Cloud computing

Cloud computer services are computer resources that are accessed remotely, via the internet, but which were once commonly accessed locally. The cloud can be divided into **cloud storage**, in which customer files or other data are stored remotely, and **cloud computing**, involving remote processing. A common characteristic of the latter category is that the provider's resources are shared dynamically among different customers; this distinguishes cloud computing from simple leasing of physical web servers (see below at Cloud Models). Use of the cloud can mean reduced risk of data loss, reduced upfront investment costs and simplified software setup; the cloud has, for example, made it easier to launch high-tech companies by reducing startup costs. On the other hand, cloud users must ensure they understand the privacy and security implications of the cloud, and their agreement with the cloud provider.

Cloud-based resources can vary considerably. Cloud storage may involve storing very large files, such as video, or relatively small files, such as calendar entries. Cloud computing may involve remote databases, web servers and other application servers, or general-purpose computers that are configured by the customer as desired.

Cloud-based services are most often offered by a third party, in which case they are classed as the **public cloud**. Alternatively, site can maintain its own **private** cloud; private cloud services can also be leased. While a self-maintained private cloud means that the owning organization must still handle all the hardware and lower-layer software issues that are avoided by public-cloud users, there can still be significant cost savings due to scale, resource sharing and the standardization of basic hardware building blocks.

Cloud services can be offered from a single site or from a geographically distributed set of sites. Many cloud-storage providers offer limited personal accounts free of charge, though such accounts rarely

have much access to technical support. Business use almost always requires a paid account, together with a very clear understanding of the terms of service.

The most basic form of cloud services is cloud storage. The issue of privacy is usually important, as the cloud provider has access to the customer's content; see below. For large files, such as media files, cloud storage is often done simply because sufficient local storage is not available. For smaller files, cloud storage offers the safety of a secure backup; many mobile-device owners have been greatly relieved to discover, after losing their device, that many of their files and records were easily restorable from the cloud. Cloud storage also allows sharing of files or data between multiple devices owned by one user; for example, video files downloaded onto a desktop computer might be uploaded to the cloud to allow viewing on a mobile device, or messages and calendar entries may be synchronized between multiple devices. Finally, files may be uploaded to the cloud to enable sharing between different users. Particularly for media files, this may lead to copyright issues; see below. It often also raises additional privacy questions about whether any users other than those intended have access to the shared files.

The term “cloud” has long been used to describe an abstract computer network, in which the details of the interconnections did not matter.

The Amazon Elastic Cloud was introduced in 2006, and is sometimes considered the first modern resource-pooling cloud service. It was followed by Microsoft Azure in 2008, and many other providers since then. Cloud providers include both broad-based general technology companies (such as Amazon and Microsoft) and also companies that specialize exclusively in cloud services.

Cloud Models

In 2011 the National Institute of Standards and Technology (NIST) released a technical report defining some cloud-computing terms. This report defined the following five characteristics of cloud computing:

- On-demand self-service: users can adjust their cloud resource usage dynamically, within the constraints of their service agreements, using online tools
- Broad network access: the cloud resources can be accessed from any place and any device
- Resource Pooling: the underlying hardware is pooled to support multiple cloud customers, with resources assigned to customers dynamically
- Rapid elasticity: customers can easily add additional resources, for both short-term and long-term use, using simple online interfaces
- Measured service: providers have a well-defined mechanism for defining the amount of service consumed

The NIST report also defined three influential models of cloud computing. The first is **Software as a Service**, or SaaS (usually pronounced “sass”), in which the provider supplies software which the clients then use. Sometimes the software is provided via a web interface, as in gmail.com or docs.google.com. In the latter case, the word-processing software is loaded via JavaScript whenever a user connects to the site. Microsoft's Office 365 provides similar functionality, but requires that a core application package be preinstalled on each customer device; documents are then stored on the Microsoft OneDrive cloud-storage platform. Important advantages of the SaaS model here include cloud-based document storage, a mechanism for secure document sharing based on authentication by the cloud provider, and limited (or no) software installation on the individual devices.

As another example of SaaS, many cloud providers offer database access. This may be used for an organization's primary database, or for archival records, or a possibly-short-term “data warehouse”, or anything in between. Important advantages of the SaaS model here are that the provider manages all the database software configuration and software updates, and is able to manage backups of user data.

