Infrastructures for High-Performance Computing: Cloud Infrastructures

Loris Belcastro
c/o DIMES – University of Calabria
Via P. Bucci 42c
87036 Rende (CS)
Italy
Email: lbelcastro@dimes.unical.it
Telephone: +39 0984 494782

Cristian Cosentino
c/o DIMES – University of Calabria
Via P. Bucci 41c
87036 Rende (CS)
Italy
Email: cristian.cosentino@dimes.unical.it
Telephone: +39 0984 494782

Fabrizio Marozzo (corresponding author)
c/o DIMES – University of Calabria
Via P. Bucci 42c
87036 Rende (CS)
Italy
Email: fmarozzo@dimes.unical.it
Telephone: +39 0984 494782

Keywords

Amazon Web Services (AWS), Cloud Stack, Eucalyptus, Google cloud platform, IBM Bluemix, Microsoft Azure, OpenNebula, OpenStack, Oracle Cloud, Private cloud and Public cloud

Abstract

This article describes the main services provided by the most popular cloud infrastructures currently in use. The systems presented in this article are either public clouds, such as Amazon Web Services (AWS), Google Cloud Platform, Microsoft Azure, IBM Bluemix and Oracle Cloud, or private clouds, such as OpenStack, OpenNebula, Eucalyptus and Cloud Stack. The formers

provide a variety of services that implement all the three main cloud service models (IaaS, PaaS and SaaS), while the latters are focused on the IaaS model, as they provide services for the management of large pools of processing, storage, and networking resources in a data center, by exploiting a variety of virtualization techniques.

Key points

- This chapter provides an overview of the main cloud infrastructures currently in use, including both public clouds and private clouds.
- The content of this chapter analyzes public cloud infrastructures (AWS, Google Cloud Platform, Microsoft Azure, IBM Bluemix and Oracle Cloud) by paying attention to the IaaS, PaaS and SaaS services they provide.
- Some examples of IaaS services made available by private cloud infrastructures (OpenStack, OpenNebula, Eucalyptus and Cloud Stack) are described.

Introduction

This article describes the main services provided by the most popular cloud infrastructures currently in use. The systems presented in this article are either public clouds, such as *Amazon Web Services (AWS)*, *Google Cloud Platform*, *Microsoft Azure*, *IBM Bluemix* and *Oracle Cloud*, or private clouds like *OpenStack*, *OpenNebula*, *Eucalyptus*, and *Cloud Stack*.

Public cloud systems provide services according to the Infrastructure as a Service (IaaS), Platform as a Service (PaaS) and Software as a Service (SaaS) models (Talia, Trunfio and Marozzo, 2015). With the IaaS model, the cloud system provides services that allow customers to rent virtualized resources like storage and servers; the PaaS model allows developers to focus on deploying of applications, since cloud providers are in charge of managing the underlying infrastructure; with the SaaS model, software and data are provided to customers as ready-to-use services.

As examples of IaaS services, AWS provides the popular *Elastic Compute Cloud* (EC2) service that allows creating and running virtual servers; Google Cloud Platform allows developers to run virtual machines through its *Compute Engine* service; Microsoft Azure provides a *Virtual Machines* service to manage virtual-machine instances; IBM Bluemix *Virtual servers* can be used to deploy and scale virtual machines on demand.

As examples of PaaS services, AWS provides *Elastic Beanstalk* to create, deploy, and manage applications using a large set of infrastructure services; Google Cloud Platform includes *App Engine*, a PaaS that allows developers to create and host web applications in Google-managed data centers; Microsoft Azure's PaaS includes a *Azure Cloud Services*, a technology that includes a *Web role* for developing Web-based applications and a *Worker role* for batch applications; IBM Bluemix includes *Mobile Foundation*, a set of APIs to manage the life cycle of mobile applications.

Finally, some examples of SaaS services are *Simple Email Service* (AWS), *Gmail* (Google Cloud Platform), *Office 365* (Microsoft Azure), *Natural Language Understanding* (IBM Bluemix) and *AI Services* (Oracle Cloud).

