Infrastructure as a Service and Cloud Technologies

Nicolás Serrano, Gorka Gallardo, and Josune Hernantes

Software and infrastructure are increasingly consumed from the cloud. This is more flexible and much cheaper than deploying your own infrastructure, especially for smaller organizations. But the cloud obviously bears a lot of risk, as our partner magazine *Computer* emphasized in its October 2014 issue. Performance, reliability, and security are just a few issues to carefully consider before you use the cloud. The perceived savings immediately turn into a large extra cost if your clients and workforce are disconnected for some minutes or critical infrastructure doesn't perform as intended— which happens more often than what brokers and providers disclose. In this issue's column, IT expert Nicolas Serrano and his colleagues examine the enterprise cloud market and technologies. They provide hands-on guidance for making the right decisions. I look forward to hearing from both readers and prospective authors about this column and the technologies you want to know more about. *—Christof Ebert*

CLOUD COMPUTING'S LOW COST, flexibility, and agility are well understood in today's corporate environment. However, to fully exploit cloud technologies, you need to understand their best practices, main players, and limitations.

The concept of cloud computing has existed for 50 years, since the beginning of the Internet.[1] John McCarthy devised the idea of time-sharing in computers as a utility in 1957. Since then, the concept's name has undergone several changes: from service bureau, to application service provider, to the Internet as a service, to cloud computing, and to software-defined datacenters, with each name having different nuances. However, the core concept is the same: providing IT services based on the Internet (the cloud).

The most-used definition of cloud computing belongs to the US National Institute of Standards and Technology (NIST):

Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction.[2]

Providers use three well known models (see Figure 1): IaaS (infrastructure as a service), PaaS (platform as a service), and SaaS (software as a service). Here, we focus on IaaS. The next step is to decide on a model for deploying cloud services. In a *public cloud*, a provider provides the infrastructure to any customer. A *private cloud* is offered only to one organization. In a *hybrid cloud*, a company uses a combination of public and private clouds.

To choose the most appropriate cloud-computing model for your organization, you must analyze your IT infrastructure, usage, and needs. To help with this, we present here a picture of cloud computing's current status.

## Cloud Computing Best Practices

As with every new architectural paradigm, it's important to design your systems taking into account the new technology's characteristics. To select a cloud provider or technology, you should understand your requirements in order to list the needed features. Here are some best practices for cloud migration.[3]

An Elastic Architecture

IaaS offers precise scalability. The cloud can outperform physical hardware's classic scale-up or scale-out strategies. To gain as much as you can from this potential, architect your systems and application with as much decoupling as possible, using a service-oriented architecture and using queues between services.

Design for Failure

High scalability has limitations. IaaS technology and architecture lead to a less robust system because you're replacing hardware with several software layers, adding obvi-
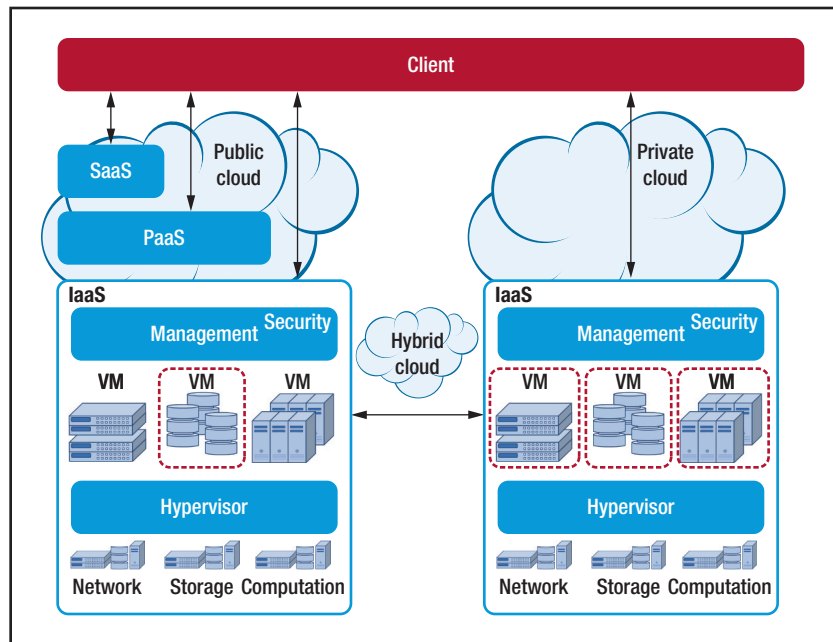
FIGURE 1. The three cloud models. IaaS is infrastructure as a service, SaaS is software as a service, PaaS is platform as a service, and VM stands for a virtual machine.

ous complexity and failure points. Redundancy and fault tolerance are primary design goals.

