

Amazon Elastic Compute Cloud

Subjects: Computer Science, Theory & Methods

Contributor: HandWiki Xu

Amazon Elastic Compute Cloud (EC2) is a part of Amazon.com's cloud-computing platform, Amazon Web Services (AWS), that allows users to rent virtual computers on which to run their own computer applications. EC2 encourages scalable deployment of applications by providing a web service through which a user can boot an Amazon Machine Image (AMI) to configure a virtual machine, which Amazon calls an "instance", containing any software desired. A user can create, launch, and terminate server-instances as needed, paying by the second for active servers – hence the term "elastic". EC2 provides users with control over the geographical location of instances that allows for latency optimization and high levels of redundancy. In November 2010, Amazon switched its own retail website platform to EC2 and AWS.

Keywords: latency optimization ; cloud-computing ; amazon

1. History

Amazon announced a limited public beta test of EC2 on August 25, 2006,^[1] offering access on a first-come, first-served basis. Amazon added two new instance types (Large and Extra-Large) on October 16, 2007.^[2] On May 29, 2008, two more types were added, High-CPU Medium and High-CPU Extra Large.^[3] There were twelve types of instances available.

Amazon added three new features on March 27, 2008,^[5] static IP addresses, availability zones, and user selectable kernels. On August 20, 2008, Amazon added Elastic Block Store (EBS)^[6] This provides persistent storage, a feature that had been lacking since the service was introduced.

Amazon EC2 went into full production when it dropped the beta label on October 23, 2008. On the same day, Amazon announced the following features:^[7]

- a service level agreement for EC2,
- Microsoft Windows in beta form on EC2,
- Microsoft SQL Server in beta form on EC2,
- plans for an AWS management console, and
- plans for load balancing, autoscaling, and cloud monitoring services.^[7]

These features were subsequently added on May 18, 2009.^[8]

Amazon EC2 was developed mostly by a team in Cape Town, South Africa led by Chris Pinkham.^{[9][10]} Pinkham provided the initial architecture guidance for EC2 and then built the team and led the development of the project along with Willem van Biljon.

2. Instance Types

Initially, EC2 used Xen virtualization exclusively.^[11] However, on November 6, 2017, Amazon announced^[12] the new C5 family of instances that were based on a custom architecture around the KVM hypervisor, called Nitro.^{[13][14]} Each virtual machine, called an "instance", functions as a virtual private server. Amazon sizes instances based on "Elastic Compute Units". The performance of otherwise identical virtual machines may vary.^[15] On November 28, 2017, AWS announced a bare-metal instance type offering marking a remarkable departure from exclusively offering virtualized instance types.^[16]

As of January 2019, the following instance types were offered:^[17]

- General Purpose: A1, T3, T2, M5, M5a, M4, T3a
- Compute Optimized: C5, C5n, C4
- Memory Optimized: R5, R5a, R4, X1e, X1, High Memory, z1d
- Accelerated Computing: P3, P2, G3, F1
- Storage Optimized: H1, I3, D2

(As of April 2018), the following payment methods by instance were offered:^[18]

- On-demand: pay by the hour without commitment.
- Reserved: rent instances with one-time payment receiving discounts on the hourly charge.
- Spot: bid-based service: runs the jobs only if the spot price is below the bid specified by bidder. The spot price is claimed to be supply-demand based, however a 2011 study concluded that the price was generally not set to clear the market, but was dominated by an undisclosed reserve price.^{[19][20]}

2.1. Cost



Amazon CTO Werner Vogels announcing price reductions at AWS Summit 2013 NYC.

<https://handwiki.org/wiki/index.php?curid=1453518>

(As of April 2018), Amazon charged about \$0.0058/hour (\$4.176/month) for the smallest "Nano Instance" (t2.nano) virtual machine running Linux or Windows. Storage-optimized instances cost as much as \$4.992/hour (i3.16xlarge). "Reserved" instances can go as low as \$2.50/month for a three-year prepaid plan.^{[21][22][23]} The data transfer charge ranges from free to \$0.12 per gigabyte, depending on the direction and monthly volume (inbound data transfer is free on all AWS services^[24]).