In general, the private cloud systems are more focused on the IaaS model, as they provide services for the management of large pools of processing, storage, and networking resources in a data center, by exploiting a variety of virtualization techniques. For example, OpenStack with its *Compute* component provides virtual servers and with its *Storage* component provides a scalable and redundant storage system; OpenNebula's *Core component* creates and controls virtual machines by interconnecting them with a virtual network environment; Eucalyptus includes a *Node Controller* that hosts the virtual machine instances and manages the virtual network endpoints, and a *Storage Controller* that provides persistent block storage and allows the creation of snapshots of volumes; Cloud Stack includes low-level services to manage compute, storage,

network, and virtualization resources.

Amazon Web Services

Amazon is a leading provider of cloud services, which it offers to developers in the form of Amazon Web Services (AWS). AWS comprises a large set of computing and storage resources that can be leveraged by users to create and run software applications with cloud capabilities. AWS is a versatile platform that can be tailored to meet the specific needs of developers, allowing them to customize their SaaS applications or integrate traditional software with cloud capabilities. This is made possible by the availability of a wide range of cloud services, including storage, compute, analytics, database, and security services. Interacting with AWS is made easy thanks to the availability of software development kits (SDKs) for the main programming languages and platforms, such as Java, .Net, PHP, and Python. This ensures that developers can seamlessly integrate their applications with AWS without having to learn a new programming language or platform. AWS includes the following main services:

- Compute: *Elastic Compute Cloud* (EC2) allows creating and running virtual servers; *Amazon Elastic MapReduce* for building and executing MapReduce applications.
- Storage: *Simple Storage Service* (S3), which is an object storage service for storing and retrieving data via the Internet; *Amazon Elastic Block Store* (EBS), a scalable, high-performance block-storage service designed to work with Amazon EC2; *Amazon Elastic File System* (EFS), a shared file storage for use with EC2 instances and on-premises servers.
- Database: *Relational Database Service* (RDS) for relational tables; *DynamoDB* for non-relational tables; *SimpleDB* for managing small datasets; *ElasticCache* for caching data.

- Networking: *Route 53*, a DNS Web service; *Virtual Private Cloud* for implementing a virtual network.
- Serverless Computation: *Lambda*, which is a serverless processing service that allows you to run code without having to manage the underlying infrastructure.
- Deployment and Management: *CloudFormation* for creating a collection of ready-to-use virtual machines with pre-installed software (e.g., Web applications); *CloudWatch* for monitoring AWS resources; *Elastic Beanstalk* to deploy and execute custom applications written in Java, PHP and other languages; *Identity and Access Management* to securely control access to AWS services and resources.
- Content delivery: *Amazon CloudFront* makes it easy to distribute content via a global network of edge locations.
- App services: *Simple Notification Service* to notify users; *Simple Queue Service* that implements a message queue; *Simple Workflow Service* to implement workflow-based applications.

Amazon is widely recognized as the pioneer of Infrastructure as a Service (IaaS) due to its ground-breaking EC2 and S3 services. However, AWS has continuously expanded its offerings to include Platform as a Service (PaaS) and Software as a Service (SaaS) solutions as well. One example of Amazon's PaaS offering is *Elastic Beanstalk*, which enables users to quickly and easily deploy, create, and manage their applications using a range of AWS services. *Elastic Beanstalk* streamlines the development process, allowing developers to focus on coding without having to worry about infrastructure management. Another notable service from Amazon's SaaS portfolio is *Amazon Machine Learning*, which provides powerful visualization tools and wizards for creating machine-learning models quickly and easily, without requiring extensive programming knowledge. With

Amazon Machine Learning, users can create accurate predictions and insights from their data, making it a valuable tool for businesses of all sizes. Finally, as an example of SaaS service, Amazon's *Simple Email Service* is a basic email-sending service that provides a reliable, scalable, and cost-effective solution for businesses that need to send transactional emails, such as order confirmations, notifications, and newsletters. The service is easy to set up and use, making it a popular choice for companies looking for a reliable and efficient way to manage their email communications. Moreover, AWS offers a variety of data processing and analytics frameworks that can be used to manage and analyze large amounts of data (Belcastro, Cantini, Marozzo et al. (2022)), such as Apache Hadoop, Spark and Hive.

