Workbook
EXIN Cloud Computing Foundation

Edition 202504

Hans van den Bent
CLOUD-consulting

Alexander Vladimirovich Esis
Sberbank

Version administration

Version	Changes
201204	Original version
201510	• Updates of content • Lay-out editing
202008	• Updates of content • Lay-out editing
202009	Added reference to the latest GDPR information regarding the US Privacy Shield adequacy decision.
202105	Corrected an error in 3.1.1.1, describing the relationship between the OSI model and the protocols
202504	• Reviewed and updated references to ISO/IEC standards • Updated to new template

Content

Introduction

Cloud computing is about providing IT related services through the Internet. Cloud computing allows flexible IT solutions to support businesses, based on clear service arrangements. This workbook will help you prepare for the EXIN Cloud Computing Foundation exam and provides you with an overview of cloud computing and its relationship with other areas of information management. The topics discussed are the fundamental concepts of cloud computing such as architecture, design, and deployment models of cloud computing, and how cloud computing can be employed by different types of organizations.

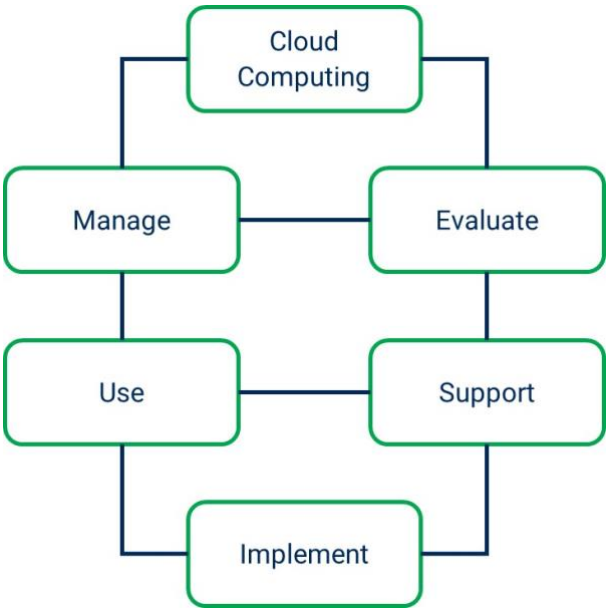

Figure 1 Cloud computing and information management

Exam preparation

For details of the exam, please refer to the Preparation Guide. You can download the Preparation Guide and a Sample Exam from the EXIN website (exin.com) for free. You can do the Sample Exam online as well.

In this workbook you will find questions to check your knowledge at the end of each chapter. This will help increase your understanding and retention of the information. The questions are different from the exam questions. For a feel of the exam questions, please use the sample exam. Additionally, you will find an overview of terms with which you should be familiar at the end of each chapter.

1 The principles of cloud computing

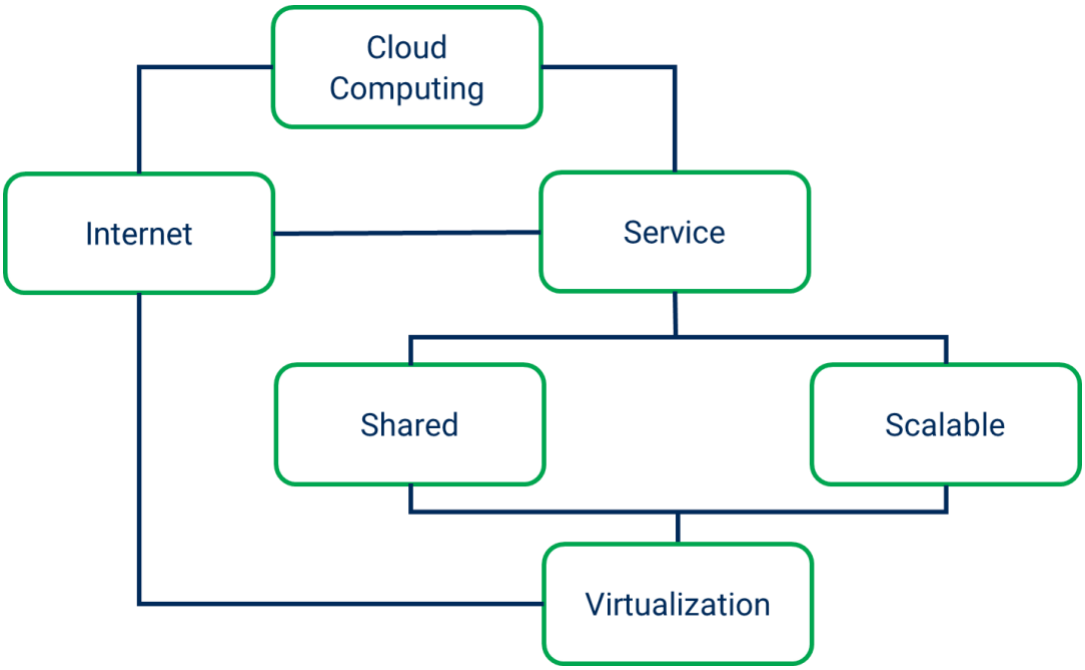

Figure 2 Overview of cloud computing

Cloud computing, method of running application software and storing related data in central computer systems and providing customers or other users access to them through the Internet.

Encyclopaedia Britannica (www.britannica.com, as viewed on 10-12-2019)

1.1 The concept of cloud computing

To understand the concept of cloud computing we need to know how it has evolved. To build a context around cloud computing, it is useful to look at another definition. According to the American 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 (for example, networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. This cloud model is composed of five essential characteristics, three service models, and four deployment models.

NIST Special publication 800-145 (September 2011)

This definition has been approved or adopted by many organizations in the industry, for example by:

- the International Standards Organization (ISO); reference:
  - ISO/IEC 22123-1 – Information technology -- cloud computing – part 1: Vocabulary
  - ISO/IEC 22123-2 – Information technology -- cloud computing – part 2: Concepts

Cloud computing is a paradigm for enabling network access to a scalable and elastic pool of shareable physical or virtual resources with self-service provisioning and administration on-demand.

Note 1 to entry: Examples of resources include servers, operating systems, networks, software, applications, and storage equipment.

Note 2 to entry: Self-service provisioning refers to the provisioning of resources provided to cloud services (3.1.2) performed by cloud service customers (3.3.2) through automated means.

ISO/IEC 22123-1

- the International Telecommunication Union (ITU); reference:
  - Recommendation ITU-T Y.3500 – Information technology – Cloud computing – Overview and vocabulary (2014)

For this workbook we will use the NIST definition as a basis. NIST states that there are

- five essential characteristics
- three service models
- four deployment models

1.1.1 Essential cloud characteristics

1.1.1.1 The client-facing perspective

The NIST publication SP 800-145, which was published in 2011, mentions five essential characteristics:

- on-demand self service
- broad network access
- multi-tenancy and resource pooling
- rapid elasticity and scalability
- measured service

Because cloud technologies have evolved since then we have updated this original list to:

- **On-demand self-service** Within an existing contract, a user/customer can for example add new services, storage space or computing power without a formal request for change.
- A wide range of **secure connectivity options** (via Internet, leased lines, and so forth)
- **Rapid elasticity and scalability** This characteristic has to do with the fundamental cloud aspects of flexibility and scalability.
  - For example, web shops need a standard amount of transaction ability during the year but need to peak around Christmas. Of-course they do not want to pay for this peak ability during the rest of the year.
- **Pay-per-use** This means monitored, controlled, and reported services. This characteristic enables a pay-per-use service model. It has similarities to the mobile telephone concept of service bundles, where you pay a standard subscription for basic levels, and pay extra for additional service without changing the contract.

- **Resource pooling** In the industry, this characteristic is also known as multi-tenancy. Many users/customers share a varied type and level of resources.

The list above is not complete without a famous quote from Bill Gates, former CEO of Microsoft. Broad network access is what he envisioned in the late nineties:

"any time, any place, and any device"

The ISO/IEC 22123-1 standard recognizes another characteristic of cloud computing:

- **Multi-tenancy** A feature where physical or virtual resources are allocated in such a way that multiple tenants and their computations and data are isolated from and inaccessible to one another.

Typically, and within the context of multi-tenancy, the group of cloud service users that form a tenant will all belong to the same cloud service customer organization. There might be cases where the group of cloud service users involves users from multiple different cloud service customers, particularly in the case of public cloud and community cloud deployments. However, a given cloud service customer organization might have many different tenancies with a single cloud service provider representing different groups within the organization.

1.1.1.2 The business perspective

From a client-facing perspective, the characteristics above drive the cloud service. From a business perspective, essential cloud characteristics are:

- predefined, comprehensive set of IT services and options
- known upfront price and quality, monitored, controlled and reported by provider
- can be ordered or changed by a 'simple click' in the self-service portal
- allows rapid, elastic and flexible capacity scale-up/scale-down over a period
- can be consumed any time, any place, any device via the Internet or other network options
- based on resource sharing of resources by many users/customers under the multi-tenancy concept
- only the ordered or consumed capacity is billed

1.1.2 The main cloud deployment models

NIST furthermore defines four deployment models: private cloud, Community cloud, public cloud and hybrid cloud, and three main Service Models: Infrastructure as a Service (IaaS), Platform as a Service (PaaS) and Software as a Service (SaaS). The deployment models will be discussed in the next paragraph, and the three Service models will be discussed in paragraph 1.1.3.

"There was a time when every household, town, farm or village had its own water well. Today, shared public utilities give us access to clean water by simply turning on the tap; cloud computing works in a similar fashion. Just like water from the tap in your kitchen, cloud computing services can be turned on or off quickly as needed. Like at the water company, there is a team of dedicated professionals making sure the service provided is safe, secure and available on a 24/7 basis. When the tap isn't on, not only are you saving water, but you aren't paying for resources you don't currently need."

Vivek Kundra, Federal CIO, United States Government

It is tempting to talk about "the cloud" as one single environment. However, studying all cloud phenomena must lead to the conclusion that there are many different clouds. For a normal end-user of web-based services like social media, webmail, online storage, collaboration software, blogs, video calling and so on, there seems to be only one cloud: the World Wide Web.

But for enterprises, public sector organizations and non-government organizations it is more of a cloudy sky. Most enterprises have their own IT infrastructure, buy IT services from service providers and use several web services. For example, an organization's e-mail infrastructure typically consists of (an array of) SMTP servers on which every employee has their own mailbox, calendar and corporate contact list. However, most employees will also have access to a webmail account connecting them to the corporate mail server through the Internet. In this case there obviously is a private and a public aspect.

In the industry there are four types of cloud deployment models that are generally accepted; most prominently by the American National Institute of Standards and Technology (NIST).

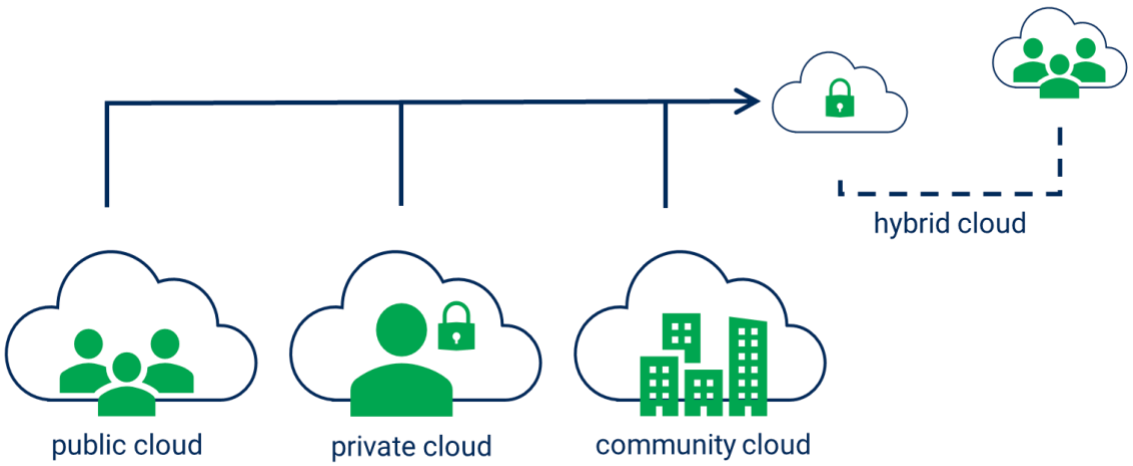

Figure 3 Four cloud deployment models

1.1.2.1 The private cloud

The main characteristic of a private cloud is that it resides on a private network that runs on (part of) a data center. The data center may be owned, managed and run by either the organization itself, a third party or a combination of the two.

Services are delivered to the different parts of the organization: for example, to its business units and internal departments like human resources and finance. The goal is to support the organization's business objectives in an economically sound way, but more important in a secure way. A private cloud solution is usually chosen when there is a need to comply with external regulations and legislation like the Sarbanes-Oxley (SOX) Act causing the need for a high degree of governance.

The downside of this solution is that there is still a high degree of total cost of ownership (TCO), resource pooling and rapid elasticity. An organization must consider if it really needs access at any time, from any place and from any device: for instance, a cloud or just a shared service center. All types of cloud have the before mentioned five characteristics. If these characteristics are missing, the services are just classical examples of hosting or ASP.

1.1.2.2 The public cloud

A public cloud is a more compelling example of what is intended with cloud computing, namely the delivery of off-site services over the Internet. Key characteristics are

- the sharing of resources, which lowers the total cost of ownership (TCO),
- and high flexibility and scalability of capacity, also called elasticity.

The downside of this sharing principle of multi-tenancy is a lower level of security and privacy making it more difficult to comply with different types of international legislation. Sharing basic

infrastructure like storage, data base servers or applications can all cause security, data protection, and privacy issues.

For example, when a multinational chain of retail stores with their head office in the Netherlands decided to migrate their Office applications to the G Suite: Collaboration & Productivity Apps for Business, their main requirement was that their data would be stored on a data center within the European Union.

For most private users or the main public these concerns are less relevant or not appreciated. For this target group the cloud is Facebook, Twitter, Skype, Dropbox, webmail and all the other examples of Internet based offerings that make life more productive or fun.

1.1.2.3 The community cloud

The community cloud has many similarities to the private cloud, because it delivers services to a specific group of organizations or individuals that share a common goal. Some examples are regional or national educational or research institutes, community centers or even commercial organizations wishing to share very high-security facilities for transaction processing (like stock exchange trading companies).

The main goals of creating a community cloud are the ease of sharing:

- data
- platforms
- applications, which otherwise would be too expensive to purchase
    - for example, university research applications

Another goal of sharing cloud facilities within your own community may be to reduce costs, improve performance, and improve privacy and security, without raising the total cost of ownership (TCO) in a significant way. Some cloud advantages can simply not be gained by running your own local computing facilities. For example:

- 24/7 access and support
- shared service and support contracts
- the economics of scale

1.1.2.4 The hybrid cloud and multicloud

It is important to notice, that many customers with long term cloud experience have recognized that there is no one size fits all solution with only one “cloud”, and prefer to follow a “multicloud” strategy which addresses different kind of needs and workloads via different cloud services and/or vendors. A multicloud deployment often consists of a mix of one or more of the other deployment models.

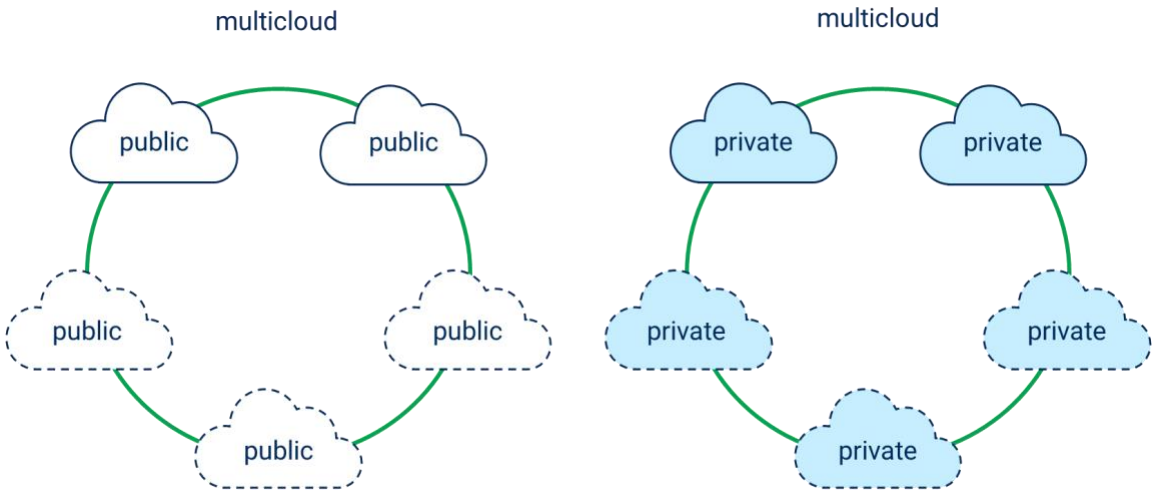

Figure 4 Multicloud implementations

The multicloud strategy[1]:

- reduces the dependency on any single vendor
- allows getting the best-of-breed market offerings for
  - different kinds of workload,
  - different geographies (to reduce latency or regulation risks, for example data location requirements for a specific country)
  - other options
- increases flexibility through its richness of choice

The multicloud usually does not assume that cloud deployment models are mixed. However, when the implementation consists of a mix of the models mentioned above, this is called a hybrid cloud.

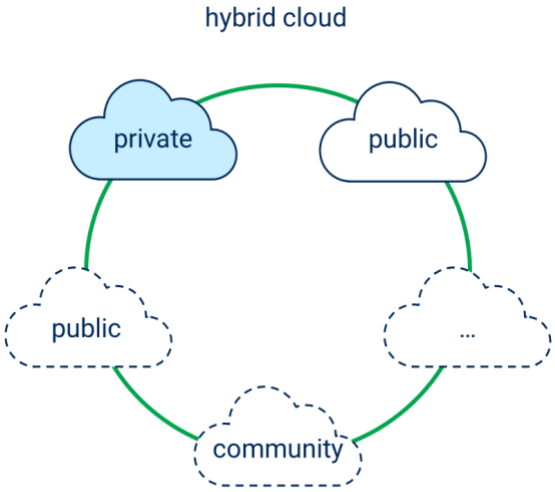

Figure 5 The hybrid cloud model

[1] If you are interested in extra details on this topic you can find several articles on the Internet. An example of an online article, evaluating multicloud details, can be found on https://www.redhat.com/en/topics/cloud-computing/what-is-multicloud.

Putting it simply, a hybrid cloud[2] is a mix of the models discussed above. It combines several private and public cloud solutions from several providers into one IT infrastructure. In this model, a clear choice must be made as to what to buy where. Choosing specific services for either private or public cloud suitability is balancing security, privacy, and compliance versus price.