2.2. Free Tier

(As of December 2010) Amazon offered^[25] a bundle of free resource credits to new account holders. The credits are designed to run a "micro" sized server, storage (EBS), and bandwidth for one year.^[26] Unused credits cannot be carried over from one month to the next.^[27]

2.3. Reserved Instances

Reserved instances enable EC2 or RDS service users to reserve an instance for one or three years. The corresponding hourly rate charged by Amazon to operate the instance is 35-75% lower than the rate charged for on-demand instances.^[28] Reserved instances can be purchased in three different ways: All Upfront, Partial Upfront and No Upfront. The different purchase options allow for different structuring of payment models. In September 2016, AWS announced several enhancements to Reserved instances, introducing a new feature called scope and a new reservation type called a Convertible.^[29] In October 2017, AWS announced the allowance to subdivide the instances purchased for more flexibility^[30]

2.4. Spot Instances

Cloud providers maintain large amounts of excess capacity they have to sell or risk incurring losses.^[31] Amazon EC2 Spot instances are spare compute capacity in the AWS cloud available at up to 90% discount compared to On-Demand prices.^[20] As a trade-off, AWS offers no SLA on these instances and customers take the risk that it can be interrupted with only two minutes of notification when Amazon needs the capacity back. Researchers from the Israeli Institute of Technology found that "they (Spot instances) are typically generated at random from within a tight price interval via a dynamic hidden reserve price".^[31] Some companies, like Spotinst, are using machine learning to predict spot interruptions up to 15 minutes in advance.^{[32][33]}

2.5. Savings Plans

In November 2019, Amazon announced Savings Plans.^[34] Savings Plans are an alternative to Reserved Instances^[35] that come in two different plan types: Compute Savings Plans and EC2 Instances Savings Plans. Compute Savings Plans allow an organisation to commit to EC2 and Fargate usage with the freedom to change region, family, size, availability zone, OS and tenancy inside the lifespan of the commitment. EC2 Instance Savings plans provide the lowest prices but are less flexible meaning a user must commit to individual instance families within a region to take advantage, but with the freedom to change instances within the family in that region.^[36] AWS Savings plans are tricky if you don't know how they work. You buy Savings Plans based on hourly commitment. AWS uses the Cost Explorer to automatically calculate how that commitment will look like as a monthly charge on your AWS bill.^[37]

3. Features

3.1. Operating Systems

When it launched in August 2006, the EC2 service offered Linux and later Sun Microsystems' OpenSolaris and Solaris Express Community Edition. In October 2008, EC2 added the Windows Server 2003 and Windows Server 2008 operating systems to the list of available operating systems.^{[38][39]} In March 2011, NetBSD AMIs became available.^[40] In November 2012, Windows Server 2012 support was added.^[41]

Since 2006, Colin Percival, a FreeBSD developer and Security Officer, solicited Amazon to add FreeBSD.^[42] In November 2012, Amazon officially supported running FreeBSD in EC2.^{[43][44][45]} The FreeBSD/EC2 platform is maintained by Percival^[46] who also developed the secure deduplicating Amazon S3-cloud based backup service Tarsnap.^[47]

Amazon has their own Linux distribution based on the Fedora and Red Hat Enterprise Linux as a low cost offering known as the [Amazon Linux AMI](#). Version 2013.03 included:^[48]

- Linux kernel version 3.4.34
- Java OpenJDK Runtime Environment (IcedTea6 1.11.4)
- GNU Compiler Collection gcc.x86_64 4.4.6-3.45.amzn1

On November 30, 2020, Amazon announced that it would be adding macOS to the EC2 service. Initial support was announced for macOS Mojave and macOS Catalina running on Mac mini.^[49]

3.2. Persistent Storage

An EC2 instance may be launched with a choice of two types of storage for its boot disk or "root device." The first option is a local "instance-store" disk as a root device (originally the only choice). The second option is to use an EBS volume as a root device. Instance-store volumes are temporary storage, which survive rebooting an EC2 instance, but when the instance is stopped or terminated (e.g., by an API call, or due to a failure), this store is lost.