Google Cloud Platform

Google Cloud Platform is a comprehensive suite of cloud services offered by Google. The platform is designed to enable developers to create a diverse range of web applications, from basic websites to complex, scalable, and multi-tenant applications. One of the key advantages of the Google Cloud Platform is that the services created by users run on the same underlying infrastructure that powers some of Google's most popular products, including Google Search, YouTube, Maps, and Gmail. This means that users can benefit from the same robust and reliable infrastructure that Google uses to support its own services. To make it easier for developers to create applications on the Google Cloud Platform, Google provides software development kits (SDKs) that support several programming languages, libraries, and frameworks. These SDKs are available for popular languages such as GO, PHP, Java, Python, and .Net, allowing developers to use the language they are most comfortable with. The platform provides several services to build and scale an application:

- Compute: *Compute Engine*, an IaaS that allows developers to run virtual machine (VM) instances; *App Engine*, a PaaS that allows developers to create and host web applications in Google-managed data centers; *Container Engine*, a management and orchestration system for software containers.
- Storage and Databases: *Cloud Storage*, a cloud storage web service for storing and accessing large and/or unstructured data sets; *Cloud Datastore*, a highly scalable NoSQL database service; *Cloud Bigtable*, a highly scalable NoSQL database service used by Google to power many core services, including Search, Analytics, Maps, and Gmail.
- Networking: *Cloud Virtual Network*, a managed networking functionality for web applications; *Cloud Load Balancing*, a load-balancer for compute resources in single or multiple regions.
- Big Data: *BigQuery*, a data warehouse enabling fast SQL queries; *Cloud Dataflow*, a data processing service for stream and batch execution of pipelines; *Cloud Pub/Sub*, a real-time messaging service to send and receive messages between applications.
- Machine Learning: *Cloud Machine Learning Engine*, a scalable machine learning service to train models and predict new data; *Cloud Natural Language API*, for extracting information mentioned in text documents, news articles or blog posts.
- Management Tools: *Stackdriver Monitoring*, to monitor cloud resources performance; *Stackdriver Logging*, for storing, monitoring, and alerting log data and events from google services; *Google Kubernetes Engine*, a container orchestration service that allows to manage and orchestrate Docker containers at scale.
- Identity and Security: *Cloud Identity & Access Management*, a service to establish a policy of access to the various services; *Cloud Key Management Service*, a key management for

managing encryption of cloud services.

Microsoft Azure

Azure is the public cloud computing platform operated by Microsoft, which includes a wide range of products and cloud services for developing cloud-native applications or enhancing existing applications with cloud-based functionalities. The platform offers flexible computing and storage resources that leverage the computational and storage capabilities of Microsoft's data centers. Azure's architecture is geared towards delivering high availability and dynamic scaling services that can be tailored to meet user requirements, while its pay-per-use pricing model ensures cost-effective usage. The platform is well-suited for storing large datasets, processing large batch computations, and building Software-as-a-Service (SaaS) applications targeted at end users. Microsoft Azure includes three basic components/services:

- *Compute* is the computational environment to execute cloud applications. Each application is structured into roles: *Web role*, for Web-based applications; *Worker role*, for batch applications; *VM role*, for virtual-machine images.
- *Storage* provides scalable storage to manage: binary and text data (*Blobs*), non-relational tables (*Tables*), queues for asynchronous communication between components (*Queues*), and virtual disks (*Disks*).
- *Fabric controller* whose aim is to build a network of interconnected nodes from the physical machines of a single data center. The Compute and Storage services are built on top of this component.

Microsoft Azure offers a range of services that can be leveraged across multiple application domains. Among these services, *Machine Learning* enables the development of cloud-based

machine learning applications, while the *IOT Hub* facilitates the connection, monitoring, and control of IoT devices. Standard interfaces provided by Azure allow developers to seamlessly interact with these services. To further streamline the process of designing and publishing Azure applications, developers can make use of Integrated Development Environments (IDEs) such as Microsoft Visual Studio and Eclipse.