As a third SaaS example, the cloud provider may provide web-hosting services and related application

services. The customer may then select from among these to supply the precise combination of online services needed for a particular customer-built application, including static web hosting, dynamic web content, authentication services, and others.

The primary advantage of the SaaS cloud model is that the customer avoids upfront hardware and software costs, and is freed from the details of securing, maintaining and configuring the application software and the underlying operating system.

The second NIST model is **Infrastructure as a Service**, or IaaS, in which customers lease individual computing hosts. Usually these hosts are virtual machines; this isolates the customer from the risk of hardware failure. The cloud provider often handles operating-system upgrades and manages overall network security. At one end, the IaaS model includes virtual private servers, in which customers lease the equivalent of a fixed number of hosts which they then configure individually. Billing is typically monthly. At the other end, IaaS includes services like the Amazon Elastic Cloud, in which users can lease multiple processors on a per-job basis (*eg* for a computationally intensive database operation), for as long as the job takes. Billing may be by the hour or even by the minute. This second category is a good example of the NIST concept of resource pooling: the cloud provider maintains a large number of virtual processors that can be allocated dynamically.

As with SaaS, the primary customer benefit of the IaaS cloud is the reduction in upfront hardware costs.

The third NIST model is **Platform as a Service**, or PaaS, in which customers lease nodes with a full suite of development tools installed. This is often used, for example, by customers developing software for mobile devices. Some web-development sites describe themselves as PaaS, although they overlap with SaaS.

Since the NIST definitions, other as-a-service models of cloud computing have been introduced, such

as “Backend as a Service” (not to be confused with “Blockchain as a Service”) and “Everything as a Service”. Some of these are more specific subcategories of the NIST models; some are hybrids.

Reliability and Security

After the savings on upfront hardware purchases, one of the primary advantages of public cloud computing to corporate customers is that the provider is responsible for maintaining the data center, the attendant physical hardware, and the basic software. This frees the customer to focus attention only on the higher levels of software configuration.

In most cases, providers supply resources in the form of virtual machines; these are mapped dynamically to commodity physical hardware and can be moved to alternative hardware on the fly if the original hardware fails. When properly implemented, this greatly increases system reliability from the perspective of the customer.

Some cloud providers also supply extensive security measures. Computer security is a highly specialized field, and outsourcing this task to the cloud provider typically leads to cost savings and, for all but the largest customers, improved security. A cloud provider generally has extensive experience with computer and network security. A provider is well positioned to extend defenses against a recent attack against one customer to all its customers. Finally, the cloud provider's larger network may be more resistant to denial-of-service attacks than a customer's own network.

On the other hand, a cloud provider can do little to protect against vulnerabilities in software managed and configured by a customer, in either an IaaS or SaaS setting. As an SaaS example, a provider may supply a very secure database, but the customer may introduce a vulnerability – for example, SQL injection – via its web interface. As another example, a provider's secure web-hosting platform may be compromised by flaws in the customer-provided software for handling dynamic web pages.

In the IaaS setting, the customer typically manages or configures a much larger portion of the software,

and the potential for customer-created vulnerabilities increases correspondingly. Customers in this situation must be prepared to shoulder a major portion of the security responsibility. The provider-customer agreement (below) must spell out exactly who is responsible for what.

Privacy

A public cloud provider has access to customer data, unless it is encrypted; it is essential for customers to have a complete understanding of the implications of this. Even if the customer-provider agreement strictly prohibits the provider from commercial use of the customer's data, access may still leak due to security problems or through government orders. In the latter case, in the United States at least the cloud customer may never find out about the release of their data.

Even if the privacy agreement with the cloud provider is completely acceptable to the customer, there may still be regulatory issues. If the customer is subject to the Health Insurance Portability and Accountability Act (HIPAA), or to the Family Educational Rights and Privacy Act (FERPA) or to the Payment Card Industry Data Security Standard (PCI DSS), or to other privacy and security regulations, it may be necessary to request special certification from the cloud provider. Such certification may involve additional costs, if it is available at all. All details must be negotiated upfront, and even then the regulatory rules may change in mid-contract.

There may also be rules about under what national jurisdictions the customer data may be stored. A European customer may face local rules discouraging or forbidding data storage in the United States, and vice-versa.

