

the Availability Digest

www.availabilitydigest.com
[@availabilitydig](https://twitter.com/availabilitydig)

@availabilitydig – Our March Twitter Feed of Outages

March 2017

A challenge every issue for the Availability Digest is to determine which of the many availability topics out there win coveted status as Digest articles. We always regret not focusing our attention on the topics we bypass. With our new Twitter presence, we don't have to feel guilty. This article highlights some of the @availabilitydig tweets that made headlines in recent days.



AWS is investigating S3 issues, affecting Quora, Slack, Trello

Cloud infrastructure provider Amazon Web Services (AWS) on 28 February confirmed that it's looking into issues with its widely used S3 storage service in the major us-east-1 region of data centers in Northern Virginia. Other services are affected as well.

<https://t.co/IGBYTTdXfi>

NBN Co admits to 'exorbitant' number of satellite failures

NBN Co has admitted to an "exorbitant" rate of failures for its Sky Muster satellite service, forecasting that it may take as long as a year to iron out software bugs. Technical issues have plagued the satellite service since the first Sky Muster satellite launched into space in October 2015. Many users have complained of poor service as well as a total inability to connect. The launch of the second satellite late last year was marred by a botched software update to the Viasat modems that help power the service.

<https://t.co/fFk5l8cDf1>

SpaceX will send 2 people on a trip around the moon next year

SpaceX is going to send two people on a trip around the moon in 2018. But don't go getting fitted for your spacesuit just yet — the two people have already been picked, and SpaceX says they paid "a significant" (but undisclosed) amount for the privilege. The identities of the two space travelers was left unannounced. SpaceX notes that training will begin later this year (and that they'll say more about who it is once the health tests are cleared). Elon noted that the passengers will be traveling in a Dragon 2 capsule, powered by SpaceX's still-in-design Falcon Heavy rocket.

<https://t.co/rY0UXHEZXI>

Tens of thousands of Chromebooks fail because of Symantec BlueCoat problem

Well, this is annoying. Maryland's (state in U.S.) Montgomery County schools are using Chromebooks. The school system is using about 120,000 Chromebooks and multiple PCs running the Chrome web browser. But when Google recently updated them to Chrome OS 56, over 30 percent couldn't log on, while many PCs running Chrome were unable to reach the web. So, was it Google's fault? Not so fast. The school system was using Symantec's BlueCoat, a man-in-the-middle (MitM) SSL web proxy. This uses ProxySG technology to examine Secure-Socket Layer (SSL) and Transport Layer Security (TLS) encrypted web content. So far, so good -- if you want to make sure your seventh graders aren't peeking into pornography. But, in this case, it turns out BlueCoat doesn't support the newest standard web security protocol, TLS 1.3.

<https://t.co/ix3D43irrn>

Google: 'Sorry for wide-scope outage but canary testing brought our cloud down'

A botched software update triggered January's two-hour outage affecting Google Compute Engine (GCE) instances, cloud VPNs, and network load balancers. While the incident wasn't as serious as a past network outage, Google had promised a full explanation due to the "wide scope" of this one, which dropped connections to all GCE instances, cloud VPN tunnels and network load balancers that were created or live-migrated on Monday, January 30. The outage was triggered by a "large set" of updates to its load-balancing gear, although the outage itself was caused by updates getting jammed during testing inside a canary deployment.

<https://t.co/15AhBmAPUe>

Backup Explained

In five questions or less, an industry expert defines and explains a technology, term or trend – with this installment seeing Richard Agnew, VP NW EMEA at Veeam, tackle backup.

<https://t.co/rd370gUEvp>

Google: We're sorry, but our cloud wiped out your Wifi and OnHub routers

Google has apologized for annoying Google Wifi and OnHub customers after accidentally triggering a mass factory reset on the routers. A bunch of Wifi and OnHub owners panicked on 23 February after discovering the devices were returned to a factory state, leaving owners offline despite the absence of any disruption on carrier networks. Frustratingly for customers, the unrequested factory reset required users to reconfigure device names and network settings, which Google admits isn't that simple to do. The outage highlighted a problem with OnHub and Wifi's reliance on Google cloud services for network connectivity.