IBM Bluemix

IBM Bluemix delivers services that can easily integrate with cloud applications without users needing to know how to install or configure them. It supports several programming languages and services for building, running and managing applications on the cloud (e.g., NET, Java, Node.js, PHP, Python, Ruby, Go). With Bluemix customers can access three types of cloud computing platforms:

- *Bluemix public*, which provides all the necessary resources for an application development available to the general public;
- *Bluemix dedicated*, which provides a own Bluemix private environment that is hosted in an isolated container managed by Bluemix;
- *Bluemix Hybrid*, which allows customers to connect their dedicated services to the Public Bluemix services by maintaining data and sensitive services on premise.

IBM Bluemix offers mainly two compute services for defining applications: *Containers* and *OpenStack VMs*. *IBM Containers* permits developers to deploy, manage and run applications on the IBM platform by leveraging open-source Docker container technology (Merkel, 2014). Docker is an open platform used for building, shipping and running applications. *OpenStack Virtual Machines* is a service that allows running specific operating systems that users can customize.

The IBM Bluemix platform provides several services to build and scale an application:

- *Compute*: *Virtual servers* to deploy and scale virtual machines on demand; *GPU Servers* to deploy high-performance GPU servers; *Containers,* to deploy, manage and run application on IBM platform by leveraging open-source Docker container technology; *Kubernetes Service*, a container orchestration service that enables you to manage and orchestrate Docker containers at scale.
- *Storage*: *Object Storage,* to access unstructured data via RESTful APIs; *File Storage*, a full-featured file storage; *Block Storage*, a scalable block storage service.
- *Network*: *Load Balancing*, a load-balancer for compute resources; *Direct Link*, to transfer data between a private infrastructure and IBM Bluemix services.
- *Mobile: Swift*, to write Swift mobile applications; *Mobile Foundation*, a set of APIs to manage the app life cycle.
- *Analytics*: *Analytics Services*, to extract insight from data; *Natural Language Understanding*, to extract meta-data from text such as sentiment, and relations.
- *Internet of Things: IOT platform,* which allows to communicate with and consume data from IOT devices.
- *Security*: *Firewalls* for protecting from malicious activities; *Security Software*, an intrusion protection system for the network and server/host level.

Oracle Cloud

Oracle Cloud is a cloud computing platform offered by Oracle Corporation that provides a wide range of cloud services. The platform offers a variety of cloud deployment models, including public, private, and hybrid clouds, to meet the needs of different organizations. In addition, Oracle

Cloud implements all the three main cloud service models:

- *Infrastructure as a Service (IaaS)*: Oracle Cloud provides a range of infrastructure services, including computing, storage, and networking, which can be used to build and run applications.
- *Platform as a Service (PaaS)*: Oracle Cloud offers a set of platform services that enable developers to build and deploy applications without having to manage the underlying infrastructure.
- *Software as a Service (SaaS)*: Oracle Cloud offers a range of software applications, such as enterprise resource planning (ERP), human capital management (HCM), and customer experience (CX) applications, which can be accessed over the Internet.

Oracle Cloud includes the following main services:

- *Compute*, a service that provides virtual servers for running applications and services.
- *Storage*, a distributed storage system for storing and accessing data.
- *Networking*, a service for managing networks and IP addresses.
- *Database*, a service for managing databases and data storage.
- *Identity and Access Management:* is a service that provides authentication and authorization for Oracle Cloud services and users.
- *Security:* is a service that provides advanced security features to protect applications and data in Oracle Cloud.
- *Analytics:* is a service that provides analytics tools and technologies to help organizations make better decisions based on their data.

OpenStack

OpenStack is a cloud system that allows the management of large pools of processing, storage, and networking resources in a data center through a Web-based interface (Wen, Gu, Li et al., 2012). The system has been designed, developed and released following four open principles:

- *Open source*: OpenStack is released under the terms of the Apache License 2.0;
- *Open design*: Every six months there is a design summit to gather requirements and define new specifications for the upcoming release;
- *Open development*: A publicly available source code repository is maintained for the entire development process;
- *Open Community*: Most decisions are made by the OpenStack community using a lazy consensus model.