Besides having an established backup strategy, to assure business continuity, ensure your system is prepared for reboots and relaunches. Automation in your deployment practice is a must, with recipes for server configuration and deployment. Providing automation requires new development practices (development and operations management, continuous integration, test-driven development, and so on) and new tools such as Chef, Puppet, or Ansible.

High Availability

IT resource disruption has a huge negative impact on any business. Lost control of the underlying infrastructure when moving to the cloud, and the fact than the service-level agreement (SLA) won't cover all the

incurred costs, should lead you to design with outages and high availability in mind. With the ease of creating virtual instances, deploying clusters of servers or services is a popular approach. In this scenario, load balancing is a well-established technique for operating with clusters; it's an important feature to consider when selecting a cloud provider.

It's also important to use several available zones or at least different datacenters to make your system as robust as possible. Amazon Web Services (AWS) experienced this in April 2011 when its systems didn't run or ran intermittently for four days. Separating clusters into regions and datacenters will increase your resources' resilience.

Performance

You need to consider the technology's limitations regarding performance—

TABLE 1

Cloud provider characteristics.

| Company | Trend 2014–2015 (%)*[4] | | | Technology | Price (US$, Aug. 2013)[5] | Functionality or features[6] | | | | | Datacenters | Vendor positioning[7] |
	Reported usage (%)*[4]	Experimenting	Plan to use			Advanced features or functionality	Hybrid cloud	Network management	Service-level agreement (%)*[4]	Average service time (yrs.)		
Amazon Web Services	54	23	9	Proprietary	66	Good	Yes	Partial	99.95	5+	US East and West Coast, Ireland, Japan, Singapore, Australia, Brazil, China	Leader
Microsoft Azure	6	20	8	Proprietary	65	Good	Yes	Partial (unavailable for Linux)	99.95	4	US, Ireland, Netherlands, Hong Kong, Japan, Singapore, China, Brazil	Leader
Rackspace	12	14	12	OpenStack	116	Adequate	Yes	No	100	5+	Central and eastern US, UK, Australia, Hong Kong	Niche player
HP	4	8	8	OpenStack	87	Good	Yes	No	99.95	2	Eastern and western US	Niche player
IBM	4	9	7	OpenStack	182	Good	Yes	No	100	5+	US, Netherlands, Singapore	Niche player
Google	4	17	13	Proprietary	42	Adequate	Yes	No	99.95	2	Central US, Belgium, Taiwan	Visionary
Others†	24	45	31	—	—	—	—	—	—	—	—	—

*In these columns, the sum of the percentages is greater than 100 because some companies use several products.
† This group includes PaaS (platform as a service) and recent providers.

mainly, lack of isolation and lost robustness. In any multitenant environment, an instance's performance can be affected by your neighbors. A usage burst in a neighbor's instance can affect the available resources, notably compute units and disks' IOPS (I/O operations per second). Your architecture should deal with these changes.

Also, bottlenecks might arise owing to latency issues, even within instances at the same datacenter. Cloud providers offer some features to deal with this (for example, AWS placement groups). However, if your architecture has servers at different regional datacenters, you should use other techniques (for example, caching).

Security

Because of a public cloud's open characteristics, designing and maintaining a secure infrastructure should be an important driver in any cloud deployment. Enforce well-established security practices: firewalls, minimal server services to reduce attack vectors, up-to-date operating systems, key-based authentication, and so on. But challenges might arise from the increased number of servers to maintain and the use of the cloud for different development environments: development, staging, and

production. In this scenario, isolating and securing each environment is important because a breach in a prototyping server can give access through the secret keys to the whole infrastructure.

Monitoring

The ease of deploying new resources can make the number of servers grow exponentially. This raises new issues, and monitoring tools are vital to system management. First, they play a basic role in automatic scaling on a cyclical basis and based on events. Second, they're part of the tools needed to ensure a robust architecture, as the Netflix Chaos Monkey showed. Finally, they're important for detecting security breaches and forensic investigation, as some security breaches have shown.

Public Clouds

The public cloud was the first type of cloud offered to the general public, when AWS offered its experience with its private cloud to the general public. When you're selecting a vendor, it's important to consider several factors, mainly cost, performance, features, data location, and availability. But because the public cloud is a fairly recent technology, you should also consider vendor positioning and future use trends (see Table 1).

Cloud providers are battling for market position, which is leading them to reduce their public IaaS cloud prices, offering attractive solutions.