<https://t.co/XQUgBfbJIA>

Mastercard Fined Over RushCard Problems

Mastercard and UniRush, the company that sells RushCard-brand prepaid debit cards, have been ordered to pay \$13 million in fines and customer restitution over issues with the cards. The total includes \$3 million in fines and \$10 million in customer restitution. The agency left it up to the companies to decide how much each will pay. The fine stemmed from a 2015 system failure that left thousands of customers unable to access the money in their accounts. The CFPB said that the problems were caused by a lack of preparation and coordination by the companies before Mastercard became UniRush's main payment processor. Customers were unable to withdraw cash, receive direct deposits, make purchases, or get accurate account balance information.

<https://t.co/nbVmqqfDTjz>

HSBC Customers Left Frustrated After Internet Banking Outage

HSBC suffered an outage to its Business Internet Banking service in the UK on 27 February, with hundreds of users taking to Twitter to vent their frustrations at being unable to access their accounts. At 10:30 am, HSBC UK Business tweeted that it was “experiencing some issues with Business Internet Banking, but we’re working to resolve this. Apologies for any inconvenience caused.” As you would expect, the bank received a fierce backlash from angry Twitter users.

<https://t.co/7zJ00qp8gK>

Barclays apologies for cash machine and debit card outage - how to get compensation if you lost out

On a Saturday afternoon in February, Barclays customers reported issues with withdrawing money from cash machines and using debit cards in shops. Online and telephone banking and in-branch payments were also affected. The bank said not all customers were impacted but that issues had been reported across the UK. Angry customers flocked to Twitter to complain, with many saying they were left red-faced after being unable to pay for shopping or restaurant meals.

<https://t.co/Kf14ptTHTq>

Westpac outage downs branches

An internal systems outage at Westpac took the bank's branch network offline around the country for about seven hours on Monday, 27 February, and forced teller staff to resort to manual processes. It's the second significant outage to hit the bank in three months following a glitch in late November that stopped online and mobile banking platforms from processing payments or showing correct balances for almost a week.

<https://t.co/qk7lg1fe4C>

Everything You Need to Know About Cloudbleed, the Latest Internet Security Disaster

Have you heard? A tiny bug in Cloudflare's code has led an unknown quantity of data—including passwords, personal information, messages, cookies, and more—to leak all over the Internet. If you haven't heard of the so-called Cloudbleed vulnerability, keep reading. This is a scary big deal.

<https://t.co/0q09sAcx5C>

Internet Outage Causes American Airlines Ground Stop at Philadelphia Airport

A ground stop halted American Airlines flights bound for the airline's Philadelphia International Airport hub for nearly two hours in February. An American spokeswoman said the airport's Division of Aviation had a planned maintenance power outage from 1 to 4 a.m. When the power came back on, American's Internet access remained down. The outage, which impacted the earliest departures Wednesday morning, prevented the airline and passengers from printing boarding passes.

<https://t.co/URQJDwefO6>

Telstra enterprise outage caused by faulty hardware

Telstra has confirmed that its enterprise voice and data outage across New South Wales was caused by a faulty piece of hardware. The outage affected only enterprise customers on Tuesday morning, 21 February, until 1pm AEDT, with consumer fixed-line and mobile services not impacted by the hours-long outage.

<https://t.co/TGv4cs9nc6>

The ransomware attack that locked hotel guests out of their rooms

Of the various things you may be worrying about these days, getting locked out of your lovely hotel room in scenic Austria is probably not very high on the list. But in late January, guests at the Romantik Seehotel Jaegerwirt in the Austrian village of Turracherhohe found themselves unable to open their hotel room doors. In the end, the hotel ended up having to pay about \$1,800 (two bitcoins) to hackers who had penetrated its systems and managed to remotely lock its doors.

<https://t.co/1wIQjAZTKC>

HPE blames tax office outage on faulty SSD drive

Hewlett-Packard Enterprise has stated that outages at the Australian Taxation Office were caused by a problem with solid state drives provided by a separate vendor. A spokesperson for HPE stated that the disruption began when an SSD failed. While HPE did not name the vendor who provided the drive in question, the spokesperson did say that a joint effort between HPE and the vendor revealed that the outages were triggered by a 'rare issue', caused by a new, unique set of circumstances.

