a salacious email. The malware stole the employee's credentials and allowed the hackers to gain access to the system. By the time the malicious code was identified and shut down ten days later, the hackers were able to access other Dept. of Revenue systems and databases. They stole millions of social security numbers and bank accounts and almost a half-million credit card numbers, most but not all of them fortunately encrypted.

South Carolina officials pointed their fingers at Russian hackers and blamed the IRS for not requiring them to encrypt social security numbers. But Mandiant placed the blame on state officials. The state's Dept. of Revenue director had declined an offer from the state's IT department for free breach-detection services. As a consequence, the state failed to spot the compromise of forty-four different systems, the installation of backdoor software, multiple instances of password hashes being dumped, the running of Windows batch scripts, and the execution of numerous commands against the databases. The attacker located and copied 74 GB of backup files to an external server outside of the state's control. Though finally detected in mid-October, the attacks had been going on for a month. The director has since resigned.

The state's bill for the data breach is expected to exceed $14 million, including Mandiant's bill, taxpayer notifications, Experian credit services offered to those whose data was breached, security improvements, and legal and public relations expenses.

[1] Windows Azure Cloud Succumbs to Leap Year, *Availability Digest*; March 2012. http://www.availabilitydigest.com/public_articles/0703/azure.pdf

United Airlines Grounded by Nationwide Computer Glitch

A computer problem in United's Unimatic system used for its ground operations brought flights across the country to a halt at 8:30 AM on Thursday, November 15th. Though the problem was resolved two hours later, it was too late for many flights. Ten flights were cancelled, and 636 flights were delayed either directly by the outage or by delays in earlier flights. The delays affected most major cities in the United States and even some foreign cities such as London.

The problem involved communication between dispatchers at the company's operations center in Chicago and planes at airports around the world. The dispatchers communicate such information as weight and fuel loads to pilots, who need the information to operate the flight.

United merged with Continental Airlines in 2010. but Continental and United Express flights were not affected. United has been struggling with technology problems associated with integrating the computing systems of the two airlines since the merger. It experienced problems in March, 2012, when it switched to a passenger information system previously used by Continental. The system needed significant reworking. United suffered another outage in May. In August, 2012, 580 United flights were delayed; and its web site was shut down for two hours because of a server failure.

GitHub Downed Twice by Network Software Errors

GitHub is a major open-source repository and collaborative site hosting millions of public repositories. In November, 2012, it suffered a major outage as it attempted to upgrade its network. The very next month, it again suffered a longer outage for similar reasons.

The November outage occurred when GitHub attempted to restructure its network from a linear network to a tree structure in order to improve transit time. Instead of traffic being routed through many network switches, the new network would be a three-level tree network; and traffic only had to pass through three switches. When it tried to bring up the network, GitHub found the network to be flooded with self-generated traffic. Working with its (unnamed) switch vendor, GitHub determined that the switches were not learning the MAC addresses held by their neighbors. Therefore, for many interswitch transfers, the destination address was being flooded to all of the switches in the network to determine the servicing port. It took two hours to determine the problem and to restore the network to operation.

Having resolved this problem, GitHub scheduled a maintenance window in late December to upgrade the switch software provided by its switch vendor. Its switching network is completely redundant, so the plan was to upgrade one side first, test it, and then to fail over and upgrade the other side. Unfortunately, there was a problem in the switch failover process; and the network collapsed for several minutes.

This caused the redundant file server clusters to lose heartbeats. Each file server assumed that its companion had failed and subsequently tried to take over. To ensure against "split brain" operation, part of the takeover involves killing power to the supposed bad server. Each server cut power to its companion, and in many clusters both servers were taken offline. It took five hours to restore the servers and their databases before GitHub could get back into service. Evidently, GitHub does not deploy an independent monitoring Quorum server to resolve such conflicts.

Summary

These outages were caused by a wide range of problems. Included were a user error (Amazon), software bugs (Microsoft Azure, United, GitHub), and hardware (United). However, in this summary, for the first time, cyber attacks were predominant (New York Times, the Gozi Trojan and the South Carolina Department of Revenue). As we see data centers become increasingly robust with respect to failures, we are now witnessing more problems with major malware infections.