Let us look at an example. A large multinational made the choice to go "the hybrid way". The mission critical logistics and enterprise resource planning (ERP) systems run on a private cloud solution, while the common office applications are fulfilled by the Google Apps for business solution. This brings savings of many millions of dollars per year and does not compromise the integrity of the core business services.

Another example could be the way insurance companies work with insurance agents. There is a lot of interaction between the two sides, but the company's and agents' infrastructures cannot and will not be integrated in the traditional way. In this scenario a public cloud extension could build a bridge between the different infrastructures.

Modern hybrid cloud solutions were born with another cloud attribute associated with multicloud solutions:

- automated cloud brokering with a unified single self-service window for end-users.

This takes away the extra challenges associated with such solutions – variations of different self-service portals, service titles, attributes, options, characteristics, and pricing across different vendors.

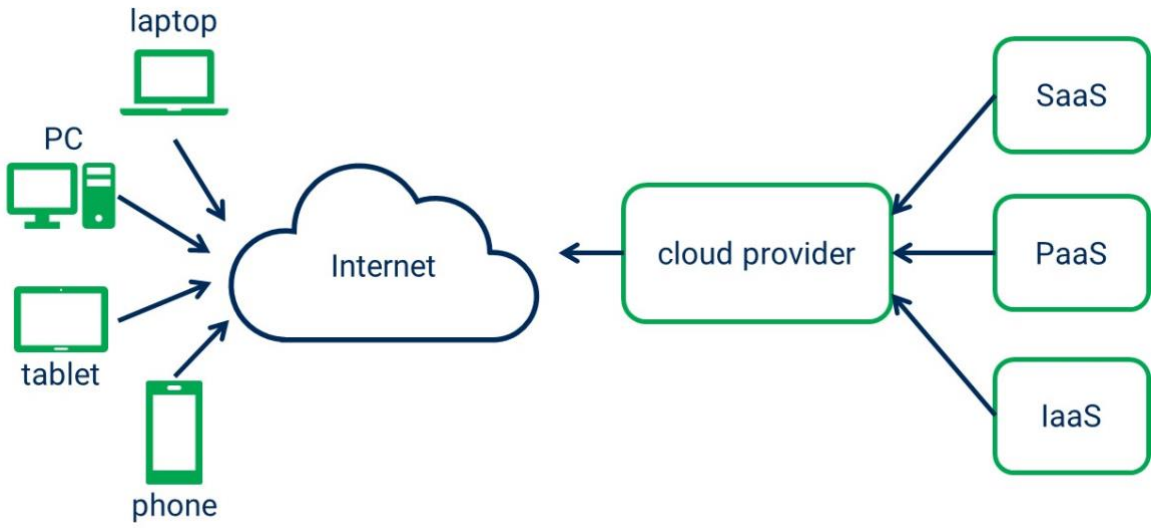

Figure 6 A new world for delivery of IT services

[2] If you are interested in this topic, an example of an online article evaluating hybrid cloud details, can be found on https://en.wikipedia.org/wiki/cloud_computing#hybrid_cloud.

1.1.3 Service models for cloud computing

There are many types of cloud services like webmail, Hosted Exchange, online storage, online backup, social media, and so forth. All these services can be grouped under three main cloud service models: Software as a Service (SaaS), Platform as a Service (PaaS) and Infrastructure as a Service (IaaS). In this section we will briefly describe these three models.

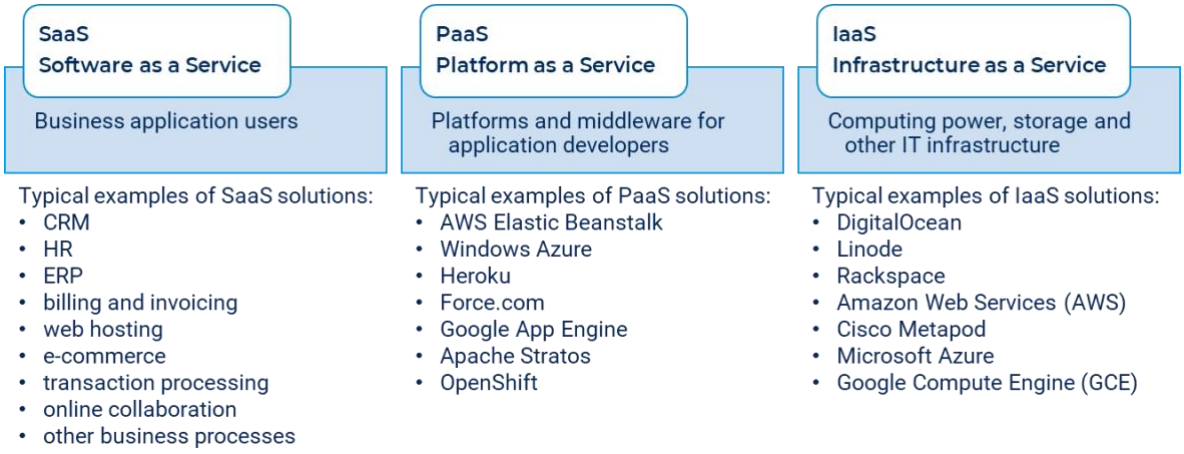

Figure 7 Service models for cloud computing

1.1.3.1 Software as a Service (SaaS)

This is the most common type of cloud service. SaaS is a software delivery methodology that provides licensed multi-tenant access to software and it functions remotely as a web-based service. SaaS is a break with the tradition that organizations buy or develop their own business applications and run and manage them on their own IT infrastructure.

Application hosting by third parties goes back to the mainframe days and came into maturity with the application service provider (ASP) industry that emerged in the early 2000s. SaaS essentially extends the idea of the ASP model. Many types of SaaS services were developed from ASP solutions, for example:

- application hosting
- pay-per-license
- emulation
- terminal services

These ASP solutions were turned into cloud solutions, for example:

- multi-tenancy
- pay-per-use
- web-based interfaces
- elasticity and elastic pricing

The key benefits of SaaS are that the customer does not need to worry about the development and management of these applications. The provider is responsible for updates and managing licenses, and most service management parameters like scalability, availability, maintenance and service continuity. A customer pays by means of a subscription or pay-per-use model.

Subtypes, or maybe simply other names for SaaS, are:

- software on demand
- hosted services
- application service provisioning (ASP)

Database as a Service (DBaaS) has emerged as a subvariety of SaaS.

SaaS key characteristics
Minimal IT organization engagement requirements (might be ordered and consumed without IT organization at all in many cases)
Software hosted offsite
Software on demand
Software package
No modification of the software
Plug-in software: external software used with internal applications (hybrid cloud)
Vendor with advanced technical knowledge
User entangled with vendor

1.1.3.2 Platform as a Service (PaaS)

There are benefits to not owning a computer platform. However, the benefits of being able to use it *on* demand for a fraction of the price are great. For example, you can use a computer platform as a cloud service when there is a need to test customized applications.

PaaS can save costs in ownership, management, and maintenance. PaaS can also inject other cloud-associated characteristics like rapid elasticity and pay-per-use into the hosted environment of PaaS applications.

Typical software development environment platforms are only used for the time the project runs. This is a benefit because a new project often has other or newer platform requirements. During some stages of the development process, for example during testing, there usually is a need for an up-scaled environment to simulate a production environment.

At the same time, application computing capacity demands may typically vary over time and in line with business activities that require an up-scaled environment to absorb extra business workloads, for example due to promo-actions, financial months/quarters/years closures, or seasonal variations.

PaaS services can offer this on-demand scalability for any case where there is only a need to pay for extra capacity during up-scaling. There are several variants of PaaS.

PaaS variant	Summary
Public PaaS	This variant is derived from Software as a Service (SaaS) and is in between SaaS and Infrastructure as a Service (IaaS).
Private PaaS	Typically, this variant can be downloaded and installed either on a company's on-premises infrastructure, or more often in a public cloud. Once the software is installed on one or more machines, the private PaaS arranges the application and database components into a single hosting platform.
Hybrid PaaS	This variant is able to “register” multiple, distinct cloud infrastructures as independent pools and merge those identifiably different pools into a single resource pool. This leads to resource normalization while still preserving identity of origin.
Mobile PaaS	MPaaS provides development capabilities for mobile app designers and developers.
Open PaaS	This variant does not include hosting, but rather it provides open-source software allowing a PaaS provider to run applications in an open-source environment.

It is important to highlight PaaS for rapid development because “Enterprise Public Cloud Platform for Rapid Development” is defined by Forrester Research (Forester report for CIOS December 29, 2014) as an emerging trend. The most common deployments of PaaS are as follows.

PaaS deployment	Summary
Software development environment	A customer can develop an application without having to buy a dedicated development environment, and without having to configure and manage the underlying infrastructure components like hardware, middleware and the different software layers. Microsoft Azure and the Google App engine are examples of such a service.
Hosting environment for applications	This service only consists of services at the hosting level like security and on-demand scalability.
Online storage	Cloud solutions, because of their architecture with Storage Area Network (SAN) servers, not only offer online storage but also extremely rapid data exchange between instances of online storage.

Some examples of PaaS service providers are Force.com (the first PaaS provider), Google with its App Engine, OpenShift and smaller players like Heroku.

PaaS key characteristics
Requires strong engagement of the IT organization for ordering and consumption (application development and management teams)
Used for remote application development and for developed applications hosting
Remote platform support
Platform may have special features and/or limited platform modification options
Low development costs
Low TCO and high agility for custom applications (platform injected benefits)

1.1.3.3 Infrastructure as a Service (IaaS)

IaaS provides access to computing resources in a virtualized environment (the cloud), across a public connection, usually the internet. The definition includes such offerings as

- physical and virtual server
- storage space
- network connections
- bandwidth
- IP addresses
- load balancers

Physically, the pool of hardware resources is pulled from a multitude of servers and networks, which are usually distributed across numerous data centers, all of which the cloud provider is responsible for maintaining. The client is given access to the virtualized components in order to build their own IT platforms.

IaaS services are sold by so-called hardware service providers from which a customer can rent physical or virtual hardware like storage, servers, or Internet connectivity. Services are sold according to a utility computing service and billing model. The background of IaaS can be found in the merger between IT and telecom infrastructure and services in the past decade.

Perhaps the best-known provider of IaaS is Amazon with their Amazon Web Services (AWS), Elastic Compute Cloud (E2C), and Simple Storage Service (S3). Amazon AWS can take over the management of your server, storage, network services, and virtualization needs.

Other examples of IaaS are Microsoft Azure, Oracle Cloud Infrastructure (OCI), Google Compute Engine, and IBM Cloud (e.g. transient virtual servers).

Many on demand IaaS infrastructures are built on components from leading vendors like Cisco, HP, NetApp, and VMware.

With the IaaS model, IT infrastructures and services that were previously only available to large enterprises are now within the reach of smaller business such as small to medium enterprises/businesses (SME/SMB).

IaaS key characteristics
Requires strong engagement of IT organization for ordering and consumption (applications and infrastructure management teams)
Dynamic scaling
Most flexible and vendor-agnostic cloud service model
May address “one stop cloud shop” strategy in many cases
Desktop virtualization
Policy-based services

1.1.3.4 Anything as a Service (XaaS)

Anything as a service, or XaaS, refers to the growing diversity of services available over the Internet via cloud computing as opposed to being provided locally, or on premises. Also known as Everything as a Service, Anything as a Service reflects the vast potential for on-demand cloud services and is already being heavily marketed and promoted by companies like VMware and HP.

Anything as a Service derives the "X" in its XaaS acronym from being a catch-all term for everything. Examples are:

- Storage as a service (SaaS)
- Desktop as a Service (DaaS)
- Disaster Recovery as a Service (DRaaS)
- Network as a Service (NaaS)

There are even emerging services such as Marketing as a Service and Healthcare as a Service.

Special attention might be paid to the following emerging technologies under the XaaS umbrella:

- Quantum Computing as a Service (QCaaS)
- Artificial Intelligence as a Service (AIaaS)
- Blockchain as a Service (BaaS)

These emerging cloud solutions can provide easy options to implement these technologies without big investments and long deployments.

Occasionally, one can come across:

- Backup as a Service (BaaS)
- Business Process as a Service (BPaaS)
- Container as a Service (CaaS); e.g. services like Docker
- Communications as a Service (CaaS)
- Data as a Service (DaaS)
- Database as a Service (DBaaS)
- Identity as a Service (IDaaS)
- Monitoring as a Service (MaaS)
- Security as a Service (SecaaS)

Depending on the point of view they can be classified within the classic three (IaaS, PaaS, and SaaS). The following table maps cloud service categories onto the three main cloud service types.

Cloud service categories	Cloud capabilities types		
	Infrastructure	Platform	Application
Compute as a Service	✓		
Communications as a Service		✓	✓
Data Storage as a Service	✓	✓	✓
Infrastructure as a service	✓		
Network as a Service	✓	✓	✓
Platform as a Service		✓	
Software as a Service			✓
Source: ISO/IEC 22123-1			

1.2 How cloud computing evolved

Cloud computing is often compared to the supply models of the electrical grid or water companies. Computing power on tap, turn it on when you need it, and turn it off when you don't. You only pay for use and for the connection to the supply.

This comparison between computing and common utilities was made as early as 1996 by Douglas Parkhill in his book *The Challenge of Computer Utility* (1996). He stated that when the electrical grid was built it quickly replaced all the small private plants and providers. Large scale economics, in combination with security and safety, won the day. This theme was discussed in more detail by Nicolas Carr in his book *The Big Switch: Rewiring the World, from Edison to Google* (2008).

1.3 Historic timeline

In the beginning there was the mainframe computer, then there was the first network, and the rest is history.

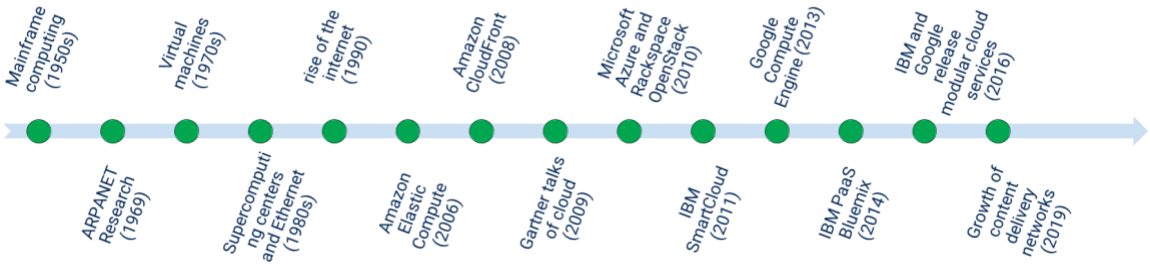

Figure 8 Timeline for cloud computing

Several key factors have contributed to the present-day existence of the cloud:

- the development of the Internet
- the move from mainframe computing to the presence of multiple personal devices with connection to the Internet
- the development of computer networks
- virtualization
- the Internet
- application service provision or hosted services (ASP, a predecessor of SaaS)

1.3.1 Networks, servers, and all that stuff

The first mainframe computers were not connected to, what we nowadays call, a network. At first, they had a point-to-point connection with a terminal. The first terminals did not even have a monitor.

When business applications started to run on mainframes, users got their own terminals, but these were not interconnected. All data processing and information sharing was done on the mainframe computer. When decentralized computing started to appear the first so-called minicomputers were developed. The first generation of these decentralized computers with their own attached terminals were called minicomputers (or sometimes front-end processor).
These decentralized computers were connected through local (LAN) or wide area (WAN) connections to the central mainframe computer.

When IBM started selling their first so-called microcomputer, the IBM PC, it was predicted that there would only be a need for a few of these per office, and most likely stand-alone.

We now know that almost immediately there came a need to connect these PCs to the central computers through a network. Instead of having a mainframe terminal plus a PC on your desk, you could have a PC connected to a network and access the mainframe with a terminal emulation program. Several network topologies were developed, many with their own protocol. Common topologies that can be found in study books are:

- point-to-point
- bus
- star
- ring or circular
- mesh
- tree
- hybrid
- daisy chain

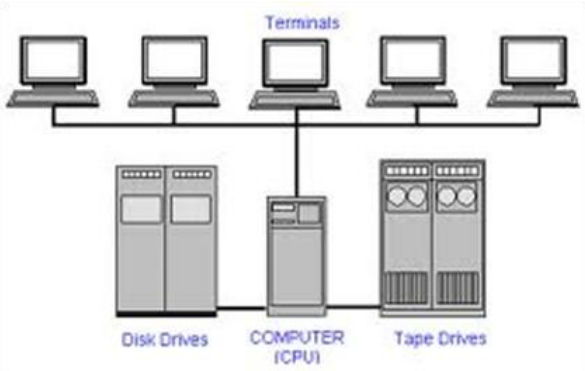

Figure 9: Mainframe with terminals and peripherals.

The most common topology today is the star topology in combination with the TCP/IP protocol. It is important to realize that the so-called Transmission Control Protocol (TCP) / Internet protocol (IP) stack is the core protocol of the Internet.

From PCs connected to the mainframe the next development was the client-server architecture. PCs were now able to connect to several different minicomputers called servers, for example file servers, or application servers.

With ever-growing bandwidth and speed of the networks, server speed and capacity, and ever-cheaper and smaller personal devices to connect to the networks, we entered the age of the Internet and application hosting by application service providers (ASP). One of the present-day variants of SaaS (Software as a Service) is a direct derivative of these ASP solutions.

Figure 10 Internet

Figure 11 Mainframe computers

1.3.2 The role of the Internet

In 1963, J.C.R. Licklider, an American computer scientist, was the Director of Behavioral Sciences Command & Control Research at the U. S. Department of Defense Advanced Research Projects Agency (ARPA). On April 25 of that year, Licklider sent a memo to his colleagues explaining his vision on a global network. The memo was called: *Memorandum for Members and Affiliates of the Intergalactic Computer Network*.

This vision was later realized in 1969 in the form of ARPANET, the direct predecessor to the Internet. ARPANET was the world's first operational packet-switching network. Since it was designed for the USA army, it was not (yet) public at that time.

The original network protocol NCP was replaced by TCP/IP in 1983 and remains the leading protocol to the present day. Based on the address restriction of the Internet protocol v4, the Internet layer has essentially to be replaced by IPv6 in the layer model. The remaining layers of this model remain relatively untouched, unless there are more addresses included in the application layer, as is the case with FTP.

The Internet is sometimes called the World Wide Web (www), but that is just the name of one of the many services that run over the Internet like FTP (data transfer) and SMTP (e-mail). To make the Internet accessible to everyone, first there needed to be two other developments:

- personal devices
- network connectivity

These developments happened in parallel and started with terminals connected to a mainframe, giving access to central computing facilities and applications.

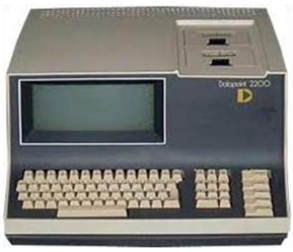

Figure 12 Computer

1.3.3 Virtualization

Virtualization refers to the act of creating a virtual (rather than actual) version of something, including (but not limited to) a virtual computer hardware platform, operating system (OS), storage device, or computer network resources. Virtualization began in the 1960s, as a method of logically dividing the system resources provided by mainframe computers between different applications.