The Amazon Elastic Block Store (EBS) provides raw block devices that can be attached to Amazon EC2 instances. These block devices can then be used like any raw block device. In a typical use case, this would include formatting the device with a filesystem and mounting it. In addition, EBS supports a number of advanced storage features, including snapshotting and cloning. EBS volumes can be up to 16TB in size. EBS volumes are built on replicated storage, so that the failure of a single component will not cause data loss. EBS was introduced to the general public by Amazon in August 2008.^[6]



EBS volumes provide persistent storage independent of the lifetime of the EC2 instance, and act much like hard drives on a real server. More accurately, they appear as block devices to the operating system that are backed by Amazon's disk arrays. The OS is free to use the device however it wants. In the most common case, a file system is loaded and the volume acts as a hard drive. Another possible use is the creation of RAID arrays by combining two or more EBS volumes. RAID allows increases of speed and/or reliability of EBS. Users can set up and manage storage volumes of sizes from 1GB to 16TB. The volumes support snapshots, which can be taken from a GUI tool or the API. EBS volumes can be attached or detached from instances while they are running, and moved from one instance to another.^[50]

Simple Storage Service (S3) is a storage system in which data is accessible to EC2 instances, or directly over the network to suitably authenticated callers (all communication is over HTTP). Amazon does not charge for the bandwidth for communications between EC2 instances and S3 storage "in the same region." Accessing S3 data stored in a different region (for example, data stored in Europe from a US East Coast EC2 instance) will be billed at Amazon's normal rates.

S3-based storage is priced per gigabyte per month. Applications access S3 through an API. For example, Apache Hadoop supports a special s3: filesystem to support reading from and writing to S3 storage during a MapReduce job. There are also S3 filesystems for Linux, which mount a remote S3 filestore on an EC2 image, as if it were local storage. As S3 is not a full POSIX filesystem, things may not behave the same as on a local disk (e.g., no locking support).

3.3. Elastic IP Addresses



Amazon's elastic IP address feature is similar to static IP address in traditional data centers, with one key difference. A user can programmatically map an elastic IP address to any virtual machine instance without a network administrator's help and without having to wait for DNS to propagate the binding. In this sense an Elastic IP Address belongs to the account and not to a virtual machine instance. It exists until it is explicitly removed, and remains associated with the account even while it is associated with no instance.^[51]

3.4. Amazon CloudWatch



Amazon CloudWatch is a web service that provides real-time monitoring to Amazon's EC2 customers on their resource utilization such as CPU, disk, network and replica lag for RDS Database replicas.^[52] CloudWatch does not provide any memory, disk space, or load average metrics without running additional software on the instance. Since December 2017 Amazon provides a CloudWatch Agent for Windows and Linux operating systems included disk and previously not available memory information,^[53] previously Amazon provided example scripts for Linux instances to collect OS information.^{[54][55]} The data is aggregated and provided through AWS management console. It can also be accessed through command line tools and Web APIs,^[56] if the customer desires to monitor their EC2 resources through their enterprise monitoring software. Amazon provides an API which allows clients to operate on CloudWatch alarms.^[57]

The metrics collected by Amazon CloudWatch enables the auto-scaling feature to dynamically add or remove EC2 instances.^[58] The customers are charged by the number of monitoring instances.

Since May 2011, Amazon CloudWatch accepts custom metrics^[59] that can be submitted programmatically via Web Services API and then monitored the same way as all other internal metrics, including setting up the alarms for them, and since July 2014 Cloudwatch Logs service is also available.^[60]


Basic Amazon CloudWatch is included in Amazon Free Tier service.

3.5. Automated Scaling




Amazon's auto-scaling feature of EC2 allows it to automatically adapt computing capacity to site traffic.^[61] The schedule-based (e.g. time-of-the-day) and rule-based (e.g. CPU utilization thresholds) auto scaling mechanisms are easy to use and efficient for simple applications. However, one potential problem is that VMs may take up to several minutes to be ready to use, which are not suitable for time critical applications. The VM startup time is dependent on image size, VM type, data center locations, etc.^[62]


4. Pricing



Linux cloud hosting: performance index below 16 (< 2CPU, 8GB RAM). <https://handwiki.org/wiki/index.php?curid=1284120>



Linux cloud hosting: performance index 16 - 1040 (2CPU, 8GB RAM - 16CPU, 65GB RAM). <https://handwiki.org/wiki/index.php?curid=1778731>



Linux cloud hosting: performance index ABOVE 1040 (> 16CPU, 65GB RAM). <https://handwiki.org/wiki/index.php?curid=1141457>

Generally, Amazon EC2 is priced on a per instance / per hour basis. However, any instance can be rented on a per month basis as well. In such case, Reserved and Spot Instances pricing can be applied, resulting in a significant discount. Instances are priced depending on their "size", namely how much CPU and RAM are included.

