the Availability Digest

www.availabilitydigest.com

Linux Leap-Second Bug Takes Down Data Centers

August 2012

What a difference a second can make. At the stroke of midnight the evening of Saturday, June 30, 2012, servers all over the world began to crash. Was this another Stuxnet virus[1] propagated by some rogue government to take down the world's IT infrastructure?

No. The crash was caused by a leap second, which is added every few years to keep the world's clocks in synchronism with the earth's rotation. A bug in thousands of unpatched versions of Linux choked on this. These servers had to be rebooted, causing hours of downtime at some of the Internet's most popular sites.

The Leap Second

Universal Coordinated Time, or UTC, is the measure of time defined by highly precise atomic clocks. Our sense of time is governed by the earth's rotation and is represented by Greenwich Mean Time (GMT). But the earth's rotational speed wobbles due to many factors. This variance is monitored by the International Earth Rotation and Reference System Service.

To account for this variation, leap seconds were introduced in 1972. Every year or two, a leap second is added to or subtracted from our GMT clocks to resynchronize the world's clocks with the earth's rotational position. So far, 25 leap seconds have been added to our clocks. Leap seconds are typically issued at midnight on January 1st or June 30th.

To keep computing systems in time synchronization, the Internet utility NTP (Network Time Protocol)[2] is used. NTP broadcasts the correct time over the Internet and uses complex algorithms to account for Internet communication-latency differences. When a leap second occurs, NTP automatically adjusts its time to account for the leap second.

At midnight on June 30, 2012, a leap second was added.

The Linux Bug

Developers had previously found that a software bug in the Linux operating-system kernel prevented it from properly handling leap-second adjustments broadcast by NTP. This bug was corrected in a March, 2012, patch; but many system administrators had not installed the patch by the time June 30th came around.

Tux, the Penguin
Mascot of Linux

[1] Stuxnet, The World's First Cyberweapon, *Availability Digest*; March 2011.
http://www.availabilitydigest.com/public_articles/0603/stuxnet.pdf

[2] Time Synchronization for Distributed Systems – Part 1, *Availability Digest*; November, 2007.
http://www.availabilitydigest.com/public_articles/0211/time_sync_1.pdf

The bug resided in Linux's high-resolution timers (hrtimer). This logic failed to reset the system time when a leap second was encountered via NTP. Therefore, the hrtimer time, which is driven by NTP, was one second ahead of the system time following the occurrence of a leap second.

The problem occurred when an application running in an unpatched version of Linux called for a kernel timeout of less than one second. The Linux high-resolution timer saw that the time already had expired (it thought the time was one second later than the system time) and returned to the program immediately. In the event of a timeout, many applications simply repeat the timeout request, creating an endless loop and causing 100% CPU utilization.

System administrators found that rebooting the application typically did not resolve the issue (the timing error was at the kernel level, not the application level). Therefore, they finally resorted to rebooting their servers. Rebooting corrected the problem, as the kernel resynchronized its time on reboot; but rebooting took an hour or more, during which time the applications were down.

Had the system administrators known, all they would have had to do was to resynchronize the system clock. This would have cured the problem immediately. But the true source of the problem was not known by many until much later.

The Consequences

Many popular Internet sites were affected, including LinkedIn, Mozilla, and Reddit. The Amadeus' ALTEA airline hosting system that many airlines such as Qantas and Virgin Australia use for passenger check-in and ground services was down for almost an hour. In some data centers, hundreds of servers had to be rebooted.

Initial conjecture was that the problem was in Java, as most affected applications were written in Java. Especially affected were the Hadoop database manager and the Cassandra database. It was only later that the bug was isolated to the Linux kernel.

Interestingly, another consequence of the bug in large data centers was a sudden spike in power utilization when a good number of servers started to run at full load.

Google was one Internet company that did not have the problem. It had prepared its systems for just such an occasion via a technique it calls "leap smear." Prior to the arrival of a leap second, Google slowly adds a few milliseconds at a time to its system clocks so that its systems are already in synchronism when the leap second is applied by NTP.

The Bug Fix

Developers continue to work on fixes for the leap second problem. In addition to their March, 2012, patch, they have incorporated fixes into Linux release 3.5 and have retrofitted these fixes into Linux versions 3.0.38, 3.2.24, and 3.4.6.

Summary

Surprisingly, just one month later, at midnight on July 31st, many servers again went down due to a leap-second addition. However, in this case, the leap second was bogus. It appears that this was some sort of a denial-of-service attack on unpatched Linux kernels.

It was just a few months ago that we reported on a Leap Day bug that took down Microsoft's Azure cloud services for a day and a half.[3] It looks like time adjustments for leap events should be approached with great caution by system administrators.

The bottom line: keep up with your patches. But never patch all of your systems at the same time. Try a new patch for a few days in a small number of select servers. Otherwise, you may be subject to the McAfee syndrome,[4] which caused millions of PC systems worldwide to lock up because of a faulty upgrade to McAfee's antivirus product.

Acknowledgements

The information in this article was taken from the following sources:

Linux is culprit in leap-second lapses, says Cassandra exec, *Computerworld*; July 2, 2012.
Leap second bedevils Web systems over weekend, *Computerworld*; July 2, 2012.
Amadeus admits the ALTEA crash was a result of Leap-Second Bug, *tnooz*; July 3, 2012.
"Leap second" bug causes Internet glitch, *GMA Network*; July 3, 2012.
Leap second bug in Linux wastes electricity, *The H*; July 3, 2012.
Qantas outage pinpointed on leap second Linux bug, *IT News*; July 3, 2012.
Hosts: Leap Second Caused Spike in Power Usage, *Data Center Knowledge*; July 3, 2012.
Bogus leap second disrupts Linux systems, *The H*; August 1, 2012.

[3] Windows Azure Cloud Succumbs to Leap Day, *Availability Digest*; March 2012.
http://www.availabilitydigest.com/public_articles/0703/azure.pdf
[4] Anti-Virus – A Single Point of Failure?, *Availability Digest*; May 2010.
http://www.availabilitydigest.com/public_articles/0505/mcafee.pdf