In his book *Virtualization: A Manager's Guide*, Daniel Kuznetzky (2011) describes the earliest examples of virtualization:

> The earliest form of application virtualization was developed by mainframe suppliers, such as IBM.

An example is the IBM VM/370 from 1972. Since the 1990s, Windows has also become the center of virtualization developments by Microsoft, Citrix, VMware and others.

Kuznetzky (2011) recognizes five different types of virtualization:

Type of virtualization	Summary
Access virtualization	Allows access to any application from any device
Application virtualization	Enables applications to run on many different operating systems and hardware platforms
Processing virtualization	Makes one system seem like many, or many seem like one
Network virtualization	Presents an artificial view of the network that differs from the physical reality
Storage virtualization	Allows many systems to share the same storage devices, enables concealing the location of storage systems, and more

For the average cloud user this means that hardware, applications and data can be located anywhere in the cloud; we only need to access and use them.

Virtualization is the solution for integration of:

- high-speed computers
- large storage capacity
- Internet

Virtualization key features
Virtualization multiplies the use of high-performance computers.
With today's modern processors there is a vast amount of processor availability. Virtualization puts extra capacity to use.
Concept of the cloud: virtualized operating environment & thin clients; web-based delivery
Multi-tenancy

1.3.4 Managed services

Managed services are the practice of outsourcing day-to-day management responsibilities and functions as a strategic method for improving operations and cutting expenses. Having your services managed basically means that you turn IT services over to a third party.

An old example, since mainframe days, is application hosting by an IT provider. In the late 1990s applications were offered by providers, meaning that they were no longer owned by the customer. This first example of shared managed services was delivered by ASP's (application service providers). As the 'Internet bubble' burst in the early 2000s ASPs never became a big thing and slowly developed into one of the major cloud service models: SaaS.

In parallel with these managed services, the need for a good service management framework arose. With IBM's so-called Redbooks (which contains operation guides et cetera) as a basis and the adoption of many industry best practices, the IT Infrastructure Library (ITIL) framework for IT service management was developed in the early 1970s.

Key internal ITIL processes for a cloud data center include

- availability management
- capacity management
- security management
- service continuity management

External processes include

- service level management
  - maintaining, managing, reporting on and improving service levels sold to and agreed with the customer
- financial management

As a result of managed services, the customer's own IT department can shift their focus away from operational issues. They no longer have to worry about constant server updates and other maintenance issues. However, this does not mean that the Chief Information Officer (CIO) can lean back and do nothing. Instead, the focus will have to be shifted to IT governance.

Key issues for IT governance are:

Issue	Related question
Performance	Can the cloud services support our business model, also when it is transforming?
Compliance	Do the services comply with our own national and international legislation?
Contingency	What happens if the cloud provider goes out of business?
Financial Performance	Can we leverage all expected financial benefits and avoid significant wastes on cloud spending?

1.3.4.1 Audit standards and guidelines

How can a customer stay 'in the driver's seat'? It seems likely that there will be an increased need for audit models focusing on the IT service management processes as well as data center performance and compliance issues. Data centers with a history of platform and application hosting often use the following ISO/IEC audit standards and guidelines for their internal and external audit mechanisms[3].

ISO/IEC 20000 'family' IT Service Management	
ISO/IEC 20000-1	Information technology -- Service management -- Part 1: Service management system requirements
ISO/IEC 20000-2	Information technology -- Service management -- Part 2: Guidance on the implementation of service management systems
ISO/IEC 27000 'family' Information Security	
ISO/IEC 27001	Information security, cybersecurity and privacy protection — Information security management systems — Requirements
ISO/IEC 27002	Information security, cybersecurity and privacy protection — Information security controls
ISO/IEC 27017	Information technology – Security techniques — Code of practice for information security controls based on ISO/IEC 27002 for cloud services
ISO/IEC 27018	Information technology – Security techniques — Code of practice for protection of personally identifiable information (PII) in public clouds acting as PII processors
ISO/IEC 27036-4	Information technology — Security techniques — Information security for supplier relationships — Part 4: Guidelines for security of cloud services

There are several other ISO/IEC standards which address cloud computing that may be useful to cloud providers.

Name	Summary
ISO/IEC 22123-1 ISO/IEC 22123-2	provide an overview of cloud computing along with a set of terms and definitions (vocabulary and concepts).
ISO/IEC 22123-3	specifies the cloud computing reference architecture (CCRA).
ISO/IEC 19944-1	extends the existing cloud computing vocabulary and reference architecture in ISO/IEC 22123-1, ISO/IEC 22123-2, and ISO/IEC 22123-3 to describe an ecosystem involving devices using cloud services, describes the various types of data flowing within the devices and cloud computing.
ISO/IEC 19941	specifies cloud computing interoperability and portability types, the relationship and interactions between these two cross-cutting aspects of cloud computing and common terminology and concepts used to discuss interoperability and portability, particularly relating to cloud services.

For customers of cloud services, good governance practices will be of increasing importance, and this will place more focus on the following international *standard*s and frameworks for corporate *governance* of information technology.

[3] Source: https://www.iso.org/home.html and https://www.isaca.org/

Standard or framework	
COBIT	Guidance for executive management to govern IT within the enterprise
ISO/IEC 38500	Information technology -- Governance of IT for the organization

Even now, the purchase of bare metal (physical servers and other hardware) can be emulated in commercial cloud. An example would be, billing by usage or physical server billing by the hour. As a cloud solution, a bare-metal server request with all the resources needed, and nothing more, can be delivered within a matter of hours.

The cloud story is not finished yet. The evolution of cloud computing has only just begun.

1.3.5 Recent and future developments

Considering the ever-accelerating pace with which cloud computing is developing it is difficult to predict any future. However, some recent developments are really compelling.

Artificial Intelligence as a Service (AIaaS); this development is revolutionizing data science and transforming business processes (for example, the IBM IA services and their Watson AI platform).

Another new development is edge computing:

> Edge computing is a distributed computing paradigm which brings computation and data storage closer to the location where it is needed, to improve response times and save bandwidth
>
> https://en.wikipedia.org/wiki/Edge_computing, as seen on 28-05-2020

1.4 Cloud computing architectures

Two key architectural principles apply to cloud computing:

- multi-purpose architecture
- multi-tenancy architecture

In the past most architectures were proprietary and single purpose. Examples of such architectures are accounting systems and systems for the storage of healthcare data.

Multi-purpose infrastructure is key to cloud computing. An example is a system on which data is not only stored, but also distributed over the Internet.

1.4.1 Multi-purpose architecture

Virtualization is one of the key factors that contribute to the multi-purpose architecture principle. Many different types of implementations, applications, and workloads can run on the same platform or type of platform in a virtual environment. Virtualization makes it easy to guarantee scalability for all customers. Re-installing a new dedicated virtual platform is much quicker and easier than installing or re-installing a physical server.

Multi-purpose architecture key characteristics
Multi-tiered (different tiers for database, application and load balancing)
Virtualization (server)
Interoperable layers
Open standards
Portability

1.4.2 Multi-tenancy architecture

Multi-tenancy is an architecture in which a single instance of a software application serves multiple customers. Each customer is called a **tenant**. Tenants may be given the ability to customize some parts of the application, such as color of the user interface (UI) or business rules, but they cannot customize the application's code.

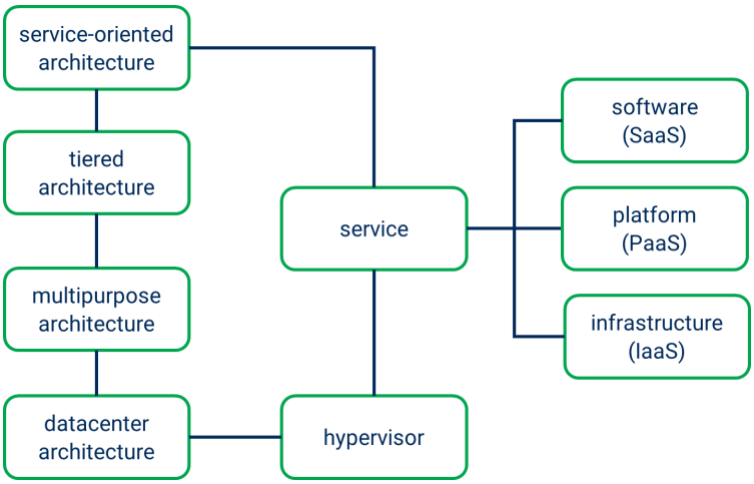

Figure 13 Multi-tenancy architecture.

Multi-tenancy can be economical because software development and maintenance costs are shared. It can be contrasted with single tenancy: an architecture in which each customer has their own software instance and may be given access to code. With a multi-tenancy architecture, the provider must make updates only once. With a single-tenancy architecture, the provider must touch multiple instances of the software in order to make updates.

In cloud computing, the meaning of multi-tenancy architecture has broadened because of new service models that take advantage of virtualization and remote access. A Software as a Service (SaaS) provider, for example, can run one instance of its application on one instance of a database and provide web access to multiple customers. In such a scenario, each tenant's data is isolated and remains invisible to other tenants.

In his article *Multi-Tenancy Misconceptions in Cloud Computing*, Srinivasan states the business case for multi-tenancy:

> …a large number of users, basically multi tenants, makes the cloud platform most efficient in terms of usability of the application and 'Do More with Less Resources.
>
> (Srinivasan 2011)

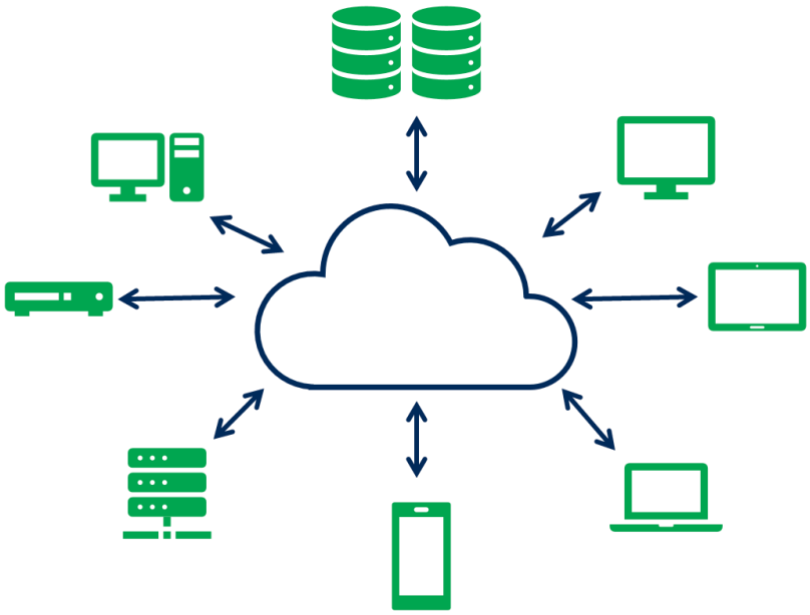

Figure 14 Multi-tenancy

Srinivasan also gives some examples of multi-tenant solutions:

- Salesforce.com: a SaaS-based CRM application for various businesses using common framework and multi tenancy model
- Microsoft Dynamics CRM Online offering
- multi-tenancy IaaS/PaaS offerings from Amazon or IBM or Microsoft Azure

A key element of multi-tenancy is **security**. If security cannot be guaranteed on all levels of the infrastructure, from the basic infrastructure to the web interface, customers will be hesitant to adopt this model.

1.4.3 Service-oriented architecture (SOA)

In his article *The Cloud-SOA Connection* (2009) Paul Krill quotes Jerry Cuomo, chief technical officer (CTO) of IBM's WebSphere business. Krill's question "Can we build a datacenter infrastructure on SOA principles?" is answered by Cuomo with:

> Yes, and that's the cloud, so it's a service-oriented infrastructure, […] It's taking that architectural principle of SOA and applying it to an infrastructure.

SOA is the instantiation of

- interoperability
- portability
- scalability

A service-oriented architecture (SOA) is basically a collection of services that communicate with each other. This communication may be simply data transferring between two or more services, or a jointly managed activity. In many cases, connecting services involves web services using XML. Historically SOA goes back to the CORBA[4] specification. One could say that there would be no cloud without SOA.

[4] CORBA stands for Common Object Request Broker Architecture and is a standard which allows objects from different programming languages and on different machines to interact.

The Open Group's SOA Work Group has come up with the following definition:

Service-Oriented Architecture (SOA) is an architectural style that supports service orientation.

Service orientation is a way of thinking in terms of services and service-based development and the outcomes of services.

A service:

- is a logical representation of a repeatable business activity that has a specified outcome (for example, check customer credit, provide weather data, consolidate drilling reports)
- is self-contained
- may be composed of other services
- is a "black box" to consumers of the service

SOA Work Group's Definition of SOA project. © The Open Group™

1.5 Benefits and limitations of cloud computing

Cloud computing is now evolving like never before, with companies of all shapes and sizes adapting to this new technology. Industry experts believe that this trend will only continue to grow and develop even further in the coming years. While cloud computing is undoubtedly beneficial for mid-size to large companies, it is not without its downsides, especially for smaller businesses.

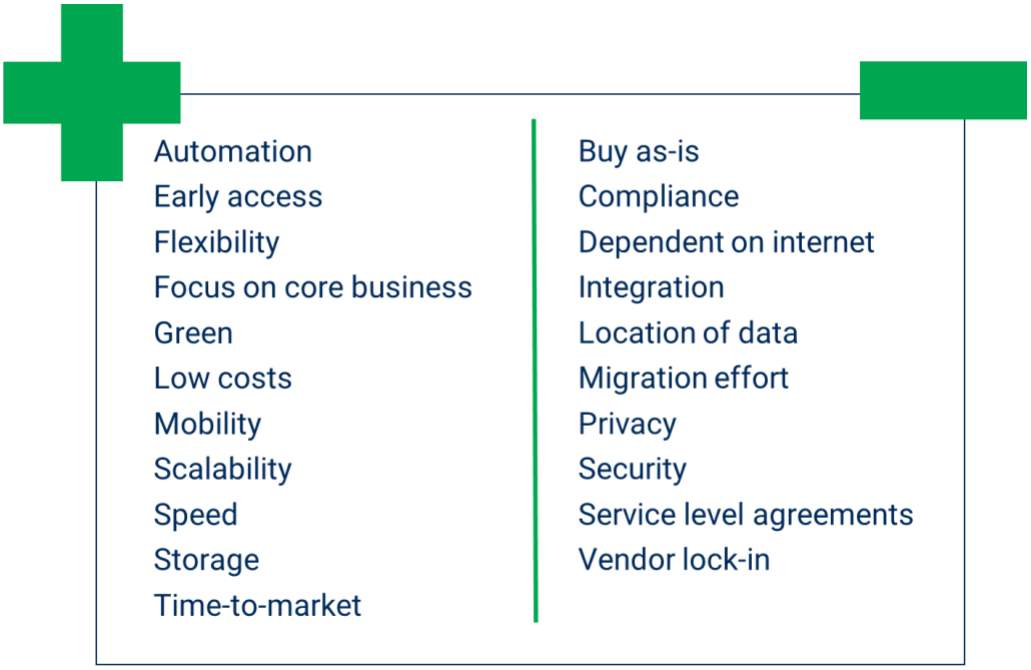

Figure 15 Benefits (left) and limitations (right) of cloud computing

1.5.1 Main benefits of cloud computing

If used properly and to the extent necessary, working with data in the cloud can vastly benefit all types of businesses. Mentioned below are some of the advantages of this technology:

Benefit	Summary
Reduced cost	Because of the pay-per-use or subscription model, organizations do not have to invest in IT infrastructure upfront. For cloud providers costs are lower because of the economics of scale and the multi-tenancy principle; no 'floor space' is left unused.
Speed and time-to-market	Unified and repeatable, fast cloud service deployments, scaling and provisioning within minutes of receiving a request seriously accelerates the speed of changes and reduces time to market for new solutions in the form of IT enabled business products.
Emerging technologies	Fast and easy access to cloud-based emerging technologies *as a Service*, without huge investment and under a pay-per-use model allows a quick and cost-efficient way for experiments with, and deployments of emerging technologies.
Automated	Updates, security patches and backups are no longer a concern of the customer. IT personnel does not need to worry about keeping software up to date.
Flexibility	Cloud computing offers more flexibility than legacy IT services. Within an existing or standard contract, a customer can change the 'cloud mix' of services in a dynamic way to support business demands and requirements.
More mobility	Data and applications can be accessed through the Internet from any type of smart computing device, anytime and anywhere.
Shared resources	Customers share resources allowing smaller organizations to have access to corporate scale IT facilities, services and supporting services. Users belonging to one or more customers can work together in a shared project environment.
Agility and scalability	Enterprises can scale their IT infrastructure up or down on demand.
Back to core business	Most businesses, especially start-up businesses, do not need to own and operate IT.
More IT functionality for a lower price	By virtue of the shared resources, more IT functionality is relatively cheap.

1.5.2 Main limitations of cloud computing

Despite the many benefits of cloud computing, there are also disadvantages. Businesses, especially smaller ones, should be aware of the limitations before using this technology.

Limitation	Summary
Internet access	Usually, no Internet access means no public cloud access. If Internet is not reliable, this is a risk.
Security	Cloud data centers can feature high security and be highly managed, but there are also low security and badly managed data centers. It is not always easy to differentiate between the two.
Privacy	In case of public or hybrid cloud offerings, it may be uncertain where your data is physically stored. This can be problematic in light of varying national and international legislation on privacy and data protection. It is difficult to know for sure who can access your data.
Buy as-is / Buy and Go	Do the provider's offering and options fit your needs? Is the provider ready to adjust the service offering for you? And for a reasonable price?
Service level agreement (SLA)	It is important to ensure that your SLA with the cloud provider allows for flexibility and scalability. If it does not, an important benefit of cloud is lost.
Performance	The SLA should also contain end-to-end performance commitments. These commitments should be both at your location and extending into the cloud. The vendor's multi-purpose cloud architecture or cloud technologies must be able to cope with the performance bottlenecks and limitations that apply to your specific workloads. If these conditions are not met, cloud does not have the full benefits for the business.
Integration	The SLA should allow for data and workflows integration across your whole IT environment. This must be done in a secure way (see security) and for a reasonable price.
Vendor lock-in	Migrating cloud services can be troublesome and risky. This could mean that the business stays with a provider that does not meet their needs, just to avoid the migration and a possible service interruption.

Exam preparation: chapter 1

'Get it' questions

1 / 7
What are the **4** types of cloud deployment models?

Give a summary or example of each of the four types.

2 / 7
Name **2** key characteristics of SaaS, PaaS and IaaS each. (6 key characteristics in total)

3 / 7
Which **3** developments enabled the emergence of cloud technology?