Amazon EC2 price varies from \$2.5 per month for "nano" instance with 1 vCPU and 0.5 GB RAM on board to "xlarge" type of instances with 32 vCPU and 488 GB RAM billed up to \$3997.19 per month.

The charts above show how Amazon EC2 pricing is compared to similar Cloud Computing services: Microsoft Azure, Google Cloud Platform, Kamatera, and Vultr.^[63]

Amazon EC2 compared to similar cloud computing services

	Amazon EC2	Microsoft Azure	Google Cloud Platform	Kamatera	Vultr
1vCPU 0.5GB RAM	\$3.29				\$2.5
1vCPU 0.75GB RAM		\$14.88			
1vCPU 1GB RAM	\$6.83			\$11	\$5

1vCPU 1.75GB RAM		\$44.64		
1vCPU 2GB RAM	\$13.14		\$17	\$10
1vCPU 3.75GB RAM			\$24.27	
2vCPU 3.5GB RAM		\$89.88		
2vCPU 4GB RAM			\$41	\$20
2vCPU 7.5GB RAM			\$48.55	
2vCPU 8GB RAM	\$52.56		\$61	
4vCPU 7GB RAM		\$178.56		
4vCPU 8GB RAM			\$86	\$40
4vCPU 15GB RAM			\$97.09	
4vCPU 15GB RAM	\$134		\$134	
6vCPU 16GB RAM			\$159	\$80
8vCPU 14GB RAM		\$357.12		
8vCPU 16GB RAM			\$184	
8vCPU 30GB RAM			\$194.18	
8vCPU 32GB RAM	\$219.64		\$280	\$160
8vCPU 49GB RAM			\$328	
8vCPU 56GB RAM		\$744		
16vCPU 32GB RAM	\$412.53			
8vCPU 65GB RAM			\$408	
12vCPU 65GB RAM			\$626	
16vCPU 32GB RAM			\$388.36	\$746
16vCPU 65GB RAM				\$320
20vCPU 65GB RAM				\$826
12vCPU 112GB RAM		\$1339.20		
16vCPU 112GB RAM		\$1450.80		
20vCPU 98GB RAM				\$986
36vCPU 60GB RAM	\$825.06			
20vCPU 131GB RAM				\$1146
32vCPU 120GB RAM			\$776.72	
16vCPU 224GB RAM		\$1935.1		
20vCPU 196GB RAM				\$1466
20vCPU 262GB RAM				\$1786
24vCPU 224GB RAM		\$2678.40		
64vCPU 240GB RAM			\$1553.44	
32vCPU 488GB RAM	\$3997.19			

5. Reliability

To make EC2 more fault-tolerant, Amazon engineered *Availability Zones* that are designed to be insulated from failures in other availability zones. Availability zones do not share the same infrastructure. Applications running in more than one availability zone can achieve higher availability.^[64]

EC2 provides users with control over the geographical location of instances that allows for latency optimization and high levels of redundancy. For example, to minimize downtime, a user can set up server instances in multiple zones that are insulated from each other for most causes of failure such that one backs up the other.

Higher-availability database services, like Amazon Relational Database Service run separately from EC2 instances.

6. Issues

In early July 2008, the anti-spam organizations Outblaze and Spamhaus.org began blocking Amazon's EC2 address pool due to problems with the distribution of spam and malware.^[65]

On December 1, 2010, Amazon pulled its service to WikiLeaks after coming under political pressure in the US.^[66] The Internet group Anonymous attempted to attack EC2 in revenge; however, Amazon was not affected by the attack.^[67]

Amazon's websites were temporarily offline on December 12, 2010, although it was initially unclear if this was due to attacks or a hardware failure. An Amazon official later stated that it was due to a hardware failure.^[68]

Shortly before 5 am ET on April 21, 2011, an outage started at EC2's Northern Virginia data center that brought down several websites, including Foursquare, Springpad, Reddit, Quora, and Hootsuite.^{[69][70][71]} Specifically, attempts to use Amazon's elastic-disk and database services hung, failed, or were slow. Service was restored to some parts of the datacenter (three of four "availability zones" in Amazon's terms) by late afternoon Eastern time that day;^[72] problems for at least some customers were continuing as of April 25.^[73] 0.07% of EBS volumes in one zone have also been lost; EBS failures were a part of normal operation even before this outage and were a risk documented by Amazon,^[74] though the number of failures and the number of *simultaneous* failures may find some EC2 users unprepared.