<https://t.co/JXkCNvO27W>

Google Outage Traced to Network Glitch

Google has tentatively traced the cause of a roughly two-hour global cloud outage to an internal software issue related to its virtual network traffic routing. The Google Compute Engine public cloud outage commenced at 10:40 p.m. Pacific time on Wednesday (Feb. 18). Traffic loss was stopped and normal traffic levels resumed at about 1:20 a.m. local time on Thursday (Feb. 19). In a status report, Google said the "preliminary" root cause of the outage was traced to the Compute Engine's virtual network for VM outgoing traffic. The company said the network stopped issuing routing information.

<https://t.co/r3KodexDLg>

Self-Driving Cars Have a Problem: Their Human Monitors Keep Dozing Off!

As Ford has been developing self-driving cars, it has noticed a problem during test drives: Engineers monitoring the robot rides are dozing off. Company researchers have tried to roust the engineers with bells, buzzers, warning lights, vibrating seats and shaking steering wheels. They've even put a second engineer in the vehicle to keep tabs on his human counterpart. No matter -- the smooth ride was just too lulling and engineers struggled to maintain "situational awareness," said Raj Nair, product development chief for the Dearborn, Michigan-based automaker.

<https://t.co/SVG8Z35dIO>

Lessons Learned from Gitlab's Massive Backup Failure

GitLab.com, a multi-million-dollar startup, lost over 300GB of data after a failed backup process. This is just one of the cases that show us that backups don't matter if you can't restore and that IT administrators need to make a solid habit out of testing their backups and recovery capacity.

<https://t.co/wxzZghkrvg>

IT glitch at NASA led to fire

A security patch that shut down monitoring equipment in a large NASA engineering oven resulted in a fire that destroyed spacecraft hardware inside it. Since the computer reboot to accommodate the software upgrade also crippled fire alarm activation, the fire in the oven wasn't discovered for three and a half hours. This is just one example of how a lack of coordination between IT and industrial control systems can wreak havoc.

<https://t.co/5EXsx7bMC9>

Particles from outer space are wreaking low-grade havoc on personal electronics

When your computer crashes, and you get the dreaded blue screen; or your smartphone freezes, and you have to go through the time-consuming process of a reset, most likely you blame the manufacturer: Microsoft or Apple or Samsung. In many instances, however, these operational failures may be caused by the impact of electrically charged particles generated by cosmic rays that originate outside the solar system.

<https://t.co/RQHR8GqGyb>

Stopping financial crashes with physics

In August 2012, an incomplete software update at Knight Capital caused around 4 million rogue stock-market orders to be executed. Over the next 45 minutes, some \$457 million was lost as a result. Whilst this event was caused by a combination of a software glitch and human error, it's very possible that a cyber-attack could have the same effect. Finsec - the cyber-security of financial systems – is a thematic focus for the Institute for Security Science and Technology. We caught up with Vincent Sebag to hear about his recent MSc. project at the ISST, applying econophysics to finsec.

<https://t.co/xlyRRUXJlg>

Suncorp system upgrade causes cash to disappear

A problem with an upgrade to Suncorp Bank's systems in February caused money to disappear from customer accounts and sent others into overdraft. The unspecified glitch occurred during an upgrade to Suncorp's core banking platform.

<https://t.co/K9FUHhy78b>

Why Wind Turbines Should Talk to Each Other

A wind turbine spinning its blades in a valley in southeast India asks a turbine on a plain in Iowa if it should slow down or speed up its rotation. Sound like the stuff of science fiction? It's not. GE has been developing software, sensors and networking technology that enable wind turbines to talk to each other not only within the confines of a particular wind farm but even across the planet.

<https://t.co/tfMFnUBzys>

Four Challenges to Building A Reliable Cloud Telephony System In India

Cloud telephony is emerging as one of the fastest-growing technologies of the decade. Its low setup cost and easy scalability makes it an easy choice for every business that must deal with customers – *every business*. Both the Internet and legacy telephone systems come with unique problems of their own, right from dependency on weather conditions to challenges with policy regulations.