4 / 7
Which **2** other developments were crucial to make the Internet accessible to everyone?

5 / 7
What is a single purpose architecture?

Does cloud computing use single purpose architectures?

6 / 7
What does SOA mean?

7 / 7
Name **4** benefits and **4** limitations of cloud computing.

Exam terms

- cloud characteristics
- cloud computing architectures (multi-purpose, multi-tenancy)
- cloud computing
- deployment models (private, public, community and hybrid)
- drivers and limitations of cloud computing
- evolution of cloud computing
- Internet
- LAN
- managed services
- service models (SaaS, PaaS and IaaS)
- service-oriented architecture (SOA)
- virtualization

Answers to 'get it' questions

1 / 7
Your answer should list all cloud deployment models and contain a correct summary or example.

Cloud deployment model	Summary or example
Private cloud	A private network that runs on a data center that is exclusively used by one organization.
Public cloud	An Internet-based cloud service that is publicly accessible. Examples are free e-mail services, social media, and free cloud storage.
Community cloud	A shared private cloud, for instance by government organizations.
Hybrid cloud	Any mix of private, public and community cloud services.

See also section The main cloud deployment models.

2 / 7
Your answer should contain 2 of the following characteristics for SaaS, PaaS and IaaS each.

SaaS key characteristics
Minimal IT organization engagement requirements (might be ordered and consumed without IT organization at all in many cases)
Software hosted offsite
Software on demand
Software package
No modification of the software
Plug-in software: external software used with internal applications (hybrid cloud)
Vendor with advanced technical knowledge
User entangled with vendor
PaaS key characteristics
Requires strong engagement of the IT organization for ordering and consumption (application development and management teams)
Used for remote application development and for developed applications hosting
Remote platform support
Platform may have special features and/or limited platform modification options
Low development costs
Low TCO and high agility for custom applications (platform injected benefits)
IaaS key characteristics
Requires strong engagement of IT organization for ordering and consumption (applications and infrastructure management teams)
Dynamic scaling
Most flexible and vendor-agnostic cloud service model
May address "one stop cloud shop" strategy in many cases
Desktop virtualization
Policy-based services

See also section Identify the impact of cloud computing on the primary processes of an organization.

3 / 7
Most important developments:
- the development of computer networks
- the development of the Internet
- the move from Mainframe computing to personal devices with connection to the Internet

See also sections How cloud computing evolved and Historic timeline.

4 / 7
To make the Internet accessible to everyone, two developments were crucial:
- personal devices
- network connectivity

See also sections How cloud computing evolved and Historic timeline.

5 / 7
In the past, most architectures were proprietary and single purpose. Examples are accounting systems and storage of health care data. This infrastructure had only one purpose. Cloud computing infrastructure are multi-purpose. An example could be a system on which data is not only stored, but also distributed over the Internet.

See also section Multi-purpose architecture.

6 / 7
SOA stands for service-oriented architecture. This is an architectural style that supports service orientation. Service orientation is a way of thinking in terms of services and service-based development and the outcomes of services. This contrasts with thinking in terms of products.

See also section Service-oriented architecture (SOA).

7 / 7

Benefits	Limitations
Automation	Buy as-is
Early access	Compliance
Flexibility	Dependent on internet
Focus on core business	Integration
Green	Location of data
Low costs	Migration effort
Mobility	Privacy
Scalability	Security
Speed	Service level agreements
Storage	Vendor lock-in
Time-to-market	

See also section Benefits and limitations of cloud computing.

2 Implementing and managing cloud computing

2.1 Building a local private cloud environment

It can be said that a local or private cloud environment is not much different from the traditional data center. However, by employing modern cloud technology and architectures, large organizations can benefit from the best of both worlds, with both a private data center and a private cloud solution. A private data center provides you with complete control and your own internal security mechanisms.

Nowadays only large companies own their own data centers, due to the high costs involved.

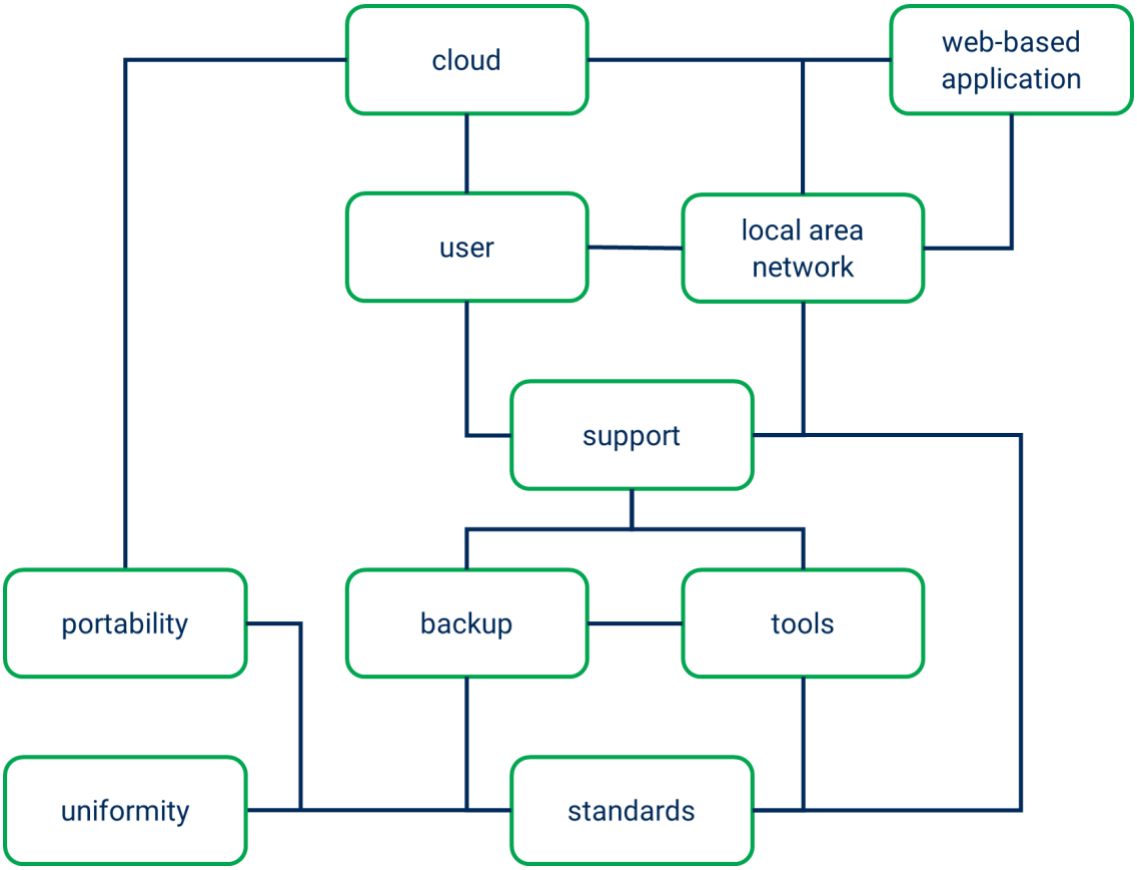

Figure 16 Local private cloud environment

2.1.1 Main components of a local cloud environment

This paragraph gives a quick overview of the main hardware and software components of a local cloud environment as well as some key performance criteria.

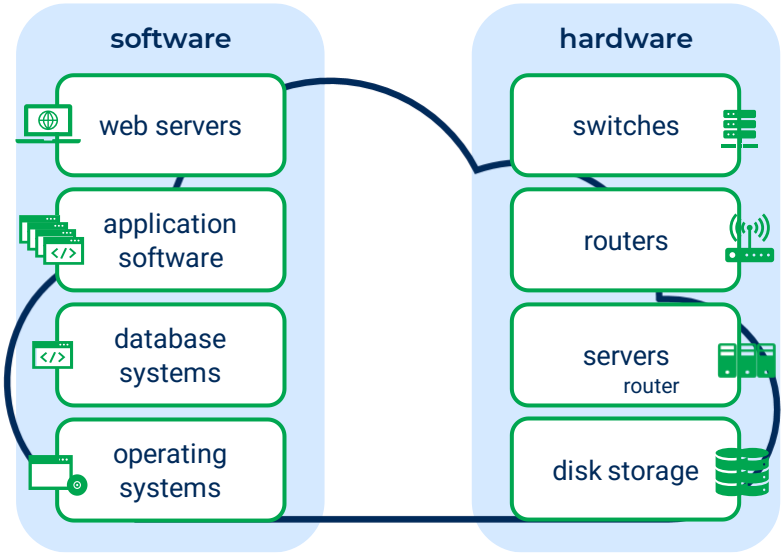

Figure 17 Components of a local cloud environment

2.1.1.1 Main hardware components of a local cloud environment

Hardware component	Summary
Blade server arrays	Server chassis housing multiple thin, modular electronic circuit boards, known as server blades. Each blade is a server (in its own right) and is therefore often dedicated to a single type of software, such as web hosting, virtualization, and cluster computing.
Local Area Network (LAN)	A network connecting local devices.
Storage Area Network (SAN)	A dedicated network that provides access to consolidated, data storage.
Network Attached Storage (NAS)	Storage that uses TCP/IP for data transfer in a network.
Backup & Restore solution	Tape libraries, usually supplemented by virtual tape libraries (often hard disk drive arrays), for cloud data backup and restoration.
Load balancer	Provides the means by which instances of applications can be provisioned and de-provisioned automatically, without requiring changes to the network or its configuration. It automatically handles the increases and decreases in capacity and adapts its distribution decisions based on the capacity available at the time a request is made.

2.1.1.2 Main software components of a local cloud environment

Software component	Summary or example
Cloud automation software	OpenStack, Red Hat Ansible
Virtualization software	VMware
Cloud application software	CRM, Office suite, ERP
Database software	Oracle, IBM DB2, Microsoft SQL
Middleware	software that connects software components or enterprise applications and is the software layer that lies between the operating system and the applications on each side of a distributed computer network
Operating systems	Microsoft, Linux, UNIX, macOS

2.1.2 Architectural considerations

To make the most use of the interoperability principle of cloud computing it is important to standardize the architecture. This can be done by using standard protocols and other building blocks that are location and vendor independent.

Some examples of these are: virtualization, SAN servers and blade servers, and load balancing. To manage the infrastructure, a central management console is needed. OpenStack is an example of a cloud operating system orchestrating all these elements.

The security and service continuity of a local cloud environment must also be considered. Ideally, the local cloud environment:

- consists of multiple sites to assist with disaster prevention and recovery (business continuity)
- uses proper backup mechanisms and data storage replication across sites
- uses common and high-security elements like firewalls, DMZ, security software, and role-based user profiles.

Successful cloud solutions require specific performance criteria. Some examples of these criteria are:

Components	Criteria
Physical components	Scalability of server and storage capacity
Storage	SAN performance (read, write, and delete times)
Internal processes	Connection speed, Deployment latency and Lag time

2.1.3 Sourcing strategy

There are two key options for cloud management or operations for the local (private) cloud.

Option	Summary
Outsourcing	Purchase a local (private) cloud service as a standard cloud offering from a cloud provider or from an independent IT company.
In-house	Have the internal IT organization manage the local (private) cloud.

2.1.3.1 Outsourcing

With the outsourcing option, all cloud service components, down to cloud assets ownership and provisioning, including procurement, shipment, and deployment, can be addressed by shared or dedicated providers organization. A shared providers organization is usually the less expensive choice. A dedicated providers organization can better focus on your specific requirements.

2.1.3.2 In-house options

With the in-house option all these components can be addressed by your own team or teams. This option also provides a choice. The business may wish to keep technical sources and tasks in-house to be performed by existing IT teams. Using existing IT teams is more suitable for small- to medium-scale private cloud installations.

The business can also choose to establish a special, dedicated, cloud-focused in-house IT organization, responsible only for the cloud environment. This choice is the most cost-effective and quality-efficient with a large-scale private cloud. This option only makes sense when there is enough workload for several teams, focused on specific cloud-embedded technologies or processes.

2.1.4 Cloud procurement and capacity reserves

With the local (private) cloud installations, special attention should be paid to the day-to-day cloud procurement process on which the cloud's fast scale-up provisioning is based. High velocity is required to procure and ship extra cloud capacity for fast and massive scale-ups. This must be supplemented by suitable capacity reserves to address minor changes and to absorb procurement lead time.

Major public cloud providers have quick procurement setups, and a huge install base[5]. When actual reserve needs (in relative figures) are very small and are scaling-up constantly in line with the cloud market growth, this should address most issues with increasing capacity. However, these providers may still have special clauses for massive scale-ups in their agreements, to transfer risks associated with procurement lead time to the customer side. A massive scale-up would for example be ordering more than 30% of the baseline capacity with less than 3 months' notice.

In case of cloud ownership and procurement by the customer, which is the usual setup for the private cloud, all such risks are at the customer side, so to get cloud-associated flexibility and speed, the following should be taken into consideration:

- **whole cloud capacity scale-up demands over time**, and speed expectations for the cloud solution, with respect to your business dynamics and expected differing business workloads
- **risk-sharing and transferring options for scale-ups**
  - for example, cloud sharing across a few businesses or companies with different scale-up demands over the time
  - leveraging cloud provider capabilities; e.g. procuring the private cloud as a service under a pay-per-use agreement with suitable terms for scale-ups in the contract
- **a special, fast cloud procurement setup**, with respect to identified demands, speed expectations, chosen risk-sharing and transferring options, and its influence on the cloud capacity reserve
  - long procurement time drives up reserve needs and vice versa
- **upfront cloud capacity reserve**, which will support the approach taken

[5] Install base, or installed based, is a measure of how many units of a product or service are used.

2.1.5 Secured access

A Virtual Private Network (VPN) is necessary to grant secure and efficient access to the local cloud environment from remote locations, and to get all known benefits of the public cloud (access anywhere, anytime, and on every device). The VPN is potentially combined with Multiprotocol Label Switching (MPLS). A VPN combined with MPLS is typically made up from

- a central computing environment
- remote office locations
- other remote employee locations, like their home office

Since network access for the cloud is similar to traditional infrastructure solutions, existing network setups can be re-used with minor adjustments. The differences can be addressed by existing network management teams. Network access should be shaped with strong engagement from your security officer, because the network plays an important role in your enterprise's security.

2.1.5.1 Key benefits of using VPN and MPLS for building a hybrid cloud

A cloud VPN and MPLS connection can help you to create a secure and fast connection between your private cloud and your remote users, external service providers and public or Community cloud providers. In order to achieve high security, you create your own virtual private cloud within the public environment and connect it to your private environment.

Multiprotocol Label Switching (MPLS) is a mechanism in high-performance telecommunications networks that directs data from one network node to the next based on short path labels rather than long network addresses, avoiding complex lookups in a routing table. The labels identify virtual links (paths) between distant nodes rather than endpoints. MPLS can encapsulate packets of various network protocols. MPLS supports a range of access technologies, for example DSL, ATM, Frame Relay, and so forth

https://en.wikipedia.org/wiki/Multiprotocol_Label_Switching, as viewed on 03-01-2020

Key benefits of combining VPN and MPLS for building a hybrid cloud
Remote secure connectivity; extends your LAN/WAN to a global scale.
Cheaper than using traditional rented network connections; it makes use of standard Internet connections through DSL, fiber or cable connections or fast mobile data connections.
More mobility for employees; improve productivity for employees working from their home.

2.1.5.2 Architectural considerations

For most VPN connections the principle of IP tunneling is used. You create a secure point-to-point connection, a tunnel, through which you transfer your data. Technically seen, tunneling is the process of placing a packet within another packet and sending it over a network. Three different protocols are needed for tunneling:

Protocol	Summary	Examples of this protocol
Carrier protocol	This is the protocol used by the network.	
Encapsulating protocol	This is a protocol used for 'wrapping' the original data.	IPSec, GRE, L2F, L2TP, PPTP
Passenger protocol	This is the original data packet.	IP, IPX, NetBEUI

Key (architectural) building blocks are:

- IP tunneling
- security (firewalls, Internet protocol Security Protocol (IPsec))
- encryption
- Authentication, Authorization, and Accounting servers (AAA)

Special attention should be paid to the cloud network and the network's security performance. Network workload may be distributed across several network paths within your existing environment. This network workload will be consolidated at a single *cloud access gate*. The cloud access gate may end up with performance bottlenecks and additional security risks. These security risks may, for example, be due to overloaded security devices.

2.1.6 Risks of connecting a local cloud network to the public Internet

The IT network security company SecPoint provides us with an interesting question concerning cloud Internet security:

> Are companies really willing to risk having all their information, data, privacy, and software handled in a virtual cloud—a place where they're most susceptible to hack attacks and cyber invasions?
>
> www.secpoint.com

Storing data in the cloud seems comparable to storing your precious goods in a bank vault. However, in the case of a bank you know where they are physically located, and in which locker your goods are stored. With cloud computing, it is difficult to know how well your data are protected against theft.

Data security is an important consideration when using cloud computing. With Internet-based cloud services, data can be stored in any place in the world without you knowing it. Since you put more responsibility on the cloud service provider, new questions must be answered. That is, if they can be answered at all.

- What happens to your data when you are not using it?
- Is data automatically put under lock and key, or are your databases still open to the outside world?

SecPoint formulates this as follows:

> Cloud providers have a responsibility in ensuring that there are no exploitable bugs in their SaaS collection by deploying penetration testing and vulnerability assessment procedures on each and every last program they have available. Packaged or outsourced application code must be sorted out, examined, and monitored from the inside out as well.

Another heavy responsibility lies with the customers of the cloud service providers. The customer must check if their provider has data security under control. For example, the customer should find out how *data protection and partitioning* are organized in the cloud environment. Several measures must be taken to keep data safe.

Some of the available options are

- having a wall between data from different clients
- zoning
- hidden storage
- role-based customer profiles and user profiles
    - this guarantees anonymity, you do not know who your neighbors are

2.2 The principles of managing cloud services

2.2.1 IT service management principles in a cloud environment

Outsourcing IT services to a cloud provider does not mean that a customer, whether they are a large corporation or a small to medium enterprise or business (SME/SMB), can sit back and leave everything to the cloud service provider. Even if infrastructure management, operational IT, and management of services are outsourced, the customer must still focus on business-IT alignment and IT governance.

In order to check the provider's performance and compliance proper IT service management (ITSM) must be set up, and preferably in an auditable way. Proven frameworks can be used, both by the provider and the customer. Examples are:

- service integration and management (SIAM™)[6]
- business information management
- application management
- IT service management standard
    - ISO/IEC 20000-1
- Information Technology Infrastructure Library (ITIL)

2.2.1.1 Cloud service provider

Just like traditional IT service providers, cloud service providers must

- provide quality
- be cost-effective
- maintain security
- ensure availability of services

The Service Provider must be in control of the complete supply chain consisting of

- data center operations
- network and Internet providers
- other providers like SaaS solution partners

[6] The SIAM™ Body of Knowledge is freely available from Scopism. If you are interested in this topic, we recommend the EXIN SIAM™ program: https://www.exin.com/business-service-management/exin-siam/.