It's important to select the most effective vendor from a performance–cost perspective. However, your comparison should also consider whether the performance level is guaranteed, startup times, scalability responsiveness, and latency. These factors might vary among

providers and impact the infrastructure's responsiveness.

The datacenters' location can affect your decision. The provider should comply with data privacy laws and corporate policies; the server locations should be based on these considerations. These restrictions might vary among countries and among companies. You might find you'll need to have all data under the same jurisdiction (for example, in Europe). In other cases, Safe Harbor principles, in which US companies comply with EU laws, can be good enough.

Understanding each player's SLA is important. But because almost every provider offers high-enough service levels (more than 99.95 percent), it's important to evaluate the accountability the SLA offers in case of noncompliance. Normally, this won't cover the costs of outages, so your infrastructure should be prepared for them.

Providers

Once you've defined your selection criteria, you can compare providers. The following are the most relevant ones.

Amazon. AWS (http://aws.amazon .com) continues being the dominant player in cloud computing, thanks to Amazon having been the first company to offer cloud services, in 2006.

AWS is cost effective. Its pay-as-you-go model lets you scale cloud capacity up or down without paying a high price. It also offers many additional IaaS services and integrated monitoring tools. It's particularly valuable for startups and agile projects requiring quick, cheap processing and storage.

Because AWS is a general provider, you can operate independently,

which is convenient for normal operations but becomes risky when problems occur. Extensive technical support is a premium feature, whereas most of AWS's competitors offer it as a standard feature.

Microsoft Azure. Azure (http:// azure.microsoft.com/en-us) entered the cloud IaaS market in February 2010. It has a large market share and is a good candidate because of its market position in other areas. It offers compute and storage services similar to those of other IaaS providers, and it allows full control and management of virtual machines. Additionally, Azure's UI is easy to use, especially for Windows administrators. However, because the Azure offering is newer than Amazon's or Rackspace's, it still has many features in "preview" mode and still has networking and security gaps.

Rackspace. Rackspace (www .rackspace.com) is a founder of OpenStack (which we describe later) and a major player in open source cloud IaaS. It hosts more than half of the Fortune 500 companies at its datacenters, while strongly focusing on SMEs (small-to-medium enterprises).

Rackspace provides an inexpensive, intuitive cloud with optional managed services and an easy-to-use control panel that suit SMEs. It also guarantees extensive support. However, it has limited pricing options, providing only month-to-month subscriptions. Also, it doesn't offer specialized services.

Google. Although Google AppEngine was a pioneer of cloud computing in the PaaS model, Google Compute Engine (https://cloud.google .com/compute) is relatively new

TABLE 2

The main products used to create private clouds.*

Solution	Reported usage (%)[4]	License	Trend, 2013–2014[4]
VMware	43	Proprietary	Down
OpenStack	12	Open source	Up
Microsoft	11	Proprietary	Down
CloudStack	6	Open source	Down
Eucalyptus	3	Open source	Down

*Some of the included solutions, as stated in State of the Cloud Report,[4] don't strictly meet all cloud-computing requirements.

to the IaaS market. Nevertheless, Google's number of physical servers and global infrastructure make it a good candidate. Moreover, Google Compute Engine is well integrated with other Google services such as Google Cloud SQL and Google Cloud Storage.

Google Compute Engine is well suited for big data, data warehousing, high-performance computing, and other analytics-focused applications. Its main limitation is that it doesn't integrate administrative features. So, users must download extra packages.

HP. HP is still relatively new in the IaaS game; it launched its service in December 2012. Its public cloud, HP Cloud Compute (www.hpcloud.com), is built on OpenStack and offers a range of cloud-related products and services. It's a good candidate owing to its positioning in the server market. Its IaaS offering supports public, hybrid, and private clouds.

HP Cloud Compute is a good solution for companies that want to integrate their existing IT infrastructure with public-cloud services and invest in a hybrid cloud.

IBM. IBM's resources, size, and knowledge of datacenters make it another player to consider. IBM Cloud (www.ibm.com/cloud-computing/us

/en) offers core computing and storage services. This IaaS is best for large enterprises with heavy data-processing needs and security concerns.

IBM Cloud provides a good combination of management, software, and security features for administrators. However, its focus is limited to medium-to-large enterprises and enterprises whose main provider is IBM.

Issues and Concerns

When considering adoption of a cloud architecture, it's important to understand what the technology can offer you and the main issues you'll have to deal with in each of these new infrastructures. Only by clearly understanding each of the approaches' business and technical opportunities and limitations will you be able to select the best option on the basis of your needs.