<https://t.co/3h3rzRNWBh>

The science of uptime: How big websites manage huge traffic loads 24x7

Websites like YouTube, Facebook and other popular names need to be running 24 x 7 and must keep up with demand for their billions of users worldwide. The time websites spend up and running is called "uptime." Maintaining uptime can be a challenge especially for many of these top high-traffic websites. It requires significant computing power to process all the requests and data being transferred from across the globe. Here are some of the things that help huge websites remain up.

<https://t.co/cOXUZw37lb>

Cameroon's Internet outage is damaging the country's economy

On January 17, around 20 percent of the people in Cameroon found their Internet blocked. Now March, the connections remain down; and the outage is generally believed to be politically motivated. But the country must face up to a cold, hard fact: if it's blocking its Internet, it's damaging its own economy.

<https://t.co/Y9lro2VUHg>

Weather Service suffers 'catastrophic' outage, stops sending forecasts, warnings

On a day when a blizzard was pasting Maine and Northern California faced a dire flooding threat, several of the National Weather Service's primary systems for sending out alerts to the public failed for nearly three hours. Two core routers for transmitting information from the Weather Service offices out to satellites, which beam the information back to public service providers, stopped working. Both the primary and the backup failed.

<http://bit.ly/2o0N1iT>

Extreme Solar Blackouts May Cost U.S. More Than \$40 Billion Daily

The daily U.S. economic cost from solar storm-induced electricity blackouts could be in the tens of billions of dollars, with more than half the loss from indirect costs outside the blackout zone, according to a new study published in a science journal. Previous studies have focused on direct economic costs within the blackout zone, failing to take into account indirect domestic and international supply chain loss from extreme space weather, according to this latest report.

<https://t.co/n0zwe7PR9E>

Earth's magnetic field 'could be about to flip - and it's now overdue'

Scientists have said that the Earth's magnetic poles could flip for the first time in 786,000 years – and that the event is now 'overdue'. Compasses would point south – and the Earth's power grid could collapse. But before you head for the doomsday bunkers, it's probably not going to happen right away – with scientists predicting a 'flip' sometime in the next 2,000 years.

<https://t.co/SHhH3h9bHy>

Brussels Grid Collapses – Solar Batteries No Help

Large parts of Brussels were without power the evening of 9 February after the City Grid ceased to function at about 10.45 pm local time. But buildings with solar panels on rooftops were not able to maintain their own power supply because the panels are all feeding into to the grid. Even solar panels with battery backup were unable to maintain a supply because EU regulations require a cutoff of the battery supply in the event of a grid outage.

<https://t.co/yI5jFQR1S3>

Blackouts after DBIS goes down

Large parts of Guyana were plunged into several hours of blackout in mid-February after the Demerara Berbice Interconnected System (DBIS) went down. A spokesperson from the Guyana Power and Light confirmed that the DBIS went down around 8 pm and resulted in a total power outage.

<https://t.co/rIVEEVY4Wu>

Booted up in 1993, this server still runs -- but not for much longer

In 1993, U.S. President Bill Clinton was in the first year of his presidency, Windows NT 3.1 and Jurassic Park were both released, the North American Free Trade Agreement (NAFTA) was signed, and Phil Hogan, an IT application architect, booted up a brand-new Stratus Technologies fault tolerant server. A lot has changed in 24 years, but one thing hasn't: The Stratus server is still in operation and Hogan -- who works at steel products maker Great Lakes Works EGL in Dearborn Mich. -- continues to keep it that way.

<https://t.co/xNv3KrcAK2>

From <http://SHADOWBASESOFTWARE.COM> - "Improving Availability via Staggered Systems Part 1: MTTF – Mean Time to Failure"

The availability of a pair of redundant systems can be significantly enhanced via a simple expedient. Simply stagger their starting times. In this way, the time corresponding to the peak probability of failure of one system will not align with the time corresponding to the peak probability of failure of the other system. When one system is likely to fail, the other system is likely to survive. In this pair of papers, we delve into the theory behind this concept. In the first part of this series, we point out a fallacy in classic availability theory.