The modular architecture of OpenStack is composed of a number of components providing APIs to access infrastructure resources. Among these, the core components are:

- OpenStack *Nova*, which offers a dynamic, on-demand virtual machine infrastructure by managing the processing resource pool available in the data center. This service is compatible with various virtualization technologies, such as VMware and KVM, and its architecture is designed to scale horizontally.
- OpenStack *Horizon*, which is the web interface to OpenStack for managing and configuring the entire infrastructure without having to use the command line interface.
- OpenStack *Swift*, which provides a highly scalable and distributed object storage system, ideal for storing binary or unstructured data of any size.
- OpenStack *Cinder*, a component that provides block storage services to users. It allows users to create and manage block storage volumes and attach them to virtual machines.
- OpenStack *Neutron*, which manages network resources and IP addresses, making it easier

to control and manage complex network topologies.

- OpenStack *Keystone*, which is the Identity Management Service for the cloud infrastructure. It maps users and their permissions for the different services they use.

OpenNebula

OpenNebula is an open-source framework that is primarily used for creating private and hybrid clouds. It was introduced by (Sotomayor, Montero, Llorente et al., 2009). The architecture of OpenNebula is based on a Core component that is responsible for creating and controlling virtual machines. The Core also creates a virtual network environment, which interconnects the virtual machines. In addition, the Core communicates with specific storage, network, and virtualization operations through pluggable components known as *Drivers*. This design allows OpenNebula to be independent of the underlying infrastructure and provide a uniform management environment. The Core also supports the deployment of *Services*, which are sets of linked components executed on multiple virtual machines. For instance, a Web server and a database can be linked together as a Service. Another important component of OpenNebula is the Scheduler, which is responsible for allocating the virtual machines to the physical servers. It communicates with the Core through appropriate deployment commands.

Overall, OpenNebula provides a powerful and flexible framework for building private and hybrid clouds. Its pluggable architecture allows for easy integration with a variety of underlying infrastructures, and its support for Services and *Scheduler* makes it a comprehensive cloud management solution. OpenNebula includes the following main services:

- *Compute:* is a service that provides virtual servers for running applications and services.
- *Storage:* is a distributed storage system for storing and accessing data.

- *Networking:* is a service for managing networks and IP addresses.
- *Security* (*AuthZ*): is a service for managing authentication and authorization for OpenNebula services and users.
- *Cloud Infrastructure Management (Sunstone)*: is a web-based interface for managing OpenNebula services and resources.
- *Orchestration (OneFlow): ia a* service that automates the deployment and management of OpenNebula resources and applications.

OpenNebula can implement a hybrid cloud using specific Cloud Drivers that allow interacting with external clouds. In this way, the local infrastructure can be supplemented with computing and storage resources from public clouds. Currently, OpenNebula includes drivers for using resources from Amazon EC2 and Eucalyptus.

Eucalyptus

Eucalyptus (Elastic Utility Computing Architecture for Linking Your Programs To Useful Systems) is an open source and paid software platform for implementing IaaS in a private or hybrid cloud computing environment (Endo, Gonçalves, Kelner et al. (2010)). Eucalyptus enables developers to use compute, storage, and network resources that can be dynamically scaled and can be used to build web-based applications. Companies can use AWS-compatible services, images, and scripts on their own on-premises IaaS environments, without using a public Cloud environment like AWS. It is simple to interact with Eucalyptus since it provides a console to self-service provision and configure compute, network, and storage resources.

Eucalyptus is based on five basic components. The *Node Controller* (NC) hosts the virtual machine instances and manages the virtual network endpoints. The *Cluster Controller* (CC) acts as the front

end for a cluster and it is responsible for deploying and managing instances on NCs. The *Storage Controller* (SC) provides persistent block storage and allows the creation of snapshots of volumes. The *Cloud Controller* (CLC) is the front-end for managing cloud services and performs high-level resource scheduling and system accounting. CLC also allows interacting with the components of the Eucalyptus infrastructure providing an AWS-compatible web services interface to customers. *Walrus Storage Controller (WS3)* offers persistent storage to all of the virtual machines in the Eucalyptus cloud and can be used to store machine images and snapshots. WS3 can be used with AWS S3 APIs.