On Sunday August 6, 2011, Amazon suffered a power outage in one of their Ireland availability zones.^[75] Lightning was originally blamed for the outage; however, on August 11, Irish energy supplier ESB Networks dismissed this as a cause, but at time of writing, could not confirm what the cause of the problem was.^[76] The power outage raised multiple questions regarding Amazon's EBS infrastructure, which caused several bugs in their software to be exposed. The bugs resulted in some customers' data being deleted when recovering EBS volumes in a mid-write operation during the crash.^[77]

August 8, 2011 saw another network connectivity outage of Amazon's Northern Virginia datacenter, knocking out the likes of Reddit, Quora, Netflix and FourSquare.^[78] The outage lasted around 25 minutes.

Another Northern Virginia datacenter outage occurred on October 22, 2012, from approximately 10 am to 4 pm PT. Edmodo, Airbnb, Flipboard, Reddit, and other customers were affected. Anonymous claimed responsibility, but Amazon denied this assertion.^[79]

References

1. Barr, Jeff (August 25, 2006). "Amazon EC2 Beta". Amazon Web Services Blog. http://aws.typepad.com/aws/2006/08/amazon_ec2_beta.html.
2. Barr, Jeff (October 16, 2007). "Amazon EC2 Gets More Muscle". <http://aws.typepad.com/aws/2007/10/amazon-ec2-gets.html>.
3. Barr, Jeff (May 29, 2008). "More EC2 power". Amazon Web Services Blog. <http://aws.typepad.com/aws/2008/05/more-ec2-power.html>.
4. Amazon EC2 Instances. [Aws.amazon.com](http://aws.amazon.com/ec2/instance-types/). Retrieved on 2013-08-09. <http://aws.amazon.com/ec2/instance-types/>
5. Barr, Jeff (March 27, 2008). "New EC2 Features". Amazon Web Services Blog. <http://aws.typepad.com/aws/2008/03/new-ec2-feature.html>.
6. Barr, Jeff (August 20, 2008). "Amazon EBS (Elastic Block Store) - Bring Us Your Data". Amazon Web Services Blog. <http://aws.typepad.com/aws/2008/08/amazon-elastic.html>.
7. Barr, Jeff (August 23, 2008). "Big Day For EC2". Amazon Web Services Blog. <http://aws.typepad.com/aws/2008/10/big-day-for-ec2.html>.
8. Barr, Jeff (May 18, 2009). "New Features for EC2". Amazon Web Services Blog. <http://aws.typepad.com/aws/2009/05/new-aws-load-balancing-automatic-scaling-and-cloud-monitoring-services.html>.
9. "Nimbula puffs up 'cloud operating system'". June 25, 2010. https://www.theregister.co.uk/2010/06/25/nimbula_cloud_os/.