<http://bit.ly/2kPbRDE>

Fair Warning! There's Math Ahead. "The Fallacy of Classic Availability Theory" by the Digest's Dr. Bill Highleyman

Classic availability theory is flawed in that the expected time to a system failure does not change with time. Clearly, as time goes on, the expected time to system failure should shorten. This flaw is corrected with the concept of Mean Time to Failure (MTTF). MTTF can be used to determine the impact on the availability of various redundant system configurations.

<https://t.co/NVaKp0I5DJ>

How environmentally-conscious data centres are increasing network uptime

As the need for new data centres rises in lockstep with concerns for their environmental impacts, engineers at a growing roster of companies are stretching their imaginations to entertain out-of-the-box strategies for saving power and keeping their data centres cool. As a result, new facilities are being built in increasingly exotic locales and with progressively innovative designs. At the same time, companies recognise the importance of making sure that these non-traditional data centres – which, by their nature, are often located at a distance from operational centres – remain remotely and reliably accessible in order to ensure uninterrupted monitoring, management, and ultimately uptime.

<https://t.co/UdITQtHIJL>

Building High Availability for Industrial and Embedded Systems

High Availability (HA) is not just for the data center. While the principles of achieving extreme uptime have been honed by enterprise IT teams, it's just as important for industrial and embedded applications, which are often deployed in mission-critical environments. By understanding and leveraging HA principles perfected in the enterprise environment, industrial and embedded servers can be made more robust, reliable, and resilient.

<https://t.co/dvetexr6qC>

ATO admits it's baffled by outages but upgrades to latest HPE hardware anyway

The Australian Taxation Office has admitted it still does not know why its online systems crashed again in early late January/early February, causing a four-day systems outage that led accountants to call for compensation. Speaking to The Australian Financial Review after most of the systems were back up online, ATO chief information officer Ramez Katf said it would jettison the Hewlett-Packard Enterprise data storage system that first failed last December and would expand the scope of an ongoing review into the first crash to include the latest one.

<https://t.co/bLoCaNwwwWi>

Blackout leaves at least 40,000 people without power in South Australia

South Australia is currently experiencing some of the highest demands for electricity ever recorded as the state swelters through a heatwave. In a process called load-shedding, the Australian Energy Market Operator (AEMO) deliberately is shutting off power to certain suburbs temporarily in order to prevent a more severe and prolonged outage.

<https://t.co/6ihmND3D0x>

Trident launch failure could have been due to faulty IT system

The nuclear-warhead carrying Trident missile that veered towards the US instead of its intended target after launch could have suffered from an IT systems glitch. The missile was launched from HMS Vengeance, which had just completed an IT systems upgrade. It is highly possible that an error or misconfiguration of this new equipment or a bug in the new IT systems was to blame for the failure. The missile in question was unarmed, but the consequences of an armed nuclear missile veering off course are potentially disastrous.

<https://t.co/LUkDWMjAjz>

For better or worse, the UK's nuclear submarines still run Windows XP

Britain owns four missile submarines: The HMS Vanguard, the Victorious, the Vigilant, and the Vengeance. They patrol the oceans to protect its citizens against a surprise nuclear attack. And while it is reassuring to know what the country does in order to protect the land, the scary part is that each submarine runs Microsoft's first commercial NT-based operating system, installed since 2008.

<https://t.co/kyhWI2eDc8>

Fatigue, Spam, and Lack of Backups Take Down <http://GitLab.com>

The production data loss and hours of downtime at GitLab is an unfortunate and fascinating story about how little things, from spam to engineer fatigue, can coalesce into something more catastrophic. Anecdotes started to trickle in on January 31st, but a single tweet confirmed that something was amiss at GitLab.com: "We accidentally deleted production data and might have to restore from backup."

<https://t.co/BKL86qcvIm>

Businesses experiencing an epidemic of certificate-related outages

Four in five businesses suffered certificate-related outages last year. Almost two-thirds (64%) said their organisations could not respond to a certificate-related security event in six hours or less. The leading cause of outages of this kind is that companies are simply missing the expiry date on their certificates and have no processes in place to automate renewal.