A well-proven standard for IT service management is ISO/IEC 20000-1. The scope covers both management and improvement of ITSM. ISO/IEC 20000-1 has both internal and external audit mechanisms. In this standard there are four process areas:

1. IT and the Business
2. Designing for service
3. Control of IT services
4. Support of IT services

2.2.1.2 Cloud customer

Customers of cloud computing and cloud-based services should expect, and demand, at least the same levels of service as is provided by traditional IT service providers and internal IT organizations. A corporate or public sector customer of cloud services must make sure that proper governance mechanisms are in place. Main elements are

- a good service level management (SLM) process on the side of the provider
- proper audit standards and instruments

A provider is preferably ISO certified to ensure that a good internal audit system, covering the complete supply and demand chain, is in place. Further business-focused standards may also be of benefit to the customer's own IT service management organization.

Standard	Name or summary
COBIT	Guidance for executive management to govern IT within the enterprise
ISO/IEC 38500	Information technology – Governance of IT for the organization
ISO/IEC 33001	Information technology – Process assessment (also known as SPICE) – Concepts and terminology

Recent cloud experience stresses a new challenge: wasted cloud spending. This means that the customer is paying too much, mostly for unused cloud services. Wasted cloud spending may take up more than 30% of total cloud spending. Reasons are, for example, not shutting down unused workloads or failing to select lower-cost clouds or regions. To avoid wasted cloud spending, there is an extra need for governance, focused on cloud cost optimization. This should include cloud usage policies and their automation.

Important questions to ask the cloud service provider:

- How are audits performed?
- Where are servers located, and what or which legislation applies to the data?
    - both now and in the future
    - How can you put this down in the agreement with the provider?
- What continuity plans are in place for recovering data, infrastructure, and applications?
- What are the provisions when a service changes or ends?
    - service life cycle and end-of-life
- What are the provisions if we want to migrate to another provider?
    - contract life cycle and end of life

2.2.1.3 Service level management (SLM)

Different cloud deployment models have different service level requirements (SLRs). Public clouds should be as standardized as possible, for instance. Private clouds need to be fine-tuned to the specific needs of the business. Since hybrid clouds are a complete mix of public and private clouds, the SLRs will also be a mix.

Cloud type	SLRs	Example of services
Public	• highly standardized • available to everybody • not used for critical data	• e-mail • productivity applications • storage
Private	• adjustable SLA's to fit your own specific requirements and needs • Operational Level Agreements (OLAs) to cover all IT services and systems and the network in the cloud stack • adherence to the service levels	• data storage • proprietary applications
Hybrid	• mix of SLRs for public and private	• mix of services

2.2.2 Management of service levels in a cloud environment

As a customer you will want the provider to comply with several quality specifications, and the corresponding processes must be in place.

The ISO/IEC 20000-1 standard contains several processes which are important for a cloud data center or provider. As a customer you will want to see proven strategies and implementation of the most essential service management processes to guarantee compliance.

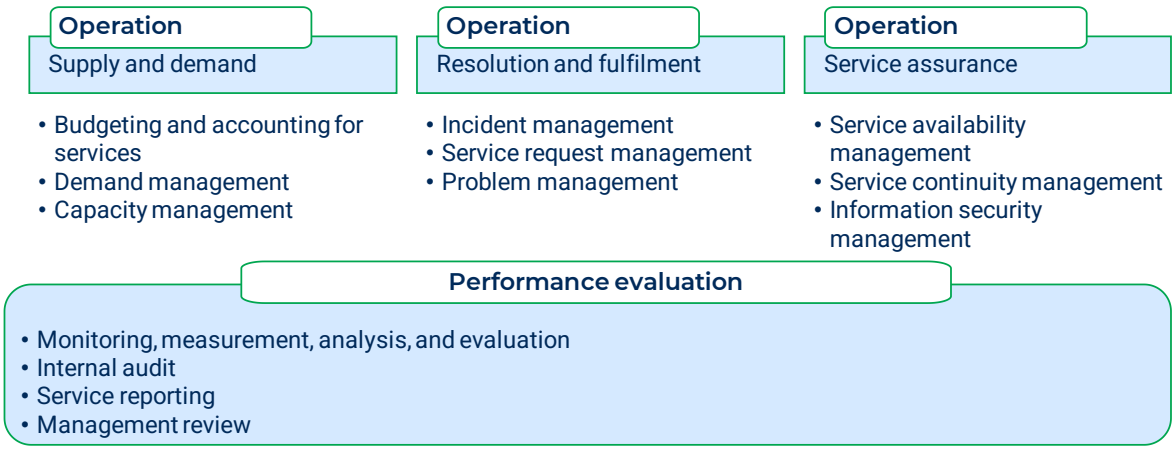

Figure 18 ISO/IEC 20000-1 processes relevant to customers of cloud service providers

2.2.2.1 ISO/IEC 20000-1 specification: processes

The goal of an ISO/IEC 20000 certification and follow-up audits is to make sure that the organization meets the ISO/IEC 20000 requirements. A number of requirements are mandatory.

- documentation of the ISO/IEC 20000 processes
- a demonstration that members of staff are familiar with these processes
- proof that the staff adheres to the procedures and working instructions

2.2.2.2 ISO/IEC 19086 (family) for management of service levels in a cloud environment

Cloud service management shares some basic principles with conventional IT service management (ITSM). Cloud management tools help providers administer the systems and applications that facilitate the on-demand service delivery model. The goal of cloud management is to improve the efficiency of the cloud environment and achieve a high level of customer satisfaction.

Given the elastic, highly virtualized nature of cloud environments, there are some key differences in approaches to cloud service management and conventional IT service management. The two disciplines have different objectives, requiring tools that emphasize their individual requirements. Automation is vital for cloud services to ensure efficiency and reduce costs.

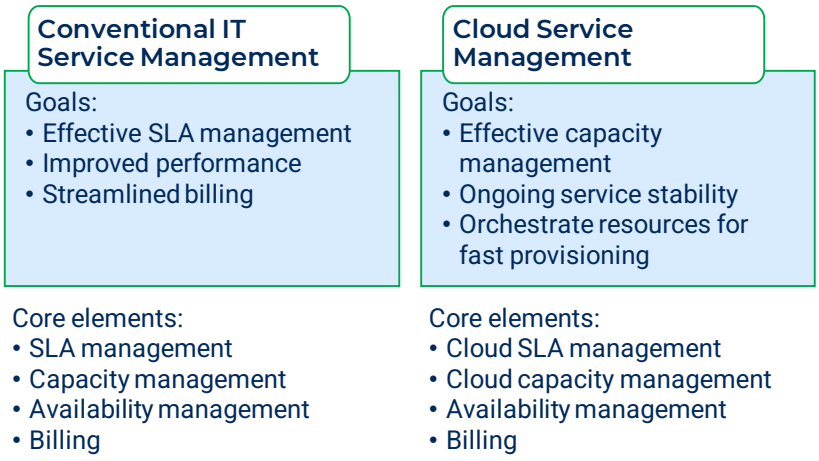

Figure 19 Objectives and core elements of conventional and cloud service management

The core elements of cloud service management mirror those of conventional ITSM and are applied to administer a cloud delivery environment in a systemic way. Currently, the cloud service industry lacks standardization in SLAs. The use of SLAs as a potential marketing vehicle have resulted in an SLA-jargon that is not always easy to understand. However, users are becoming more demanding in terms of service requirements, offered and guaranteed levels of quality, data protection, and so forth.

ISO/IEC 19086-family of standards	
ISO/IEC 19086-1	Information technology — cloud computing — Service level agreement (SLA) framework — Part 1: Overview and concepts
ISO/IEC 19086-2	Cloud computing — Service level agreement (SLA) framework — Part 2: Metric model
ISO/IEC 19086-3	Information technology — cloud computing — Service level agreement (SLA) framework — Part 3: Core conformance requirements
ISO/IEC 19086-4	Cloud computing — Service level agreement (SLA) framework — Part 4: Components of security and of protection of PII

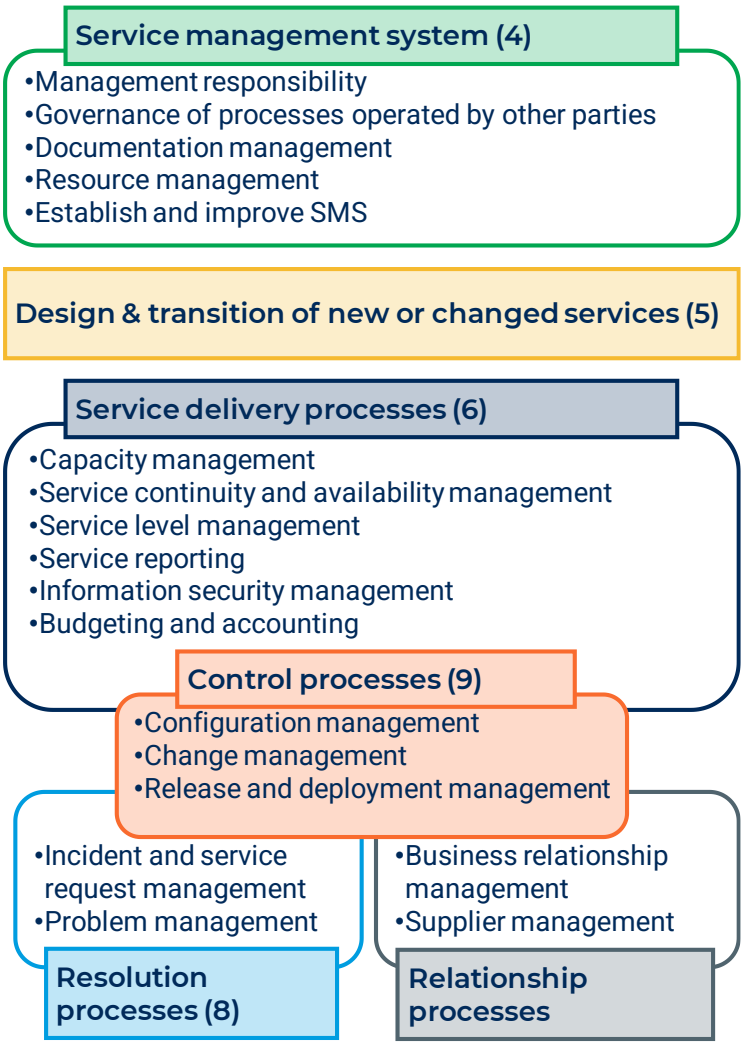

Figure 20 ISO/IEC 19086 (family) for management of service levels in a cloud environment

Exam preparation: chapter 2

'Get it' questions

1 / 4
Draw a local cloud environment with its main components and how they are interconnected.

2 / 4
What are the key benefits of using a VPN combined with MPLS?

3 / 4
There are risks of connecting a local cloud network to the public internet. One of the risks is that the burden of responsibility lies heavily on the customers. They need to check if their provider has everything under control.

For example, the customer will need to look how data protection and partitioning is organized in the cloud environment. Several measures must be taken to keep data safe.

What are the available options?

4 / 4
The goal of an ISO/IEC 20000 certification or a follow-up audit is to check if the organization fulfills the ISO/IEC 20000 requirements. Several requirements are mandatory.

One of these is the documentation of the ISO/IEC 20000 processes.

What are three process groups?

Include the processes per group.

Exam terms

- cloud provider
- hardware
- implementing cloud computing
- managing cloud computing
- local cloud environment
- MPLS
- risks
- software
- service management
- SLA
- VPN

Answers to 'get it' questions

1 / 4
Your drawing must show that the main hardware and main software components are part of the local cloud environment. For details, see section: Main components of a local cloud environment. For instance:

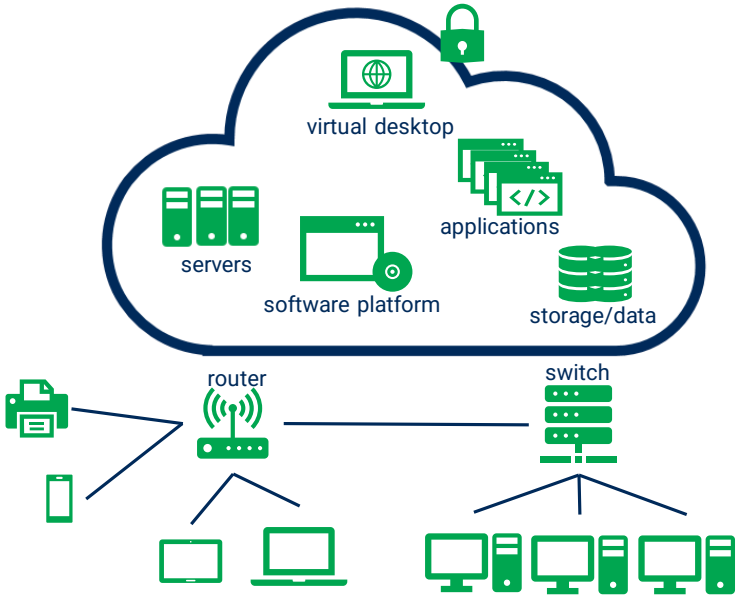

Figure 21 Local cloud environment

See also section Main components of a local cloud environment.

2 / 4
The key benefits of using a VPN combined with MPLS are:

- VPN secures remote connectivity. It extends your LAN/WAN environment over the Internet.
- VPN is cheaper than traditional rented network connections, because it uses standard Internet connections through DSL connections at home or fast cellphone network data connections.
- VPN gives employees more mobility without compromising security, which improves productivity for employees working from home and in the field.

See also section: Secured access.

3 / 4
The available options separate data from different clients, by using

- zoning
- hidden storage
- role-based customer and user profiles

This ensures anonymity, and guarantees your data are not accessed by unauthorized persons.

See also section: Risks of connecting a local cloud network to the public Internet.

4 / 4

All process groups are described in the table below. Your answer should include 3 groups and at least 1 process per group. See also section: ISO/IEC 20000-1 specification: processes.

.

Process group	Process
Service delivery processes	• Service level management • Service reporting • Service continuity and availability management • Budgeting and accounting for IT services • Capacity management • Information security management
Relationship processes	• Business relationship management • Supplier management
Control processes	• Configuration management • Change management
Resolution processes	• Incident management • Problem management
Release process	• Release management

3 Using the cloud

3.1 Accessing the cloud

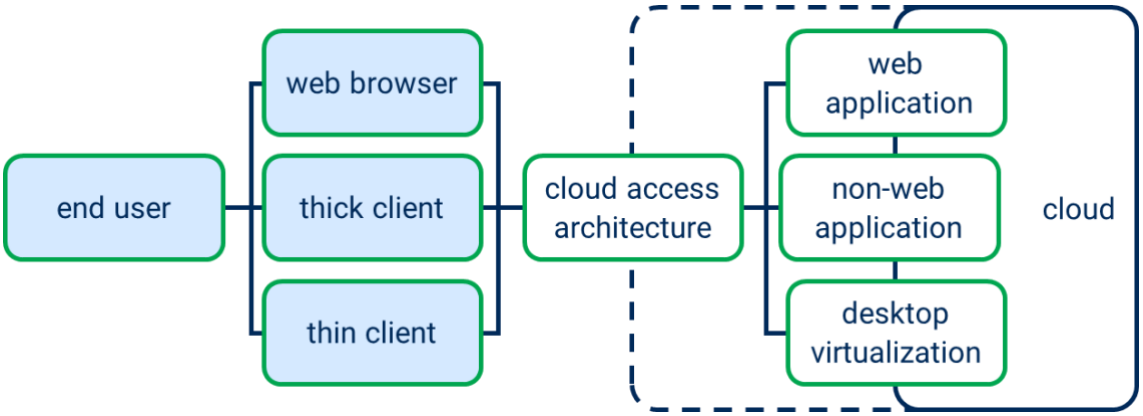

Figure 22 Accessing the cloud

3.1.1 How to access web applications through a web browser

Getting connected to the cloud to access cloud-based SaaS solutions, such as Microsoft Office, seems to be very simple. All you need is a PC or laptop, an Internet browser, an Internet connection, and a solution from a cloud service provider.

Pretend that you decided you want a free web-based productivity solution. As a private user you:

- buy a laptop
- connect it to the wireless router provided to you by your Internet service provider (ISP)
- choose a browser of your liking
  - for example, Google Chrome, Microsoft Internet Explorer, or Mozilla Firefox
- get onto the Internet
- find a provider
- subscribe for a free (but maybe limited) account

And that would be all. You can now use the web-based productivity solution. However, for businesses, small or big, it is a little more complicated. As a business, you will first have to determine:

- why you want to start using cloud solutions
  - price, flexibility, scalability/elasticity, access any time, any place from any device
- what you want to use
  - Office solution, file storage and sharing, online collaboration, professional applications like CRM, or a virtual desktop for your sales staff
- how you want to access these services
- what basic infrastructure suits your business best

In addition, you need to decide a few things:

- Do you want to have your own server farm and LAN, or do you want to get rid of all that equipment from your office?
- And what about privacy and security?

Accessing the cloud is still a straightforward and simple concept. All you need is access to the Internet! Or maybe, in case of private cloud plans, the intranet. To help understand how and why we need the Internet, it is good to have a quick look at a few basic concepts. The first one is the Internet. Let us start with a definition 'from the cloud', courtesy of Wikipedia.

The Internet is the global system of interconnected computer networks that uses the Internet protocol suite (TCP/IP) to communicate between networks and devices. It is a network of networks that consists of private, public, academic, business, and government networks of local to global scope, linked by a broad array of electronic, wireless, and optical networking technologies. The Internet carries a vast range of information resources and services, such as the inter-linked hypertext documents and applications of the World Wide Web (WWW), electronic mail, telephony, and file sharing.

http://en.wikipedia.org/wiki/Internet as viewed in May 2020

One of the key architectural components of the Internet is the standard Internet protocol suite (TCP/IP). Every device that wants to access the Internet requires an address to be identified and recognized; this is called the IP address. It is important to realize that the Internet is a global system of interconnected computer networks, meaning that you will be connected to the rest of the world. The computer networks are "linked by a broad array of electronic, wireless and optical networking technologies." In order to get these concepts into perspective we need to look at the Open Systems Interconnection model (OSI model), and the TCP/IP model.