A cloud customer can prevent provider access by encrypting the files it stores on the cloud site. This raises problems of key management, though: if employee Alice encrypts file budget.ods and shares it via the cloud with employee Bob, then either Bob must have a pre-existing public key accessible to Alice, or else Alice must transmit a temporary key to Bob.

Cloud-storage users who share files with other users need to be concerned about just who has access to a given file, and how straightforward it is to view and modify that access. If Alice shares a file with Bob, it should be easy, for example, for Alice to verify that Charlie does *not* have access.

Organizations who build and maintain their own private cloud have the fewest external privacy concerns. For customers who lease a single-client private cloud, though, the privacy situation is quite similar to that of the public cloud.

Customer Service Agreements

For cloud applications beyond simple storage, a good agreement between customer and provider is essential. It should explicitly state the provider's obligations in terms of data access, security provisions, downtime and data backups. It should also spell out explicitly all support costs and time-lines in the event of a problem. For example, a customer database might grow so large that it no longer functions efficiently as a monolithic unit. It may need subdivision or restructuring, quite possibly on an expedited basis. The agreement should specify who pays for what.

The agreement should also spell out fees. Providers have, in principle, an ethical responsibility to create a fee structure that is easily understood by customers. In practice, there is often confusion: customers may not be well placed to estimate total costs accurately, and unanticipated surges in demand for bandwidth, storage or technical support may lead to significant cost spikes. Organizations maintaining their own private cloud do not escape these issues: on the one hand, a site with its own private cloud may be immune from some cost spikes, but on the other hand the needed additional bandwidth, storage or support resources may simply be unavailable.

Other terms of service may also cause misunderstandings. For example, some customer content may be regarded by the provider as inappropriate, or legitimate internet traffic associated with a customer's servers may fit the provider's automated profile of malicious traffic. Such circumstances may trigger a

sudden account shutdown. The customer agreement should define a clear appeals process, and, ideally, contain a provision that services will be kept running while the dispute is negotiated.

An agreement should ensure that data access will not be entirely withheld for non-payment. A related question is what happens to the customer's data if the provider goes out of business, or is sold.

Copyright

Cloud providers have an ethical responsibility to keep their customers' data reasonably secure. The customer may also have ethical responsibilities regarding data security, particularly when it comes to copyrighted content stored in the cloud. Many cloud file-sharing providers have strict rules about unauthorized sharing of copyrighted content (typically music and video files), but enforcement is another matter. In the United States, the Digital Millennium Copyright Act (DMCA) absolves providers of liability provided they abide by takedown requests from the content owner, even when the provider is generally aware that unlawful file-sharing is a frequent occurrence.

Some unethical cloud-based file-sharing sites have taken advantage of this situation by tacitly encouraging unlawful file sharing while at the same time abiding by the letter of the law. For example, a site might allow free downloads of users' publicly shared video files, but charge a fee for sufficient bandwidth to permit video streaming. The site may then reward users who uploaded content that was later frequently downloaded by other users, perhaps with the elimination of these bandwidth fees. The site has thus created an incentive for users to upload popular copyrighted videos, and a mechanism to charge other users for these videos. Under the DMCA, this may be legal. Furthermore, it often takes weeks or months for content owners to discover the shared content and send takedown requests. The ultimate legal responsibility should fall on the uploading and downloading users, but these are often anonymous. One strategy to prevent this kind of copyright abuse is for the cloud provider to require users to provide verifiable names and addresses, though this is nontrivial to implement for commodity

cloud providers. Another strategy is to provide a cap on the number of downloads by other users.

Conclusion

Ultimately, cloud computing is simply a new way of leasing computer services. It can offer many benefits to both new and long-established businesses. Cloud storage can offer similar benefits to ordinary users. Careful diligence, however, is required in understanding the many details.

Peter Lars Dordal

See also Data Mining, Data Privacy, Privacy, Cybersecurity, Copyright

Further Reading

Mell, P. and Grance, T. (2011), *The NIST Definition of Cloud Computing*, National Institute of Standards and Technology Special Publication 800-145, accessed at http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf

Amazon Web Services (2015), *Overview of Amazon Web Services (Whitepaper)*, accessed at https://d0.awsstatic.com/whitepapers/aws-overview.pdf

Microsoft Azure website, accessed at https://azure.microsoft.com

Linode Cloud Hosting website, accessed at https://www.linode.com