

More Never Again VI

April 2012

The first quarter of 2012 has had its share of catastrophic outages. We have already reported on the discovery of Oracle's ticking time bomb,¹ on the continuing string of outages suffered by Australia's "big four" banks,² and on the multiday failure of Microsoft's Azure public cloud.³ We summarize below some of the others that have made headlines in the past few months.

Zoho SaaS Service Downed for Ten Hours by Abrupt Power Failure

Over five-million users take advantage of SaaS-provider Zoho's online collaborative, business, and productivity applications; but applications were not available during the day on Friday, January 20, 2012. Zoho's collocated data center, managed by Equinix in Silicon Valley, California, suffered an "abrupt" power failure about 8AM PST.⁴ This caused every database cluster and every server to instantly fail. The manual effort to recover and test these systems overwhelmed its recovery team. Zoho runs over one-hundred database clusters, and 40% of them had synchronization errors. Consequently, service was not fully restored until after 6 PM.

An "abrupt" power failure evidently meant that many levels of redundancy failed. Specifically, the data center's UPS (uninterruptible power supply) did not kick in and give servers and databases time to shut down gracefully.

Fortunately, no data was lost as a result of the outage. Zoho mirrors all data in a secondary Equinix data center in New York to protect against data loss. However, there is not enough capacity in the New York data center to run all Zoho services.

To prevent such a problem in the future, Zoho is installing its own UPS and is moving to database servers with internal batteries so that they can shut down gracefully even in the event of an abrupt power outage. It is also considering increasing the capacity of its secondary data center so that it can take over application services in the event of a total outage of its primary data center.

¹ Oracle's Ticking Time Bomb, *Availability Digest*, February 2012.

http://www.availabilitydigest.com/public_articles/0702/oracle_scn.pdf

² Australia's Painful Banking Outages, *Availability Digest*, March 2012.

http://www.availabilitydigest.com/public_articles/0703/australian_bank_outages.pdf

³ Windows Azure Cloud Succumbs to Leap Year, *Availability Digest*, March 2012.

http://www.availabilitydigest.com/public_articles/0703/azure.pdf

⁴ Our Friday Outage and Actions We Are Taking, *Zoho Blogs*, January 23, 2010.

Tokyo Stock Exchange Suspends Trading

Just before the market opened on Thursday, February 2, 2012, trading in 241 securities for 222 companies was halted at the Tokyo Stock Exchange (TSE), the world's third largest exchange.⁵ The suspension included blue-chip stocks such as those for Sony and Hitachi as well as other instruments such as exchange-traded funds, real-estate investment trusts, and convertible bonds. Trading in all suspended shares resumed for the afternoon session.

The outage hit during the middle of the earnings season, a particularly active time for trading. February 2 turned out to be the third-heaviest trading day of the year.

According to the TSE, one of the eight servers the exchange uses to distribute stock data failed, and its backup did not take over. The failure affected about 10% of the exchange's 2,459 listed securities.

In early 2010, the exchange replaced its aging trading system following a series of disruptions in 2005, 2006, and 2008. The new Arrowhead high-speed trading platform has had generally favorable reviews from the investment community. This was its first outage since it went into service.

FOUR African Undersea Cables Chopped

Over the years, we have reported on many undersea data-cable breaks that have taken down Internet access for days, particularly on the African continent. It happened again in February, 2012. But this time it was four cables cut in two different regions.⁶ Internet transfer rates were brought to their knees in nine countries in Eastern and Southern Africa as traffic was rerouted over surviving cables.

On February 15th, three cables in the Red Sea - the Europe India Gateway (EIG), the Southeast Asia Middle East Western Europe-3 (SMW-3), and the Eastern Africa Submarine Cable System (ESSAY) - were all severed at the same time. Ten days later, on February 25th, it was suspected that a ship dropped anchor off the coast of Kenya, cutting The East African Marine Systems (TEAMS) cable. TEAMS was already carrying much of the traffic normally routed over the three previously broken cables.

With four cables cut in such a short time, there were suspicions of sabotage. However, the three Red Sea cables lay in relatively shallow water about 650 feet below the surface and could have been cut by a passing ship similar to the fate suspected of the TEAMS cable off Kenya. Repairs took about three weeks before full Internet service was restored to the affected countries.

Server Provisioning Disabled in Ninefold Cloud

Ninefold is a major Australian cloud provider. On Thursday, February 17, 2012, it suffered a major outage that caused all virtual-server provisioning in its cloud to be disabled for five hours.⁷

According to Ninefold, a Network File System (NFS) server experienced a failure; but it automatically restarted and resumed normal operation in a few minutes. However, a number of physical host servers with their virtual machines (VMs) that were performing operations with the NFS server became unresponsive to further provisioning. Eventually, some customer VMs on these host servers became unavailable; and they could not be restarted on other physical hosts.

Though the failure did not affect the majority of physical hosts or virtual servers, Ninefold determined that the most expedient way to return service to the affected customers was to terminate provisioning services for all customers. Consequently, no customer could provision new VMs for the five-hour period of the outage.

⁵ [Tokyo Tackles Trading Glitch](#), *Wall Street Journal*, February 2, 2012.