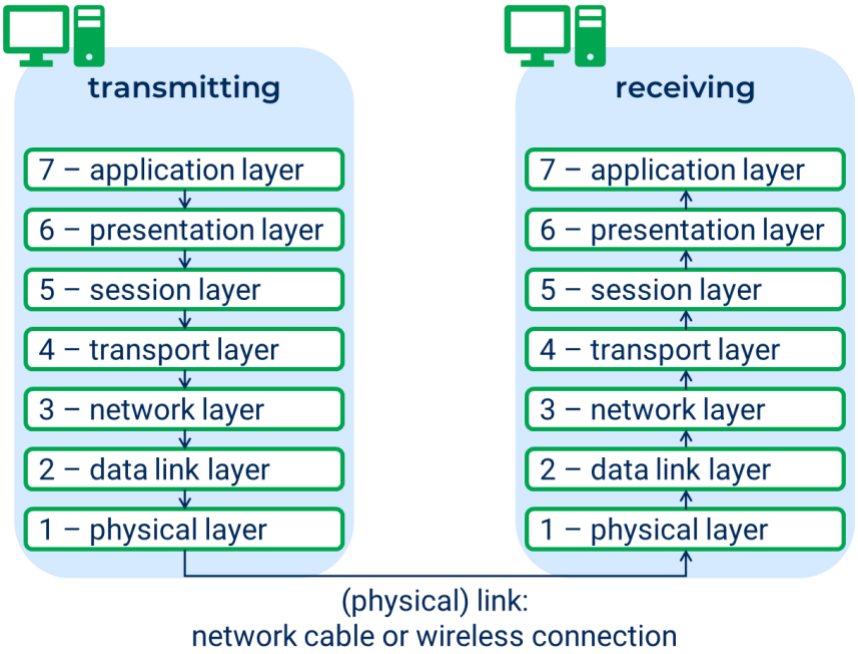

Figure 23 A simple representation of the OSI-model

(read more: https://en.wikipedia.org/wiki/OSI_model)

3.1.1.1 Open Systems Interconnection model (OSI model)

The OSI model (see Figure 23 A simple representation of the OSI-model) is a product of the Open Systems Interconnection project at the International Organization for Standardization (ISO), and is maintained by two organization:

- The International Organization for Standardization (ISO) – ISO/IEC 7498-1
- ITU Telecommunication Standardization Sector – ITI-T X.20

The OSI model was developed to help standardize the functions of communications systems in terms of layers. Cloud computing, because of its characteristics of sharing resources and interconnectivity, needs a high degree of standardization of its components, both on the provider and on the customer side. We can map standards to each layer of the OSI model to see what is needed in the chain between transmitting and receiving data.

Starting on the **physical layer (1)**, you will need some kind of physical or wireless connection to the Internet. This can be done in a few ways:

- with a wired connection accessing a LAN connection, which requires:
  - a network interface card (NIC), for which there are many so-called IEEE standards,
  - an Ethernet cable (RJ45-cat. 4 or 5)
- with a wireless connection
  - at home, in the office, or via public hotspots
- with a mobile Internet connection (3G, 4G, or 5G)
- through a satellite connection

Now you have established an Ethernet connection for the **data link layer (2)**.

The Internet protocol (IP) address provides the **network layer (3)**.

Next you need a transport medium for the **transport layer (4)**. This is the Transmission Control Protocol (TCP).

This covers the hardware side of things. Next, we need to establish a session with an application. The **session layer (5)** takes care of opening, closing, and managing applications. The NetBIOS protocol and sockets facilitate this layer.

In the **presentation layer (6)** we take care of presenting the data, for instance by translating EBCDIC[7] into ASCII. We use our web-based or cloud applications through the presentation layer:

- access a www website (http)
- make secure payments with a safe connection (https)
- receive e-mail through the post office protocol (POP)
- send mail with the simple mail transfer protocol (SMTP)
- transfer files with the file transfer protocol (FTP)

This leaves the **application layer (7)** where the application actually runs. It is defined as the user interface where the transferred data are displayed to the user.

In advanced cloud computing, layers 1 through 6 are all connected by unified protocols enabling interoperability.

[7] extended binary-coded decimal interchange code

From a business perspective, we need to establish our business requirements concerning

- speed and capacity
  - Do we need a simple cable connection or a secure rented line?
  - Do we want to use a copper cable or dark fiber?
  - Do we use multiprotocol label switching (MPLS)?
- mobility and mobile security
  - Do we want to allow the use of free unsecured hotspots?
  - Do we need personal connections through cellphone data connections?
- security
  - Do we set up a VPN?
  - What encryption do we need?

In the next section we will have a closer look at the architectural building blocks necessary to access the cloud.

The following figure shows how the ISO/OSI and TCP/IP models compare and differ.

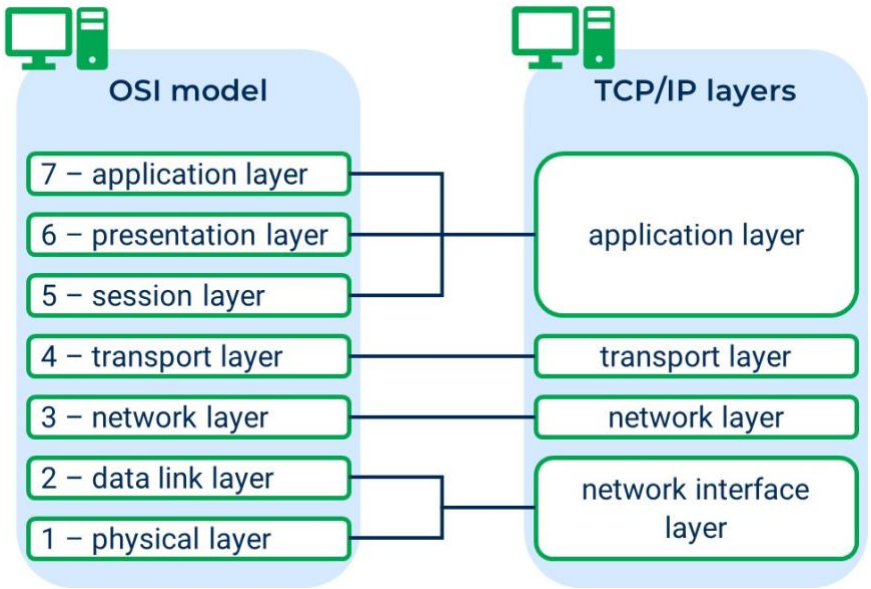

Figure 24 Comparison between OSI model and TCP/IP layers

3.2 Cloud access architecture

One of the key architectural features of cloud computing is the use of standard protocols and other design standards. The following table relates the OSI-model layers we discussed in the previous paragraph to some of these standards and what they are used for in the chain between transmission of data over the Internet, or your enterprise network, to receiving them on your network-enabled device (PC, thin client, smart phone, tablet, and others.)

OSI model layer	Protocol or standard	Description
7 – application layer	HTTP or HTTPS	• Hypertext transfer protocol (HTTP) • Hypertext transfer protocol secure (HTTPS) These protocols are the foundation for the world wide web (www).

6 – presentation layer	VT, RTSE	• Virtual terminal (VT)
		  ○ On the SASE sub layer (specific application service element)
		• Reliable Transfer Service Element (RTSE)
		  ○ On the CASE Sub layer (common application service element)

5 – session layer	API-sockets or sockets	• Application programming interface (API) An API allows application programs to control and use network sockets.
4 – transport layer	TCP, SSL	• Transmission control protocol (TCP) • Secure sockets layer (SSL) SSL is a protocol for encrypting information over the Internet.
3 – network layer	IP	• Internet protocol (IP)
2 – data link layer	Ethernet, IEEE 802.3	• Carrier Sense Multiple Access/Collision Detection (CSMA/CD) Defines the physical layer and data link layer's media access control (MAC) or wired Ethernet. (Wi-Fi equivalent: 802.11a,b/g,n)
1 – physical layer	100BASE-T, 1000BASE-T	Fixed line (100, 1000 megabits per second cable, LAN, WAN) or wireless connection

3.2.1.1 Security services

Cloud access architecture is not only concerned with the data transmission itself, but also with the security services which protect and influence the way you access the cloud. The cloud services' user experience depends on both data transmission and security services. The settings can limit cloud data transmission performance or even block your data traffic completely. Examples of cloud security services can be found both on the client side and on the internet service provider (ISP) side.

3.2.1.2 Virtualization

One or more layers of the OSI model, including network security services, can be virtualized[8].

3.2.2 Using a thin client and desktop virtualization

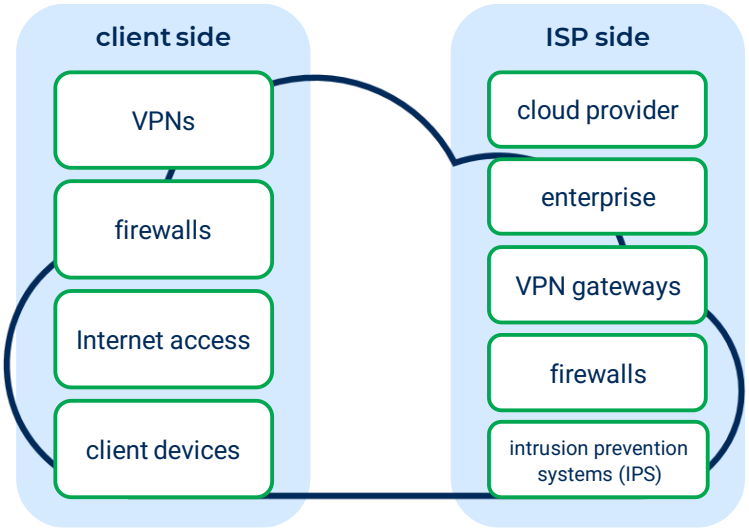

Figure 25 Cloud security services on client and internet service provider (ISP) side

Initially, central computer facilities and functions were accessed from a simple device called a *terminal*. However, this was still far from being connected to the whole world via the Internet. The connection was point-to-point between the terminal and the mainframe.

As early as the 1960s a form of virtualization was used to divide the powerful mainframe processor or processors into several separate machines, each performing their own task. The next milestone was ever more powerful personal computers and client server systems. Even with more computing power, a part of the actual application, the so-called client application, still runs on the local device.

With the emergence of cloud computing, most of the necessary computing power now runs on the systems in the cloud which is comparable to how the computing power used to run on the mainframes. This enables us to replace the expensive powerful PCs from our offices with a thin client, a kind of modern terminal, or other devices, such as mobile devices. Using a thin client with desktop virtualization is a very popular solution for businesses and enterprises of all sizes: it combines graphically powerful thin client terminals with cloud computing resources.

Traditionally, a thin client is a bare bones computer that allows users to access programs, files, and functionality that is hosted on another computer. This is typically a physical server located somewhere else in the office. The server pushes the operating system, programs, and information to the thin client when a user logs in. This is called *desktop virtualization*. The thin client is little more than a computer terminal that acts as a vessel to access information and functionality. The desktop virtualization service, with which your thin client interacts, can be obtained from many providers as a part of their cloud services. This is a Desktop as a Service (DaaS) offering. It is best to use the same cloud provider as for your major data and application hosting, because they use the same datacenter location. This will ensure that your virtualized desktop will have data access and application access through a high-speed, cloud datacenter network. Therefore, it will be able to achieve the highest performance for your data processing and applications interactions.

[8] You can do more research on this topic using https://en.wikipedia.org/wiki/Network_virtualization or https://virtualizationreview.com/articles/2014/10/14/7-layer-virtualization-model.aspx.

On top of being able to use a thin client, desktop virtualization allows access to your data and applications in different ways at the same time. You can use a web browser or software client for accessing a virtualized desktop on your PC, laptop, or even on a mobile device.

Desktop virtualization provides great flexibility. For example, you can use an enterprise thin client at the office and your own tablet or device at home to work on the same enterprise desktop anytime, anywhere, and in a secure way. This works for any kind of application, whether they are web-based, not web-based, or even legacy applications.

For businesses seeking enhanced security, higher user satisfaction, cost savings, and improved accessibility for remote users, combining thin clients with cloud computing offers an effective solution. To summarize the benefits:

- greater security and enhanced resource management
- higher end-user satisfaction
- more cost effective and green
- greater remote accessibility

3.2.2.1 Greater security and enhanced resource management

Thin clients offer network administrators peace of mind by eliminating the chance for physical data loss and lessening concerns about client integrity. The data in your cloud (datacenter) does not actually leave the secure datacenter perimeter. It will be presented at the virtual desktop, for example located inside of the same datacenter, and then it will be transmitted to your thin client in the form of a set of desktop screens.

3.2.2.2 Higher end-user satisfaction

Bringing a thin client into a cloud computing environment gives end-users all the benefits of a desktop PC without the headaches and challenges that typically accompany “fat” computers.

3.2.2.3 More cost effective and green

Because thin clients have few or no moving parts, there are few physical items that can break, which means a longer lifespan for every terminal. Since these machines let the cloud do most of the computational work, such as saving documents and sharing files, their power consumption is very low, which usually leads to lower costs and reduced energy needs.

3.2.2.4 Greater remote accessibility

Desktop Virtualization, through combining cloud computing with thin clients, allows for much greater flexibility in remote access than more traditional computing environments. Users can access their files, programs, and e-mail anywhere, from any device. This increases productivity without increasing information technology challenges.

3.2.3 Using mobile devices in accessing the cloud

The world can no longer be imagined without smartphones. Since the emergence of mass cellphone usage, the ‘phone’ seems to have come of age. We say ‘phone’, because the smart devices are hardly ever used for actual calling anymore.

Cellphone services using servers first appeared when we started to send short text messages (short message service, SMS). This was quickly followed by the transmission of pictures and short videos (multimedia message service, MMS).

Now, receiving e-mail on a phone is a standard function, using SMTP and Exchange protocols. The hardware platforms and operating systems for smartphones and other mobile devices have transformed them into web-enabled personal devices. Currently, the leading mobile operating systems are Google’s Android and Apple’s IOS.

Mobile phones are no longer the only smart devices. There are 5G enabled tablets, folded 7" screen smartphones, external monitors, smart TVs, and smart devices (traditional and innovative peripherals which connect to the Internet). In addition, there is easy access to desktop virtualization connecting devices to full-blown enterprise applications, graphic processing units (GPUs), and operating systems like Windows or Linux.

Microsoft is testing Windows 10 running on ARM processors (used in smartphones, tablets, and laptops), but a real compelling solution that competes with the desktop operating systems is still a way off. Another future development has been triggered by global politics. Because of limitations in accessing Android, Huawei was forced to start development of their own operating system for mobile devices. Whatever will come next, the future will tell.

Although there is a high degree of interoperability between the different cellphone networks this is not true for cellphone content like apps and services. These apps for one specific platform cannot be used on an operating system they were not created for. To increase their market share, software providers often develop apps for all leading mobile operating systems.

Developing apps for the cloud would mean universal accessibility. An example that already exists is using Exchange mail in combination with the e-mail push function on a smartphone. The user can update their calendar, contact list, notes, and task list on their phone, and receive as well as send e-mail. The changes are automatically updated on the Exchange server. When the user logs in to the Exchange account on a desktop application all changes immediately synchronize from the server to the desktop. When logging in to the web-based version, all changes are already synchronized.

For business users, there are some downsides to these new technologies. The new technologies bring security and privacy issues, such as business data that can be easily transported on removable media like SD cards or transferred over the Internet to inappropriate places. These new challenges require solutions to deal with the security and privacy issues. Possible solutions are device data encryption, and desktop virtualization. Working outside of the office may bring the additional problem of *shoulder surfing*, where unauthorized people literally look over an employee's shoulder at their screen. This may compromise sensitive data and may be tackled by a user-awareness campaign.

Business managers should consider the way smart devices are used by their staff. Many employees will have a corporate smart device and are used to being constantly online and connected to social media, music services, and streaming media. Most smartphone users have so many apps installed, that use of them may distract them from work. It may be worth looking into restricting the installation rights of users. In addition, mobile data usage outside of the agreed mobile plan could become a high cost for your organization.

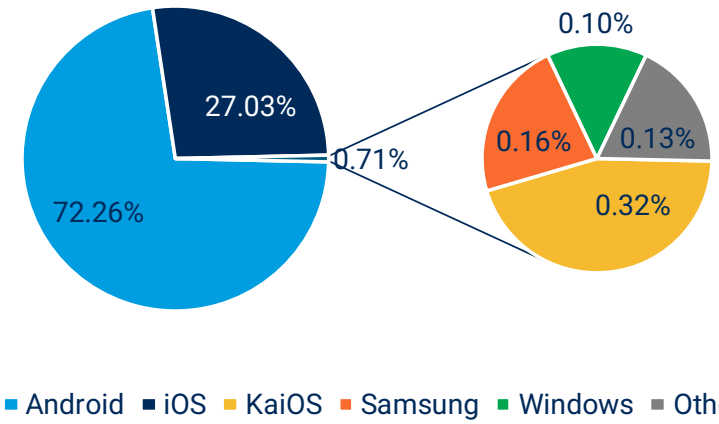

Figure 26 Mobile operating systems market shares worldwide as per March 2020 (source: https://gs.statcounter.com/os-market-share/mobile/worldwide)

However, smart devices may also lead to gains in productivity and flexibility, since employees carry their phones around with them all the time. This benefit may outweigh the aforementioned downsides, especially with establishment of the modern bring-your-own-device (BYOD) concept. More and more cloud services are compatible with BYOD, which will surely increase the use of the concept, since employers usually do not pay for the devices.

3.3 How cloud computing can support business processes

3.3.1 Identify the impact of cloud computing on the primary processes of an organization

First, we need to determine which cloud service model we *need* to support our business. Depending on this choice we can determine how to integrate our own infrastructure, present or future, with both the Internet and our cloud provider's infrastructure.

3.3.1.1 IaaS

If we only need a virtual desktop for our sales staff in the field, we can go for the Infrastructure as a Service model (IaaS). Typical IaaS services cover the first four layers of the OSI model.

3.3.1.2 PaaS

Or maybe we need many types of software development platforms with the elasticity between low capacity during the design and programming stages, and high capacity during the testing stages. In that case we can opt for the Platform as a Service (PaaS) model. Since we are programming our own application programming interfaces (APIs), we only need OSI coverage up to layer five, and maybe part of layer 6.

3.3.1.3 SaaS

For most customers the Software as a Service (SaaS) model is the most logical choice. In the SaaS model the customers arrange their own Internet connection, set up intranet or a local network if necessary, and supply devices that can connect to the networks. The cloud service provider takes care of the rest. Given all the connection options available, this means business can continue anytime, anywhere, and from any device.

Whether SaaS is a good solution depends on the scale of your business requirements and IT capability needs. Not all types of workloads and capabilities are covered by SaaS. Not all market offerings may fit your requirements. There may be conflicts between your requirements for mobility, and the requirements for security and privacy, or a conflict between your data and application capacity, and your integration needs.

When multiple SaaS services are necessary, this usually means using different SaaS vendors with different datacenters. This may cause challenges to connections and integrations. Using multiple vendors will affect the costs and requires you to think about requirements like speed, capacity, and elasticity.

Special attention should also be paid to new green-field cloud options offering new product lines based on new IT capabilities and emerging technologies. These new offerings should be evaluated against your time-to-market requirements, investment requirements (CAPEX, skills/capabilities, infrastructure), and risks. Often, really new technologies or solutions are high-risk. They may introduce failure when combined with current in-house setups.