Besides the economic advantages from a cost perspective, the main competitive advantages are the flexibility and speed the cloud architecture can add to your IT environment. In particular, this kind of architecture can provide faster deployment of and access to IT resources, and fine-grain scalability.

A recent survey indicated the issues that beginner and experienced enterprise cloud users face.[4] For beginners, the main issues are se-

curity, managing multiple clouds, integration with current systems, governance, and lack of expertise. Experienced companies face issues of compliance, cost, performance, managing multiple clouds, and security.

The differences are understandable. Different problems arise on the basis of the degree of advancement of cloud architecture adoption. Early on, the main issues are resource expertise and control because the company hasn't acquired enough knowledge of and experience with the architecture. For more experienced companies, performance and cost are important because the architecture's limitations might have started emerging.

Both groups must deal with security, compliance, and managing multiple clouds. Regarding security and compliance, some problems might arise from the multitenant architecture. Some of these problems might not be solved, which might tip the balance toward a private or hybrid cloud. Such a decision is plausible, in keeping with the issue of managing multiple clouds.

Private and Hybrid Clouds

To solve the issues with public clouds, cloud-computing providers introduced the private cloud. This cloud might be in the organization's buildings, in the farm of the organization's provider, or in another provider's datacenter. Usually, it will be virtualized, but other combinations are possible. The important element is that only the customer's organization can operate it.

Because all private-cloud products allow integration with public clouds, we discuss both private and hybrid clouds here. Table 2 shows the main products used to create private clouds.

Eucalyptus

Eucalyptus (www.eucalyptus.com) released its first product in 2008. Nowadays the company provides its software as open source products and services. (Recently, Eucalyptus was bought by HP, a supporter of OpenStack.) From the company's download area, you can install a private cloud on your computer. From its product area, you can contract servers for your private cloud.

Eucalyptus software's main advantage is its AWS compatibility (see Table 3), based on a partnership with Amazon. So, some features that AWS makes available for the public cloud are applicable to Eucalyptus services.

Eucalyptus software's weak points are the limited GUI and the risk of uncertainty generated by AWS's private-cloud strategy: AWS offers Amazon Virtual Private Cloud and a connection to a hardware VPN (virtual private network).

OpenStack

OpenStack (www.openstack.org) is the other main player in the private-cloud field. It's also open source, and its greatest strength is its support from companies such as AT&T, AMD, Cisco, Dell, HP, IBM, Intel, NEC, Red Hat, VMware, and Yahoo.

OpenStack is complex, with different components and multiple command-line interfaces. Competitors say it's not a product but a technology. This can be a barrier for nontechnical companies but not for public- and private-cloud providers, which are OpenStack's main users. For them, an open source product is attractive because, just as with using Linux in server computers, there are cost and portability advantages for the end user.

Portability is another important feature of OpenStack because end

TABLE 3

Eucalyptus compatibility with Amazon Web Services (AWS).

AWS services	Eucalyptus components
Amazon Elastic Compute Cloud (EC2)	Cloud Controller
Amazon Elastic Block Storage (EBS)	Storage Controller
Amazon Machine Image (AMI)	Eucalyptus Machine Image
Amazon Simple Storage Service (S3)	Walrus Storage
Amazon Identity and Access Management (IAM)	IAM
Auto Scaling	Auto Scaling
Elastic Load Balancing (ELB)	ELB
Amazon CloudWatch	CloudWatch

users don't want to be locked into a particular provider. However, providing the option of portability can be an issue for providers that want to offer differentiated proprietary features.

CloudStack

Citrix purchased CloudStack (http://cloudstack.apache.org) from Cloud.com. Citrix donated it to the Apache Software Foundation, which released it after it spent time in the Apache Incubator.

Unlike OpenStack, CloudStack offers a complete GUI and a monolithic architecture that simplifies installing and managing the product. Like OpenStack, most installations belong to service providers. CloudStack also offers AWS compatibility through an API translator.

Proprietary Solutions

VMware (www.vmware.com) and Microsoft (www.microsoft.com/enterprise/microsoftcloud) emphasize the hybrid nature of their offerings. They have products for both public and private clouds and provide on-premises servers.

VMware products include vCloud Hybrid Service, vCloud Connector, and vSphere virtualization. Microsoft has Windows Azure, Windows Server, and Microsoft System Center.

These two providers offer a more integrated solution because they own their products, but the disadvantage is lack of portability.

The public-cloud market has some years of history and well-known players. But remember that the cloud-computing market is growing. Newcomers are always entering, and the leaders in public- and private-cloud services can change.