10. "Amazon's early efforts at cloud computing? Partly accidental". June 17, 2010. <http://itknowledgeexchange.techtarget.com/cloud-computing/2010/06/17/amazons-early-efforts-at-cloud-computing-partly-accidental/>.
11. "Citrix and Amazon Web Services (AWS)". citrix.com. <http://www.citrix.com/global-partners/amazon-web-services.html>.
12. "Now Available – Compute-Intensive C5 Instances for Amazon EC2". 6 November 2017. <https://aws.amazon.com/blogs/aws/now-available-compute-intensive-c5-instances-for-amazon-ec2/>.
13. "Amazon EC2 FAQs". <https://aws.amazon.com/ec2/faqs/#compute-optimized>.
14. "AWS re:Invent 2017: NEW LAUNCH! Amazon EC2 Bare Metal Instances (CMP330)". https://www.youtube.com/watch?v=o9_4uGvbnk.
15. J. Dejun, G. Pierre and C.-H. Chi. EC2 Performance Analysis for Resource Provisioning of Service-Oriented Applications. In Proceedings of the Workshop on Non-Functional Properties and SLA Management in Service-Oriented Computing, November 2009. http://www.globule.org/publii/EPARPSOA_nfpsla2009.html
16. Lardinois, Frederic (November 28, 2017). "AWS launches bare metal instances". TechCrunch. <https://techcrunch.com/2017/11/28/aws-launches-bare-metal-instances/>.
17. "Amazon EC2 Instance Types – Amazon Web Services (AWS)" (in en-US). <https://aws.amazon.com/ec2/instance-types/>.
18. "Amazon EC2 Pricing". <http://aws.amazon.com/ec2/pricing/>.
19. Ben-Yehuda, Orna Agmon; Ben-Yehuda, Muli; Schuster, Assaf; Tsafir, Dan (2011). Deconstructing Amazon EC2 Spot Instance Pricing. The Taub Faculty of Computer Science, Technion (Technical report). Retrieved 2020-12-27. <https://www.cs.technion.ac.il/users/wwwb/cgi-bin/tr-info.cgi/2011/CS/CS-2011-09>
20. Wu, Xiaohu; Pellegrini, Francesco De; Gao, Guanyu; Casale, Giuliano (2019-12-24). "A Framework for Allocating Server Time to Spot and On-Demand Services in Cloud Computing". ACM Transactions on Modeling and Performance Evaluation of Computing Systems (Association for Computing Machinery (ACM)) 4 (4): 1–31. doi:10.1145/3366682. ISSN 2376-3639. <https://dx.doi.org/10.1145%2F3366682>
21. \$109 for a three-year Heavy Utilization Reserved t2.micro Instances reservation amortized over thirty-six months plus one month at \$0.002/hour.
22. Amazon Elastic Compute Cloud (Amazon EC2), Cloud Computing Servers. Aws.amazon.com (2014-07-01). Retrieved on 2014-07-01. <http://aws.amazon.com/ec2/#pricing>
23. EC2 selector tool. Simplified and comprehensive EC2 price list updated in real time. <https://ec2-selector.ec2-cli.com>
24. Barr, Jeff (June 29, 2011). "AWS Lowers its Pricing Again! – No Inbound Data Transfer Fees and Lower Outbound Data Transfer for All Services including Amazon CloudFront". Typepad. <http://aws.typepad.com/aws/2011/06/aws-lowers-its-pricing-again-free-inbound-data-transfer-and-lower-outbound-data-transfer-for-all-ser.html>.
25. "Announcing AWS Free Usage Tier". <https://aws.amazon.com/about-aws/whats-new/2010/10/21/announcing-aws-free-usage-tier/>.
26. "AWS Free Tier". <https://aws.amazon.com/free/>.
27. "FAQs". <http://aws.amazon.com/free/faqs/>.
28. "Additional Reserved Instance Options for Amazon EC2". Amazon. December 2011. <http://aws.typepad.com/aws/2011/12/reserved-instance-options-for-amazon-ec2.html>.
29. Santos, Alan (September 29, 2016). "A Quick Start Guide To AWS Regional & Convertible Reservations". CloudHealth Technologies. <https://www.cloudhealthtech.com/blog/quick-start-guide-aws-regional-convertible-reservations>.
30. Hanley Frank, Blair (March 14, 2017). "AWS Follows Google With Reserved Instance Flexibility Changes". PC World. <https://www.pcworld.com/article/3180821/cloud-computing/aws-follows-google-with-reserved-instance-flexibility-changes.html>.
31. Technion - Israel Institute of Technology. "Deconstructing Amazon EC2 Spot Instance Pricing". Technion University. <https://www.cs.technion.ac.il/~ladypine/spotprice-ieee.pdf>.
32. Hall, Susan (December 26, 2017). "Spotinst: Making the Most of Cheaper Excess Compute Capacity". The News Stack. <https://thenewstack.io/spotinst-making-cheaper-excess-compute-capacity/>.
33. Miller, Ron (October 25, 2017). "Spotinst delivers spot cloud infrastructure services at discount prices". TechCrunch. <https://techcrunch.com/2017/10/25/spotinst-delivers-spot-cloud-infrastructure-services-at-discount-prices/>.
34. "Savings Plans – Amazon Web Services". <https://aws.amazon.com/savingsplans/>.