⁶ [Epic net outage in Africa as FOUR undersea cables chopped](#), *The Register*, February 28, 2012.

⁷ [Ninefold outage disables server provisioning](#), *CRN*, February 21, 2012.

Ninefold had previously suffered a host server outage in August of last year and another incident a few months earlier in May. It is expecting to launch a second availability zone in another data center in May of 2012 and is working out the details to see if the two data centers can be run in an active/active configuration.

U.S. Defense Information Systems Agency Down Again

The U.S. Defense Information Systems Agency (DISA) provides IT services for the Department of Defense (DoD) and as such has a responsibility for ensuring network reliability. It will soon be responsible for hosting the U.S. Army's entire email and collaboration system in a private cloud. However, DISA is facing its own resiliency problems.

On Thursday, March 1, 2012, DISA users in Washington, D.C. and in some Midwest regions lost access to the Internet for three hours because of problems at three of the DoD (Department of Defense) gateways that interface the DISA web site with commercial Internet access points.⁸ DISA worked with commercial vendors and partners to reroute traffic until the network was restored.

Network issues are not new to government agencies. Last November, a Storage Area Network fault took down DISA's web site and blocked DoD personnel from accessing applications. It took five days to resolve the problem. CIO.gov, the web site for the White House's CIO Council, had several outages in the past few months. In October, a series of faults with the USAJobs.gov web site prompted a Congressional investigation.

BATS Halts Trading in its Own Shares

BATS Global Markets is the third-largest equities exchange operator both in the U.S. and globally. In Europe, BATS Chi-X Europe is the largest pan-European equities market operator.

On Friday, March 23, 2012, BATS encountered a software bug that caused an infinite loop and made open customer orders inaccessible for more than two hours for stocks with symbols in the range of "A" to "BFZZZ."⁹ Not only did this halt trading in Apple (AAPL) and BATS, but BATS stock plunged from \$16 a share to 2 cents in just nine seconds during its IPO auction, BATS' long-awaited initial public offering, which had just begun.

BATS resolved the problem by noon for all stocks except its own. BATS management determined that the delay of more than two hours had eroded investor confidence, and it elected to cancel its IPO.

We talk about the penalties of outages in terms of cost of downtime, erosion of confidence of customers and suppliers, and bad publicity. Here is a new and devastating cost – the aborting of an IPO.

California's Child Services Loses 800,000 Social Security Numbers

The California Department of Child Services needed to test its ability to cope with a disaster. Part of the test was to ensure that it could manage its data remotely. To run this test, it shipped data regarding 800,000 adults and children from Sacramento, California, to Boulder, Colorado, on four computer-storage devices. The test went well but ended up in a disaster of its own making.

The storage devices were supposed to be shipped back to Sacramento via Iron Mountain at the end of March; but since Iron Mountain did not offer air services, the devices were shipped by FedEx instead. FedEx picked them up but never delivered them. Suddenly, names, addresses, social security numbers,

⁸ DISA Suffers Another Network Outage, *InformationWeek*; March 5, 2012.

⁹ BATS Global Markets does damage control, doesn't plan new IPO, *The Business Journals*; March 26, 2012.

drivers' licenses, health insurance providers, employers, and other information on parents and their children were gone.¹⁰

According to a state official, the state's belief is that the container holding the devices was not properly secured and fell out during transport. So far, there has been no indication that any information contained on the devices has been improperly accessed or used. Everyone who may possibly have been affected by this incident has been notified; and credit reporting agencies, the state Attorney General, and the state's Office of Privacy Protection are closely monitoring the situation.

Dutch Vodafone Network Destroyed by Fire

A fire in a Vodafone network facility in Rotterdam terminated cell-phone service for more than a million customers for a week.¹¹ The fire on Wednesday, April 4, 2012, destroyed the local network infrastructure and put 700 transmission towers out of commission. Mobile users were unable to place calls, send messages, use voice mail, or access mobile Internet.

Affected were customers in Rotterdam, The Hague, and surrounding urban areas. Vodafone, the largest mobile-phone operator in the world and the second largest in The Netherlands, said in a statement that restoring service after the fire was one of the most comprehensive and complex operations in its history. It was not until the following Tuesday, April 10th, that Vodafone confirmed that its network was "almost stable."

Fortunately, no one was in the building at the time; so there were no injuries. The government plans to hold talks with all of the country's mobile operators for possible network sharing in the face of future outages such as this.

Summary

Software bugs seem to be the big problem in this series of outages. Oracle had its ticking-time-bomb bug. Azure had a leap-year bug. BATS went down with an infinite-loop bug.

Equally important were recovery faults. At Zoho, the UPS failed following a power failure. At the Tokyo Stock Exchange, a backup data-distribution server failed to take over. At Ninefold, host servers failed following the recovery of an NFS server.

These outages have a common characteristic – testing. It seems that no matter how much system testing we do, there are always remaining problems. How much testing is justified? That, of course, depends upon the cost associated with an outage.

Then there were the outages that no testing would have prevented – a fire and a FedEx loss.

¹⁰ Misplaced data leave 800,000 Californians exposed, *CNET News*; March 30, 2012.

¹¹ Vodafone Network In Holland Almost Stable After Interruption Caused By Fire, *Wall Street Journal*; April 10, 2012.