So, your selection of a cloud-computing model and provider must take into account the factors listed in Tables 1 and 2, a service's specific purpose, and the elements of the application you want to migrate to the cloud. The approach and reach of your cloud adoption efforts will be limited by each situation. For example, your application architecture and the technology involved won't be the same if you're migrating an application not yet developed or an existing legacy system.

Regarding a new application,

BUNTPLANET

BuntPlanet (http://buntplanet.com) is a small-to-medium enterprise focusing on software engineering. It was founded in 2000 in San Sebastian, Spain. It offers a range of applications mainly for utilities, and it develops custom applications using agile practices.

Although BuntPlanet had some experience with cloud services, it was only in 2009 that it decided to use the cloud for its internal applications. The availability of Amazon's European datacenter was a determining factor because BuntPlanet could comply with Spanish data protection laws. During the selection process, the company compared other alternatives but chose Amazon mainly for the extensive features and competitive price.

At first, BuntPlanet replicated its server architecture in the cloud with Amazon EC2 (Elastic Compute Cloud) servers. Since then, it has incrementally refactored its applications to realize the full potential of a cloud architecture. It now uses other Amazon services such as S3 (Simple Storage Service) for data storage, Glacier for backups, RDS (Relational Database Service) for relational databases, SNS (Simple Notification Service) for push notifications, and SQS (Simple Queue Service) for queues.

To achieve the required service level, BuntPlanet uses reserved instances and a second redundant region in Europe. It's also self-monitoring its cloud resources and has set alerts for failures. This approach lets the company achieve its desired reliability and availability.

BuntPlanet's experience with the public cloud has been positive from economical and technical viewpoints. Using a public cloud has allowed BuntPlanet to simplify processes and minimize the need to support a big hardware infrastructure and a high bandwidth in its installations. As a result, it's promoting this architecture in its customer projects, using the public cloud for load tests and production environments.

might choose the Microsoft solution. However, AWS, a market leader and proven feature-rich platform, is always an option.

As you can see, because of the variety of choices, different customers might choose different platforms. For example, HP and Rackspace (service providers), Cybercom (a consulting company), and eBay (an end user) use OpenStack, whereas VMware and Microsoft customers use their provider's solution. For a look at how one company (BuntPlanet) chose its cloud provider, see the sidebar.

References

1. M. Vouk, "Cloud Computing—Issues, Research and Implementations," *J. Computing and Information Technology*, vol. 16, no. 4, 2008, pp. 235–246.
2. P. Mell and T. Grance, *The NIST Definition of Cloud Computing*, US Nat'l Inst. of Standards and Technology, 2011.
3. F. Fehling, F. Leymann, and R. Retter, "Your Coffee Shop Uses Cloud Computing," *IEEE Internet Computing*, vol. 18, no. 5, 2014, pp. 52–59.
4. *State of the Cloud Report*, RightScale, 2014; www.rightscale.com/lp/2014-state-of-the-cloud-report.
5. T. Rodrigues, "Top Cloud IaaS Providers Compared," *Enterprise Cloud*, 27 Aug. 2013; www.techrepublic.com/blog/the-enterprise-cloud/top-cloud-iaas-providers-compared.
6. "Vendor Landscape: Cloud Infrastructure-as-a-Service," Info-Tech Research Group, 2014; www.infotech.com/research/ss/select-the-right-cloud-infrastructure-server-partner.
7. *Magic Quadrant for Cloud Infrastructure as a Service*, Gartner, May 2014; www.gartner.com/technology/reprints.do?id=1-1UKQQA6&ct=140528&st=sb.

NICOLÁS SERRANO is a professor of computer science and software engineering at the University of Navarra. Contact him at nserrano@tecnun.es.

GORKA GALLARDO is a professor of information systems at the University of Navarra. Contact him at ggallardo@tecnun.es.

JOSUNE HERNANTES is a professor of computer science and software engineering at the University of Navarra. Contact her at jhernantes@tecnun.es.

you should develop it with an elastic architecture and best practices in mind. Decouple the presentation, business, and logic layers in several services and use a queue system to communicate between them. A high number of servers, a fault-tolerant design, and automatic provisioning will require high-level features from the cloud provider or technology.

Regarding a complete legacy system, refactoring the application to achieve decoupling isn't feasible. A pure cloud architecture is impossible, and a reduced list of features is required. Your priority should be virtual instances' robustness and reliability.

Other scenarios, such as disaster recovery or using the cloud when the demand spikes (cloud bursting), require specific cloud technology features.

If you're dealing with a new application and provider independence is a priority, you might prefer an OpenStack provider. If you're migrating a legacy system and you have IT experience with VMware, you might select VMware for your cloud. Regarding cloud bursting in a Microsoft Server IT environment, you