Cloud computing is a great solution when introducing new technology. Cloud computing requires fewer large and risky upfront investments, and cloud handles unknown upfront capacity needs easily. In-house deployment speeds may be low, although a short time-to-market is necessary to reach the desired business impact. There may be limited to no possibilities for integration of new technologies with existing enterprise data or application landscapes, which causes a number of

unmitigated risks. Market-proven commercial products may not carry these risks or may be able to mitigate them faster.

There are many typical business cloud solutions available. The table below gives an overview of SaaS business solution examples with possible target audiences.

Running cloud solutions will result in additional benefits to only running solutions on your own data center. You can combine solutions from your own private (cloud) environment with public solutions, to create your own tailor-made hybrid cloud.

For example, in purchasing and manufacturing, cloud can enable collaboration with suppliers.

Using exchange and sharing platforms for sales, advertising, and marketing, you can promote and enable interaction with potential customers and the market through social media platforms. Many starting artists in the musical industry, but also established ones, use this mechanism as their primary sales channel.

Cloud CRM systems make it easy for field staff to update leads and contacts.

Category	Target audience	Examples
Customer relationship management (CRM)	Large businesses	• Salesforce.com • SugarCRM • NetSuite • Zendesk
Identity management		• Okta • Centrify • Onelogin
Enterprise resource planning (ERP)	Large businesses	• NetSuite • Compiere (open source) • Microsoft Dynamics 365
Manufacturing execution systems (MES)	Large businesses and small manufacturing businesses	• SAP Digital Manufacturing Cloud
HR solutions	All businesses	• Oracle Taleo • FinancialForce • Human Capital Management (HCM) • Natural HR • Workday • Namely • Synergy
IT service management	Large businesses	• ServiceNow • Splunk • Microsoft Flow • Snow
Finance and accounting	Large businesses	• Intacct • NetSuite • Zuora
Web design and management	All businesses	• Joomla • Adobe Creative cloud • XSitePro
E-mail services	All businesses	• Microsoft Office 365 • Gmail for business • Cisco WebEx Mail • IBM SmartCloud
Messaging	All businesses	• Slack • Microsoft Teams, Skype for Business • Google Hangouts • HipChat
Office suites	All businesses	• Microsoft Office 365 • Google Apps for Work • Zoho Office Suite
E-business	All businesses	• Capgemini Immediate • Oracle RIO • Gogiro • Google OpenEntry.com
E-commerce		• Lightspeed • Magento 2 • Shopify
Online storage and file sharing	Small businesses Medium businesses	• Google Drive • Dropbox • Amazon S3 • Microsoft SharePoint 365 • CloudShare
Video conferencing	Small businesses Medium businesses	• Skype for Business • Microsoft Teams • WebEx Meeting Center • GoToMeeting • Zoom
Code		• Bitbucket • GitHub

3.3.2 The role of standard applications in collaboration

Social media like Facebook, LinkedIn, and Twitter have led the way to modern forms of collaboration. Sharing and working together on documents is now easy because of services like Google Drive and Dropbox. Video conferencing and free Internet calls are possible because of the race that was started by Skype, and things have moved on quickly from there.

New free Cloud services emerge in rapid succession, and the more successful of these are further developed into professional Cloud services. For example, Google's Gmail now has a business equivalent which can be combined with other Google Apps for business. As a result, Google has become a serious business contender. However, the big two (at this moment in time) are still Microsoft Azure with Office 365 for business and Amazon Web Services (AWS) - Cloud Computing Services that include Hybrid cloud solutions.

3.4 How service providers can use the cloud

3.4.1 How cloud computing changes the relation between vendors and customers

With more IT services being moved to the cloud, IT providers will have to rethink their service models. The relationship with their customers will change. Because providers take over a large part of the customer's IT systems and service management, they must be ready to rewrite their SLAs to demonstrate to their customers that they can and will deliver most of the value chain.

Since service providers are now essentially running their customer's business processes, customer privacy must be considered. This calls for a proper and transparent audit trail. Especially large and corporate customers will demand proof of compliance to standards like ISO/IEC 20000 and COBIT.

Recent trends for multi-cloud strategies show that enterprise IT service providers are extending their offering and starting initiatives through collaboration with other providers. This is evolving to cloud service ecosystems that provide cloud services with a single customer relationships management function. One party will act as a service broker for external service providers. This essentially creates a one-stop-shop solution from the customer's perspective. Through service brokers, uniform hybrid cloud services are created with underlying hybrid cloud automation. There is a single point of service responsibility. Multi-vendor management leads to data and application integration across the value chain. This benefits end-users.

Cloud ecosystems have a relatively small presence, although there is a big growth in cloud market offerings and ample competition. Currently, customers must frequently change between different cloud vendors to fulfill the business requirements. This creates a demand for easy cloud migrations, which happen as part of business as usual. Easy cloud migrations require suitable services, automation, and clear contractual commitments from vendors. This should lead to quick, smooth, and supportive migrations for an affordable price.

More and more products and services outside of traditional IT are complemented by cloud services or connected to the Internet of Things (IoT).

- Cameras can be connected to cloud video surveillance services with a one-click ordering or subscription option.
- Cars, like Tesla, are cloud-connected, which places different demands on dealerships and garages, enables maintenance updates to be pushed via the cloud, and allows anti-theft and navigation services to work.
- Robot vacuum cleaners are connected to the cloud to ensure optimal algorithms to clean your home.
- Smart lights are connected to the cloud to ensure the perfect lighting for any moment of the day and automated lighting when your smartphone comes within reach of your front door.
- Air conditioners are giving users feedback on temperature and air quality by connecting to a smartphone app.

Almost everyone uses cloud services. There are few products left that cannot be made 'smart'. As a consequence, customers deal with any number of different cloud vendors. Is this cloud heaven or a cloud nightmare?

In addition, almost any business will, sooner or later, become a cloud service vendor to support their own products or services. They could set this up as a part of their business or cooperate with native cloud vendors.

The increase in use of many different cloud service vendors may be the biggest driver for cloud ecosystems we have seen so far. How this will ultimately affect relationships between vendors and customers can only be guessed at.

3.4.2 Benefits and challenges of providing cloud-based services

Service providers must ask themselves if their business is future proof without cloud offerings. The first question any provider must ask themselves is:

- Can the business thrive in the current or future market without cloud services in the portfolio?

Even if a business is **not** a service provider or if the core business seems to have little to do with IT, cloud services must be considered. Questions any business should ask, include:

- Will my product survive without native cloud services attached? Especially if competitors offer that service.
- Will my product survive without cloud services attaching them to the Internet of Things (IoT)?

Even if the business survives without cloud services, the business may not thrive without. In any case, it is a good idea to evaluate the benefits of adopting cloud services.

3.4.2.1 Benefits

With the development of the cloud business model, new business opportunities, as well as challenge, for IT service providers. Traditional data center providers are quickly rebuilding their traditional infrastructure services into IaaS, PaaS, and SaaS services. Due to the multi-tenancy principle, this makes their existing resources regain business value. Multiple users on a single platform adds the economy of scale principle[9]. Software developers developing on PaaS platforms decrease the time in which their applications can be designed, built, and tested.

An even greater benefit may be the new flexible solutions from the cloud, which support rapidly changing enterprise workplaces and the recent smart workspace trend. Smart workspaces may include a more flexible office space, innovative information flows, 3D printers, business analysis automation, and face recognition technologies.

With their long experience in service management and IT, leading hosting providers have an opportunity in the cloud market. Hosting providers are uniquely positioned to capitalize on the managed services market, because not every business is ready for a self-service model and not many can afford the investments for emerging services research and development.

Cloud benefits may explode across many areas that, at first glance, have nothing to do with cloud. For example, the cloud-enabled sales cycle is a simple online click and buy transaction. But consumers spend increasingly more time selling and buying online. With low costs of sales and few personal and social sales skills required, anyone can become a salesman. Cloud offerings support sales platforms like Amazon and eBay, that support both business-to-consumer and consumer-to-consumer transactions. They also offer independent web shops:

- easy set-up of their digital stores (any e-commerce platform)
- an integration of checkout and payment possibilities (Ingenico, Stripe, PayPal)
- different shipping options (Sendcloud)

3.4.2.2 Challenges

Of course, cloud-based services also bring a number of challenges. Not every cloud service provider is ready to deal with:

- adherence to standards
- complex security issues
- high performance demands in combination with flexibility (elasticity)
- privacy and data protection legislation
- availability and continuity issues

Cloud expectations are high: the cloud is a 24/7 business!

Scaling the cloud up or down in a private or on-premise cloud environment may still offer challenges. The following questions must be addressed:

- Is a pay-per-use service with affordable pricing possible when assets are located, dedicated to, or owned by the customer?
- Is it possible to provision extra capacity on demand in a timely manner, but without a costly oversized hardware reserve?
- Can freed-up assets be re-used in an efficient and financially viable way?

[9] Read more about this principle here: https://en.wikipedia.org/wiki/Economies_of_scale.

It has never been easier to do price comparisons for cloud services between vendors. Most services look almost identical and seem to have similar pricings. Vendors offering *differentiated services* sell unique services or products. However, most cloud service providers sell a *commodity service*: something that is not unique at all, such as SaaS or PaaS platforms. Vendors will have to determine:

- what their unique selling points (USPs) are
- what value they can add to your business
- how to justify different pricing from other vendors
- what their strategy is to future-proof their business

Exam preparation: chapter 3

'Get it' questions

1 / 6
List the 7 OSI-model layers and describe the hardware, software, protocols and standards necessary to access web applications through a web browser.

2 / 6
What are the **2** key architectural features of cloud computing?

3 / 6
What is a thin client? What is a benefit of using thin clients over more traditional solutions in the office?

4 / 6
There are many business solutions available from the cloud. Name at least 7 SaaS business solution categories (such as e-mail services) and give an example of a service for each.

5 / 6
Using cloud computing changes the relationship between vendors and customers. What aspects change?

6 / 6
What are **two** benefits and **two** challenges of providing cloud-based services?

Exam terms

- accessing the cloud
- OSI model
- cloud web access architecture
- thin client
- mobile devices
- cloud computing in the business
- service providers

Answers to 'get it' questions

1 / 6

OSI model layer	Protocol or standard
7 – application layer	HTTP or HTTPS
6 – presentation layer	VT, RTSE
5 – session layer	API-sockets or sockets
4 – transport layer	TCP, SSL
3 – network layer	IP
2 – data link layer	Ethernet, IEEE 802.3
1 – physical layer	100BASE-T, 1000BASE-T

The hardware needed is:

- a device capable of accessing the Internet: desktop, laptop, mobile device
- router and switch

The software needed is:

- a web browser
- a web application

See also section: Accessing the cloud.

2 / 6

1. The use of standard protocols and other design standards
2. Virtualization

See also section: Cloud access architecture.

3 / 6

A thin client is a simple network-enabled computer without a hard disk (it boots from the network) or any other moving parts (no DVD drive).

Benefits are (you should have listed one of these):

- Costs for thin clients are lower than for fully capable devices.
  - lower initial price and running costs
- Simple devices without moving parts do not break often.
- Thin clients are better for the environment, because they produce less heat and need less cooling.
- Thin clients provide heightened security, because they
  - boot from the network with controlled access
  - do not store data locally
- User error is almost impossible with thin clients, because they are very simple.

4 / 6

Please use the table in the section SaaS to check your answer.

5 / 6

The following aspects of the relationship between provider and customers change:

- Providers are more intimately involved with the customer's business.
    - Providers are running a critical part of the customer's business.
- Providers are pressed to show their value to the business as part of the business' value chain.

See also section How cloud computing changes the relation between vendors and customers.

6 / 6

Benefits	Challenges
Reutilization of old resources (IaaS)	Compliance risks
Better use of resources because of multi-tenancy	Standards, legislation and regulations
Economics of scale	Performance issues
Quickly develop and run applications in the same environment (PaaS)	Availability, capacity, flexibility, scalability
	Security risks
	Privacy and data protection risks

See also section Benefits and challenges of providing cloud-based services.

4 Cloud security, identity, and privacy

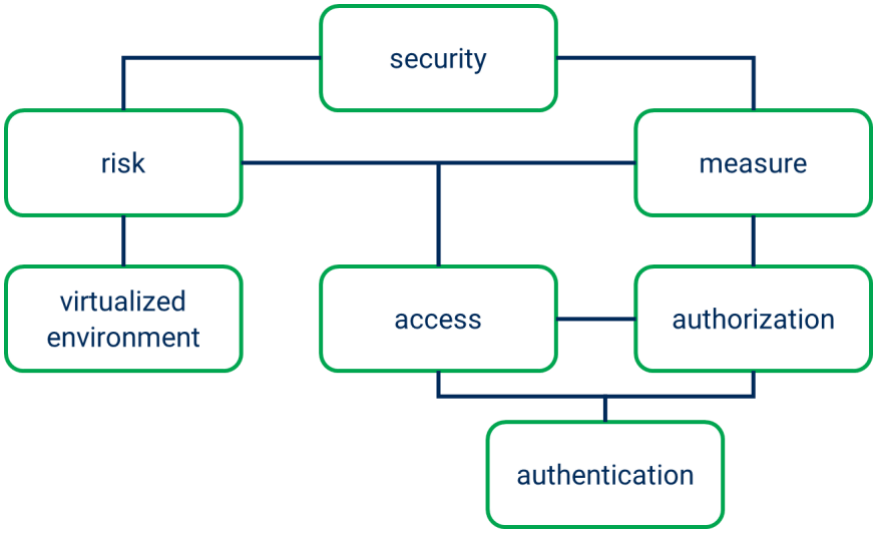

Figure 27 Relationship between security, risk and mitigations

Public cloud computing means

- sharing resources
- using the Internet
- multi-tenancy
- a mix of free and non-free services
- storing data anywhere in the world
- dealing with anonymous customers
- unclear SLAs
- many technical standards

In light of these elements, security was always high on most providers' agendas. The 2018 introduction of the European General Data Protection Regulation (GDPR)[10] poses an extra incentive to pay attention to privacy and data protection. Huge penalties have already been imposed for lack of proper handling of personal data. Private cloud computing solutions bypass some of these challenges, for instance challenges with date storage locations, but still need to address the other security, privacy, and data protection challenges.

By realizing that there are security risks, customers can assess prospective providers and choose services that will not compromise their own compliance with legislation and regulations. Cloud security is still growing. However, modern enterprise cloud developments and the introduction of the GDPR have caused cloud security to mature quickly.

Many experts claim that cloud inevitably brings security risks. However, it is not always clear which "cloud" they are talking about. In many cases, the claimed risk is only relevant for public clouds and not for private clouds. By definition, private clouds must address security risks most urgently. After all, private clouds are most likely to contain sensitive and personal data.

In unmanaged private clouds, security measures are limited to in-house measures. This is rarely enough to manage the risks properly. To manage cloud security risks properly, first a proper security assessment must be done.

[10] See also the section: European Union and European Economic Area

The steps for a security assessment are:

- define which cloud deployment model (public, private, hybrid) must be assessed
- define this deployment model and their specific characteristics within your business
- determine which part of your deployment model is actually a public cloud
- determine which part of your deployment model is a type of private cloud
    - for example, virtual private cloud
- determine which setup is used for your private cloud, because this will seriously impact the risk level
    - check operations, datacenter options, and so forth

You should perform separate, independent assessments for all different deployment models used. When you identify a hybrid cloud, also focus on the areas where the different models connect. This will help you find a proper balance for the hybrid setup.

Cloud services provided by a mature, well-established enterprise IT service provider may prove to be more secure than your in-house cloud deployment or traditional IT setup can be. The following questions may identify when that is the case by assuming your business is a cloud service provider. They also help with a fast validation or comparison of cloud providers.

- Does the provider's security incident tracking match up to benchmarks?
    - You may base this on the provider's market reputation and available benchmarks.
- Has the provider been conducting independent security compliance audits?
- Which information security and physical security measures are in place for the datacenter?
- Which staff policies and screening practices are in place to guarantee security?
    - For example, are new employees screened during recruitment?
- Has the provider's staff had appropriate security training?
- Do both you and the provider have a proper cloud security officer?
- Do both you and the provider have a security incident response procedure and a security incident response team?
- Do both you and the provider have dedicated security artifacts in operations and for all infrastructure?
    - separate management network
    - special infrastructure access solutions for administrators
    - role-based access control (RBAC)
    - automated user and user rights provisioning for administrators
    - password management automation
    - operating systems hardening practices
- What are your and your provider's security SLA practices?
    - policies updating speed
    - patching level targets
    - updating speed and automation
- What is the level of unintended IT service use and which policies address the risks?
    - unauthorized access
    - public services used in operations
    - non-authorized communication tools

4.1 Security risks of cloud computing and mitigating measures

Cloud computing has many new architectural and design features. These features create different types of security risks. A number of leading IT providers (HPE, Oracle, Qualys, Microsoft, and Rackspace) have joined forces with customers such as the Bank of America, to form the Cloud Security Alliance (CSA)[11].
The Cloud Security Alliance (CSA) is a not-for-profit organization with a mission to

- promote the use of best practices for providing cloud security
- provide education on the uses of cloud computing to help secure all other forms of computing

4.1.1 Cloud security alliance top threats and their risk mitigation measures

The top threats to cloud computing according to the CSA are[12], ranked in order of significance:

1. Data Breaches
2. Misconfiguration and inadequate change control
3. Lack of cloud security architecture and strategy
4. Insufficient identity, credential, access and key management
5. Account hijacking
6. Insider threat
7. Insecure interfaces and APIs
8. Weak control plane
9. Metastructure and applistructure failures
10. Limited cloud usage visibility
11. Abuse and nefarious use of cloud services

4.1.1.1 Data breaches

A data breach is any incident where data is seen or manipulated by an unauthorized individual. This includes both business data, and personal data or personally identifiable information (PII). Storing and accessing data in the cloud has many advantages but can also be compromised in many ways. It can be altered or deleted without a backup; it may be unlinked from its context, or accessed by unauthorized people.

Risk mitigation
Data breaches can be prevented by addressing all other top threats. Data security should be a priority for any business. Setting up solid security, both for data-in-rest and for data-in-transit, is important. Access management is key to protecting data.

4.1.1.2 Misconfiguration and inadequate change control

Data breaches often occur because of misconfiguration. Misconfiguration can happen when cloud computing is not set up for secure access to data, permissions are excessive, and default settings are left unchanged. Change control in cloud computing sometimes lacks the necessary approvals for changes, due to automation.

Risk mitigation
Solid change control mechanisms and adapting application configurations to the businesses' needs increase data security.

[11] Find out more on their website: https://cloudsecurityalliance.org/.
[12] This section is based on the CSA's *Egregious Eleven* threat list of August 2019.