35. "Savings Plans vs Reserved Instances". 7 November 2019. <https://www.gorillastack.com/news/aws-savings-plans-reserved-instances/>.
36. "Savings Plans FAQ | Amazon Web Services". <https://aws.amazon.com/savingsplans/faq/>.
37. Santis, Peter (July 14, 2021). "Running on AWS Savings Plans? You can still reduce your cloud bill". CAST AI. <https://cast.ai/blog/running-on-aws-savings-plans-you-can-still-reduce-your-cloud-bill/>.
38. Shankland, Stephen (October 23, 2008). "Amazon's Linux cloud computing out of beta, joined by Windows". CNET. http://news.cnet.com/8301-17939_109-10073696-2.html.
39. "Amazon Elastic Compute Cloud (EC2) Running Microsoft Windows Server and SQL Server". October 23, 2008. <http://aws.amazon.com/windows/>.
40. Migeon, Jean-Yves (March 13, 2011). "NetBSD now runs under Amazon EC2". NetBSD Blog. http://blog.netbsd.org/tnf/entry/netbsd_now_runs_under_amazon.
41. "Windows Server 2012 Now Available on AWS". November 19, 2012. <http://aws.typepad.com/aws/2012/11/windows-server-2012-now-available-on-aws.html>.
42. Barr, Jeff. "Interview - Live from AWS re:Invent Colin Percival, FreeBSD Developer". Amazon Web Services 28.11.2012. https://www.youtube.com/watch?v=SK7EeV_GD3M.
43. "AWS Marketplace – Additional EC2 Operating System Support (FreeBSD, Debian, CentOS)". Amazon. 23 November 2012. <http://aws.typepad.com/aws/2012/11/aws-marketplace-additional-operating-system-support.html>.
44. Percival, Colin. "How to build FreeBSD/EC2 images". Colin Percival. <http://www.daemonology.net/blog/2014-02-16-FreeBSD-EC2-build.html>.
45. Percival, Colin. "FreeBSD on EC2". Colin Percival. <http://www.daemonology.net/freebsd-on-ec2/>.
46. "Colin Percival". Amazon Web Services, Inc.. <https://aws.amazon.com/marketplace/seller-profile?id=4a931b8a-95f6-4bb2-8b7e-53507851c459>.
47. Percival, Colin. "About Tarsnap". Tarsnap and C. Percival. <http://www.tarsnap.com/about.html>.
48. "Amazon Linux AMI 2013.03 Release Notes". AWS web site. <http://aws.amazon.com/amazon-linux-ami/2013.03-release-notes/>.
49. error
50. "Introduction to EBS Volumes". 1201 Restart.com. December 2010. http://1201restart.com/Articles/Introduction_to_EBS_Volumes.
51. "What is an Elastic IP". June 2, 2022. <https://riyanchristy.goseeq.net/what-is-an-elastic-ip/>.
52. "Working with Read Replicas of MariaDB, MySQL, and PostgreSQL DB Instances - Amazon Relational Database Service". https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReadRepl.html.
53. "Amazon CloudWatch introduces a new CloudWatch Agent with AWS Systems Manager Integration for Unified Metrics and Logs Collection". 14 December 2017. <https://aws.amazon.com/about-aws/whats-new/2017/12/amazon-cloudwatch-introduces-a-new-cloudwatch-agent-with-aws-systems-manager-integration-for-unified-metrics-and-logs-collection/>.
54. "Monitoring Memory and Disk Metrics for Amazon EC2 Linux Instances - Amazon Elastic Compute Cloud". <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/mon-scripts.html>.
55. "AWS Developer Forums: User Profile for Henry@AWS". <https://forums.aws.amazon.com/profile.jspa?userID=202407>.
56. "GorillaStack CloudWatch vs CloudTrail". 2019-09-17. <https://www.gorillastack.com/news/cloudtrail-vs-cloudwatch/>.
57. "Welcome - Amazon CloudWatch". <https://docs.aws.amazon.com/AmazonCloudWatch/latest/APIReference/Welcome.html>.
58. AWS in Action & Wittig (2016), pp. 372–375.
59. "AWS Release Notes". <https://aws.amazon.com/releasenotes/item/>.
60. "Introducing Amazon CloudWatch Logs". <https://aws.amazon.com/about-aws/whats-new/2014/07/10/introducing-amazon-cloudwatch-logs/>.
61. "AWS Auto Scaling" (in en-US). <https://aws.amazon.com/autoscaling/>.
62. Mao, Ming; M. Humphrey (2012). "A Performance Study on the VM Startup Time in the Cloud". Proceedings of 2012 IEEE 5th International Conference on Cloud Computing (Cloud2012): 423. doi:10.1109/CLOUD.2012.103. ISBN 978-1-4673-2892-0. <https://dx.doi.org/10.1109%2FCLOUD.2012.103>