Cloud Stack

Cloud Stack is an open-source cloud computing platform that enables the creation, deployment, and management of private and hybrid cloud environments. It provides a complete stack of cloud computing infrastructure services, including computing, storage, networking, and virtualization. The main benefits of Cloud Stack include:

- *Flexibility:* it allows businesses to choose the resources they need and to scale up or down as their needs change.
- *Cost savings:* by using a pay-as-you-go model, businesses can avoid upfront infrastructure costs and only pay for the resources they use.
- *Security:* it provides a secure and reliable environment for businesses to store and process their data.

Cloud Stack includes the following main services:

- *Compute:* it enables businesses to create and manage virtual machines (VMs) and

containers to run their applications.

- *Storage:* it provides storage services for businesses to store and manage their data, including object storage, block storage, and file storage.
- *Networking:* it enables businesses to create and manage their network infrastructure, including virtual networks, load balancers, and firewalls.
- *Virtualization:* it supports multiple hypervisors, including KVM, Xen, and VMware, to enable businesses to virtualize their infrastructure.

Closing Remarks

This article described the main services provided by the most popular cloud infrastructures currently in use. The systems presented in this article were either public clouds Amazon Web Services (AWS), Google Cloud Platform, Microsoft Azure, IBM Bluemix and Oracle Cloud or private clouds OpenStack, OpenNebula, Eucalyptus and Cloud Stack. As discussed above, the public cloud systems provide a variety of services that implement all the three main cloud service models, i.e., Infrastructure as a Service (IaaS), Platform as a Service (PaaS) and Software as a Service (SaaS). On the other hand, the private cloud systems are more focused on the IaaS model, as they provide services for the management of large pools of processing, storage, and networking resources in a data center, by exploiting a variety of virtualization techniques.

See also

Computing Languages for Bioinformatics: Java. Dedicated Bioinformatics Analysis Hardware. Models and Languages for High-Performance Computing. Parallel Architectures for Bioinformatics. Text Mining Applications.

References

Belcastro, L., Cantini, R., Marozzo, F., Orsino, A., Talia, D. and Trunfio, P. (2022). Programming Big Data Analysis: Principles and Solutions. *Journal of Big Data* **9**(4).

Endo, P.T., Gonçalves, G.E., Kelner, J. et al. (2010). A survey on open-source cloud computing solutions. In Brazilian Symposium on Computer Networks and Distributed Systems, vol. 71.

Merkel, D. (2014). Docker: Lightweight Linux containers for consistent development and deployment. Linux Journal, 2014(239).

Sotomayor, B., Montero, R.S., Llorente, I.M. et al. (2009). Virtual infrastructure management in private and hybrid clouds. IEEE Internet Computing, 13, 14-22.

Talia, D., Trunfio, P. and Marozzo, F. (2015). Data analysis in the cloud. Elsevier, The Netherlands.

Wen, X., Gu, G., Li, Q. et al. (2012). Comparison of open-source cloud management platforms: OpenStack and OpenNebula. In Proceedings of the 9th International Conference on Fuzzy Systems and Knowledge Discovery (FSKD), pp. 2457-2461.

Relevant Websites

https://aws.amazon.com/ – Amazon Web Services.

https://cloudstack.apache.org/ - Cloud Stack.

https://www.docker.com/ – Docker.

http://www.eucalyptus.com/ – Eucalyptus.

https://cloud.google.com/ – Google Cloud Platform.

https://www.ibm.com/cloud-computing/bluemix/ – IBM Bluemix.

https://www.microsoft.com/azure – Microsoft Azure.

https://www.openstack.org/ – OpenStack.

https://opennebula.org/ – OpenNebula.

https://www.oracle.com/it/cloud/ - Oracle Cloud.