<https://t.co/xmDbFDhmp6>

Surviving a cloud-based disaster recovery plan

Traditional corporate disaster-recovery programs consider primarily on-premise problems such as power outages, server failures, and data loss. These days, disaster-recovery plans ought to (but frequently do not) include situations where you can't rely on *any* kind of connectivity to the Internet. Even when the company has a plan to address each element that can go wrong, organizations rarely test those scenarios to make sure the backup actually works.

<https://t.co/5WutYGIInGB>

How do you build Africa's newest tech ecosystem when the government shuts the Internet down?

Since mid-January, Cameroon's northwest and southwest regions have had no Internet—affecting businesses like Internet cafes, microfinance institutions, and money transfer agencies. The Internet shutdown has also had an adverse effect on the budding tech industry in the country. Known as the Silicon Mountain, a play on Silicon Valley and Mount Cameroon, most of these start-ups are located in Buea, a small town that has been described as the country's home of innovation and Africa's next tech hub.

<https://t.co/UNWB7IDKJj>

Telstra suffers major outage due to exchange fire

Australia's largest telecoms company, Telstra, recently suffered a major outage throughout its network after a core exchange near Sydney caught fire. The issues affected fixed line and mobile services across New South Wales, Victoria, Queensland, South Australia and Western Australia. The incident disrupted nearly 30 percent of overall Telstra traffic, causing delays for flights and train journeys. Some government offices and schools had to be closed, and digital businesses lost their revenue. A corruption of the internal IT systems also meant that some of the text messages sent by Telstra customers were delivered to wrong recipients.

<https://t.co/kHHYuOTWpx>

"The Best Disaster Recovery Advice for Failover and Failback Success"

Your organization's disaster recovery plan (DRP) is an essential component of its business continuity plan (BCP) and must include not only the physical setups for data replication and protection at a secondary location, but also effective failover and fallback strategies to ensure you will be able to execute your DRP successfully.

<https://t.co/rdf8PKKeGQo>

Tesla's New Powerpack Grid in LA Could Be the Future of Power Plants

Tesla recently installed a Powerpack utility grid that has enough energy to power 2,500 homes all day and reduce the Los Angeles area's reliance on fossil fuels at the same time. Tesla has posted a video on Twitter of a grid it has installed at Southern California Edison's substation in the Los Angeles Basin. Installing all of these sleek white Powerpacks took 94 days.

<https://t.co/V9LdpvYCJC>

Innovative Direct-Current Microgrids to Solve India's Power Woes

In the industrialized world, the power grid is so reliable that we take it for granted. But in India, where blackouts are a sad fact of daily life, being connected to the grid is no guarantee of reliable electricity. Roughly a quarter of a billion Indians, or one-fifth of the population, live without access to any electricity at all. In recent years, the Indian government has invested heavily in electricity generation (including solar- and wind-power plants), state-of-the-art high voltage transmission lines, and a multitude of household electrification projects. And yet these efforts have made only a modest dent in the problem.

<https://t.co/ScDczLdFUr>

Atomic clocks failure: ISRO decides to replace IRNSS-1A with backup navigation satellite

Three atomic clocks on board ISRO's IRNSS-1A satellite have failed. The Indian Space Research Organisation will be launching the backup navigation satellite this year, and it will replace the IRNSS-1A.

<https://t.co/FulwHVgUky>

GitLab suffers major backup failure after data deletion incident

Khosla Ventures and Y Combinator-backed open source Git repository [GitLab](#) went offline in early February after suffering what appears to be a major backup restoration failure after accidentally deleting production data.

<https://t.co/Nje4OfDJ48>

Delta Outage Spotlights Technology Risks

Delta's computer outage on 29 January was over by midnight, but its effects extended well into the week. Such outages can be costly. Jim Corridore, an analyst at CFRA Research, says that Delta's computer outage puts a "spotlight on risks of airline technology infrastructure, much of which is old and patched with differing systems." He said that airlines build new programming over old software, especially after a merger, when computer languages may differ. Programmers' assumptions about how software will work are sometimes wrong. While large companies such as Delta would have fewer outages with more testing of their systems, this is an expensive proposition.

<https://t.co/f7HInslqDC>