63. "Amazon AWS vs Microsoft Azure vs Google Cloud vs Kamatera vs Vultr - 5 Best Cloud Services" (in en-US). 5 Best Cloud Services. 2017-12-02. <https://5best.cloud/aws-vs-azure-vs-google-cloud/>.
64. Wittig, Andreas; Wittig, Michael (2016). Amazon Web Services in Action. Manning Press. pp. 290–294. ISBN 978-1-61729-288-0.
65. Krebs, Brian (July 1, 2008). "Amazon: Hey Spammers, Get Off My Cloud!". Security Fix (The Washington Post). http://blog.washingtonpost.com/securityfix/2008/07/amazon_hey_spammers_get_off_my.html.
66. MacAskill, Ewen (December 2, 2010). "WikiLeaks website pulled by Amazon after US political pressure". The Guardian (London). <https://www.theguardian.com/media/2010/dec/01/wikileaks-website-cables-servers-amazon>.
67. Weaver, Matthew; Tran, Mark (December 9, 2010). "WikiLeaks cables: Shell, Operation Payback and Assange for the Nobel prize – as it happened". The Guardian. <https://www.theguardian.com/news/blog/2010/dec/09/wikileaks-us-embassy-cables-live-updates>.
68. Arthur, Charles (December 13, 2010). "Amazon says outage in Europe due to hardware failure, not hacking attack". The Guardian. <https://www.theguardian.com/technology/blog/2010/dec/13/amazon-failure-not-hacking-wikileaks>.
69. Pepitone, Julianne (April 21, 2011). "Amazon EC2 outage downs Reddit, Quora". CNN. https://money.cnn.com/2011/04/21/technology/amazon_server_outage/index.htm.
70. Lohr, Steve (April 22, 2011). "Amazon Malfunction Raises Doubts About Cloud Computing". The New York Times. <https://www.nytimes.com/2011/04/23/technology/23cloud.html>.
71. "Cloud storm, EC2 startup takedown". DoesWhat. April 21, 2011. <http://www.doeswhat.com/2011/04/21/cloud-storm-amazons-ec2-takes-down-startups/>.
72. "6:18 pm PDT April 21, 2011, entry, says "all Availability Zones except one have been functioning normally for the past 5 hours."" <http://status.aws.amazon.com/>.
73. "April 25 entry." <http://status.aws.amazon.com/>.
74. Amazon Elastic Block Store (EBS) section "Volume Durability" <http://aws.amazon.com/ebs/>
75. Clark, Jack (August 8, 2011). "Lightning strikes Amazon's European cloud". ZDNet. <http://www.zdnet.co.uk/news/cloud/2011/08/08/lightning-strikes-amazons-european-cloud-40093641/>.
76. Kunert, Paul (August 11, 2011). "Amazon web outage not caused by lightning". The Register. https://www.theregister.co.uk/2011/08/11/esb_web_outage/.
77. Clark, Jack (August 10, 2011). "AWS cloud accidentally deletes customer data". ZDNet. <http://www.zdnet.co.uk/news/cloud/2011/08/10/aws-cloud-accidentally-deletes-customer-data-40093665/>.
78. Musil, Steven (August 9, 2011). "Amazon cloud outage downs Netflix, Quora". CNET. http://news.cnet.com/8301-1023_3-20089866-93/amazon-cloud-outage-downs-netflix-quora/.
79. Tibken, Shara (October 22, 2012). "Amazon cloud outage impacts Reddit, Airbnb, Flipboard". CNET. http://news.cnet.com/8301-1023_3-57537499-93/amazon-cloud-outage-impacts-reddit-airbnb-flipboard/.

Retrieved from <https://encyclopedia.pub/entry/history/show/76279>