4.1.1.3 Lack of cloud security architecture and strategy

Businesses can simply transfer all their data to a (public) cloud solution, without first deepening their knowledge of the risks involved. Often, the knowledge necessary to put a solid cloud security architecture and strategy in place is missing in the business. This leads to unmanaged security risks.

Risk mitigation

It is wise to invest in training in cloud security architecture to provide ample knowledge on the business side. Devising a strategy before migrating data may reduce the risks.

4.1.1.4 Insufficient identity, credential, access and key management

We all deal with many passwords and ways to verify our identity and credentials. Because this is bothersome, people tend to re-use passwords and avoid two-factor authentication.

Risk mitigation

Businesses should ensure that their employees rotate passwords often, use two-factor authentication, and choose sufficiently strong passwords to manage this threat.

4.1.1.5 Account hijacking

Most e-mail users are aware of fraudulent tactics like phishing, password hacking, and identity theft. Passwords that grant access to cloud services are kept outside your own company's IT domain and can be compromised. For businesses this can mean they are vulnerable to industrial espionage or can lose important business data or processes.

Risk mitigation

Examples of mitigating measures that can be taken are strong authentication techniques and monitoring of user behavior.

4.1.1.6 Insider threat

If cloud providers are a cross-section of our society, statistically seen, some employees or sub-contractor staff may be untrustworthy. However, it is not just *malicious* insiders that pose a threat. Negligence by employees and contractors is much more often the cause of data breaches. Examples include leaving laptops unlocked in public spaces, storing sensitive information on a regular USB stick, and clicking on spoofing links in e-mails.

Risk mitigation

Two examples of mitigating measures that can be taken are good HR vetting procedures and strong policies and procedures regarding information security. Additionally, awareness campaigns and security training will reduce the threat of negligence.

4.1.1.7 Insecure interfaces and APIs

Application interfaces (APIs) are key components for most cloud services. They connect different services to each other. If these interfaces are not properly designed for security, they can pose a security risk. User interfaces (UIs) are the interaction mechanism for the end-users. Since both APIs and UIs are in the public space, they get attacked often.

Risk mitigation

Examples of mitigating measures are:

- designing for security
- proper testing methods
- understanding how APIs and UIs interact with other interfaces and software
- strong authentication and access control

4.1.1.8 Weak control plane

Data in cloud environments must be protected against data breaches. The combined architecture and security measures that form this protection is called the *control plane*. In a solid control plane, the cloud architects have full control over security measures and insight in all data flows.

Risk mitigation

Choose a cloud service provider (CSP) that provides ample security controls to their customers, so that you can adapt the controls to match legislation requirements and your business' needs. Research the security controls provided before deciding on a cloud service provider.

4.1.1.9 Metastructure and applistructure failures

The applistructure is the infrastructure of the cloud-deployed applications. The customer takes care of cloud management in the metastructure. If management is not done well, metastructure and applistructure failures may lead to insecure APIs as well as poor identity and key management.

Risk mitigation

Find a cloud service provider (CSP) that provides transparency of their cloud services, and shows penetration testing results to their customers.

4.1.1.10 Limited cloud usage visibility

This security issue means that the organization has trouble to determine whether cloud applications are safely used. It has two parts:

- unsanctioned app use
- sanctioned app misuse

Unsanctioned app use means that employees are using cloud applications that were not sanctioned and supported by the company's IT provider. This may lead to employees using insecure cloud solutions, which in turn may cause data breaches. The business may not see this happening at all.

Sanctioned app misuse happens when employees use approved applications, but do not adhere to company policies that prevent unauthorized access. In addition, approved applications may suffer from a cyber-attack resulting in data loss. Since actual credentials are used, it is difficult to analyze this problem.

Risk mitigation

Risk mitigation strategies include:

- training employees to reduce negligence and increase awareness
- making non-approved apps difficult or impossible to install
- using cloud access security brokers (CASBs) or software-defined gateways (SDGs) to monitor cloud use
- investing in web application firewalls (WAFs)

4.1.1.11 Abuse and nefarious use of cloud services

Many cloud providers have easy access to their services, which are sometimes even free for a trial period. Registration is relatively anonymous which can and will attract malicious actors, such as spammers and hackers. In addition, cloud providers may (unknowingly) not only host your data and applications, but also malware (malicious software). Cloud computing resources are easily leveraged to back DDoS attacks as well as massive spam and phishing campaigns.

Risk mitigation

Examples of mitigation measures that can be taken are the validation of credentials and increased monitoring of traffic between customers and known suspicious sites.

4.1.2 Other cloud security threats

4.1.2.1 Data loss

Data in the cloud has many advantages but can be compromised in many ways. It can be altered or deleted without a backup, it may be unlinked from its context, or accessed by unauthorized people.

Risk mitigation

Examples of mitigation measures that can be taken are authentication, audit (in other words ISO/IEC 27001, Data Security) and authorization, as well as use of encryption and the application of a proper backup strategy.

4.1.2.2 Denial-of-service

Denial-of-service attacks are attacks that aim to prevent users of a cloud service from being able to access their data or their applications. They force the target cloud service to consume inordinate amounts of finite system resources such as processor power, memory, disk space or network bandwidth. In doing so, the attacker (or attackers, as is the case in Distributed Denial-of-service (DDoS) attacks) causes an intolerable system slowdown and leaves all of the legitimate service users confused and angry as to why the service isn't responding.

Risk mitigation

Mitigation measures consist of:

- developing a Denial of Service (DoS) response plan
- combining firewalls, VPN, anti-spam, content filtering and load balancing
- fixing identity management
- keeping some redundant server power
- responding quickly to signs of a DDoS attack, such as suddenly slowing networks, connectivity, or websites

4.1.2.3 Insufficient due diligence

Moving into the cloud may make it more difficult for organizations to prove their compliance to legislation and regulations during external audits, such as an EDP or IT-audit.

Risk mitigation

Examples of mitigation measures are assessing the financial health of the CSP or determining the length of time the cloud service provider has been in business

4.1.2.4 Shared technology vulnerabilities

A multi-tenant architecture has its own challenges. Some components may not have been developed for this type of use and may cause security issues.

Risk mitigation

Examples of mitigation measures that can be taken are:

- enhanced operations procedures for monitoring or escalations when security breaches occur
- application of good security practice for installation, configuration, and application of patches

4.1.3 Measures for mitigating security risks

An objective way to ensure a provider's compliance with security best practices is to demand ISO certification.

- The ISO/IEC 20000-1 standard contains a security paragraph.
- The ISO/IEC 27001 standard may be even more important.
    - Full name is: *Information security, cybersecurity and privacy protection — Information security management systems — Requirements*

- The ISO/IEC 27002 standard describes best practices.
    - Full name is: *Information security, cybersecurity and privacy protection — Information security controls*
    - It is important to check how the provider has implemented these best practices.

4.2 Managing identity in the cloud

4.2.1 Main aspects of identity management

The effectiveness of authentication processes in non-cloud situations depends on whether you are connected to a domain or not.

If you are not connected to a domain, authentication is **local**:

- Your username and password are validated against account information stored on your machine.
- Only if the files you want to access are further protected, will you be asked for further credentials.

If you are connected to a domain, domain verification (usually **active directory authentication),** is performed:

- You log on with your active directory (AD) account.
- The Kerberos protocol authenticates your credentials without transmitting a password in either clear or scrambled form.
    - This ensures that the password cannot be cracked by a so-called password engine.
- You gain access to the machine, even if your local account is not known on your machine.

Kerberos is a network authentication protocol. It is designed to provide strong authentication for client/server applications by using secret key cryptography.

A free implementation of this protocol is available from the Massachusetts Institute of Technology. (Source: http://web.mit.edu/kerberos/#what_is)

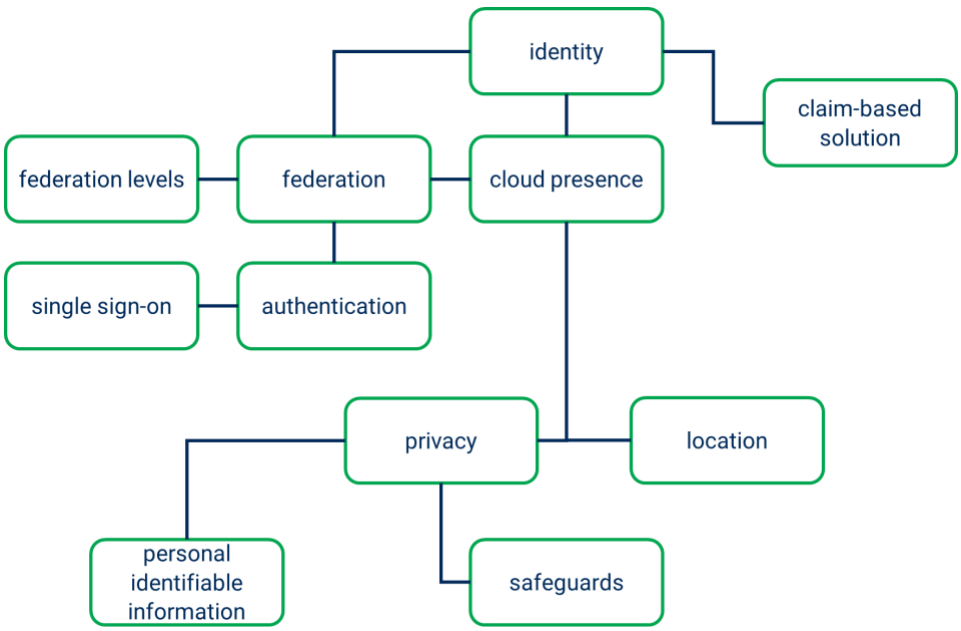

Figure 28 Identity management in cloud environments

4.2.1.1 Authentication in the cloud

Cloud models add a layer of virtualization. The hardware that supplies your computing power is most likely no longer yours. Instead, your computing power will come from a provider like Amazon AWS or Microsoft Azure.

When using a *private cloud model*, authentication is performed by a domain controller or authentication servers running on the virtual machines. Businesses need to protect their data and resources running on the virtualization servers from unauthorized access, for instance through a VPN. *Public cloud models* pose more serious authentication problems.

- It is uncertain if the public cloud solution you are working with uses identity management at all.
- If identity management is in place, there are different ways of verifying identity:
  - user-ID and password lookup in a database
  - using the Lightweight Directory Access Protocol (LDAP)
  - Kerberos verification
  - two-factor authentication (2FA) with a pass code sent over SMS or by using an authentication app.
- A single sign-on system is probably not feasible.
- There is no standardization of Internet-based security yet.

Since hybrid cloud models have both private parts and public parts, authentication must be set up separately for each part.

Regardless of the cloud model used, **two-factor identification** will always provide more security.

4.2.1.2 Authentication, Authorization, and Accounting

Authentication, Authorization, and Accounting (AAA, pronounced *triple A*) is a functionality to manage electronic access. The three concepts are the security cornerstones of IP-based network management and policy administration.

Concept	Related question	Summary
Authentication	Who or what are you?	Authenticates identity of a user or device, for instance by using a password, user ID, or security token.
Authorization	What are you allowed to do?	Determines whether an entity is authorized to take the requested action, for example access to restricted data, or a log-in outside office hours.
Accounting	What have you done?	Tracks resource usage by users and devices, for example as part of an audit trail, costing or billing, or capacity monitoring.

4.2.1.3 Identity and access management

Whether you work on the company network, the intranet or are using cloud solutions, organizations need proper identity control and access governance. Identity and access management is the management of:

- authentication
  - determining who is accessing
- authorization
  - determining what they are allowed to do

Identity must be managed both within and across system and enterprise boundaries. The goal of identity management is to increase security and productivity, while decreasing cost, downtime and repetitive tasks.

One way to think about identity management is by imagining an enormous blueprint of an office building. It shows the rooms into which each person who works in the building can enter. The blueprint also shows what kind of key each person would need to open the door to get into that room, and what that person can do once they are there.

From: Aitoro, Jill R. (2008); The Basics: A Glossary of Federal Technology – Identity Management, www.nextgov.com

In her article, Aitoro continues the metaphor and says:

A computer network is like the building, and each room represents a file, database or application on that network.

Typical characteristics of an identity management system are:

Characteristic	Summary
Role management	IT implementation of a business role
Role hierarchy	a representation of an organization chart
Separation of duties	sometimes called 'segregation of duties', this means that more than one person must be involved to complete a task
Group management	permissions are not given to people but to roles
Self-service functions	a place where users can perform simple tasks, like resetting a password, without help-desk intervention
Password synchronization	ensures that the same password can be used on multiple devices
Digital identity	presence and location determine available services and capabilities
Federation Identity Management	enables single sign-on

There are several identity management solutions available, for example, Microsoft Forefront, Azure Active Directory, and IBM Security Identity and Access Manager. We will look at two examples of identity management solutions. These solutions are single sign-on (SSO) and a special type of SSO called federation identity management.

4.2.1.4 Single sign-on (SSO) for web services

A cloud-based security infrastructure is a distributed infrastructure: security features and algorithms are spread over a certain domain. This makes it necessary for users to log in to a number of different applications. This is very inefficient and unproductive.

A solution for this problem is the single sign-on (SSO) principle. All distributed security elements are consolidated on a single SSO server. Users only have to log in once using a security measure like a smart card, a security token, or an active directory (AD) account. The SSO server ensures secure access to all services.

SSO architecture uses the SOAP protocol[13], a protocol for the exchange of information in the implementation of web services in the cloud or any other network.

4.2.1.5 Federation identity management

Federation identity management, or the "federation" of identity, describes the technologies, standards, and use cases, which serve to enable the portability of identity information across otherwise autonomous security domains.

The goal of identity federation is to enable users of one domain to securely access data or systems of another domain seamlessly, and without the need for redundant user administration.

[13] You can find more information here: https://simple.wikipedia.org/wiki/SOAP_(protocol).

Many different varieties of identity federation exist, including

- user-controlled or user-centric scenarios
- enterprise-controlled or Business-to-Business (B2B) scenarios

Federation is enabled through the use of open industry standards and openly published specifications. The openness enables multiple parties to work together and ensure interoperability.

Typical use cases involve

- cross-domain single sign on
- web-based single sign-on
- cross-domain user account provisioning
- cross-domain entitlement management
- cross-domain user attribute exchange

4.2.1.6 The future of identity management with blockchain

Blockchain technology is likely to revolutionize identity management. Blockchain technology is very resistant to hacking. This resistance reduces the threats of insufficient identity, credential, access/key management, and account hijacking.

Blockchain has no central repository of data. All data is stored in a decentralized and self-verifying way. This makes blockchain tamper proof. Tamper proof identity management and verification sound very appealing and are most likely in our future.

In addition, blockchain technology allows natural persons to federate their own identity. Right now, we are relying on trusted companies to deal with identity management for us. That means that our personal data lives on many different servers that are not in our control. Blockchain technology[14] enables self-federated identity which secures our privacy.

4.3 Privacy and data protection in cloud computing

4.3.1 The concept of privacy

The first definition of the term *privacy* concerning law can be found in an article written by Samuel Warren and Louis Brandeis (an attorney and a lawyer, respectively, both American) in 1890. In the USA, it is regarded as the first time that 'a right to privacy' was advocated.

"The right to be left alone" - Warren & Brandeis

Source: The right to Privacy, Harvard Law Review, 15 December 1890. The first definition of privacy in legal terms.

A more current definition comes closer to the modern concept of privacy.

Someone's right to keep their personal matters and relationships secret.

Source: Cambridge Dictionaries online

The overarching document for any 'privacy' legislation in the world is the Universal Declaration of Human Rights (1948, UDHR). The United Nations organization describes it as follows:

[14] If you are interested in learning more about blockchain technology, we recommend the EXIN Blockchain program: https://www.exin.com/qualification-program/exin-blockchain.

The Universal Declaration of Human Rights (UDHR) is a milestone document in the history of human rights. Drafted by representatives with different legal and cultural backgrounds from all regions of the world, the Declaration was proclaimed by the United Nations General Assembly in Paris on 10 December 1948 (General Assembly resolution 217 A) as a common standard of achievements for all peoples and all nations. It sets out, for the first time, fundamental human rights to be universally protected and it has been translated into over 500 languages.

Source: Universal Declaration of Human Rights https://www.un.org/en/universal-declaration-human-rights/

In the UDHR there are two articles with reference to the concept of privacy.

Article 12:

- No one shall be subjected to arbitrary interference with his privacy, family, home or correspondence, nor to attacks upon his honor and reputation. Everyone has the right to the protection of the law against such interference or attacks.

Article 8:

- Everyone has the right to respect for his private and family life, his home and his correspondence.

The following picture shows the hierarchical relationships of UDHR with the most important EU documents that refer to the protection of personal data.

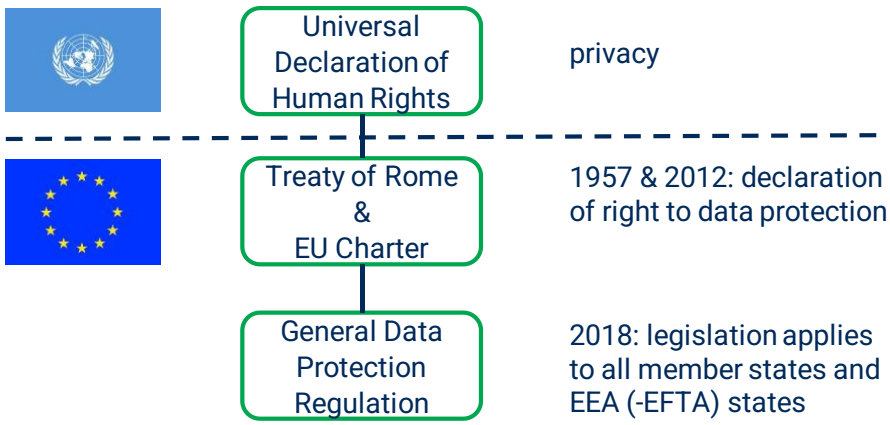

Figure 29 The concepts of 'privacy' and 'data protection' in the EU context

At the highest EU legislative level, protection of personal data is described in the Treaty on the Functioning of Europe (1957) and the Charter of Fundamental Rights of the European Union (2012, Article 8: Protection of personal data). The General Data Protection Regulation (2016), which came into effect in 2018, underpins the declaration of the right to data protection by calling it a 'fundamental right'.

4.3.2 High-impact legislation (EU and USA)

4.3.2.1 European Union and European Economic Area

GDPR

The most important piece of European legislation for privacy and data protection is the General Data Protection Regulation 2016/679 (GDPR)[15]. This is a regulation by which the European Commission intends to strengthen and unify data protection for individuals within the European Economic Area (EEA).

The regulation was adopted on 27 April 2016 and came into effect on 25 May 2018. The GDPR applies to all countries of the European economic area which includes

- all European Union (EU)countries
- non-EU countries
  - Iceland
  - Liechtenstein
  - Norway

The word privacy is not mentioned in the law text of the GDPR at all. It only occurs in the GDPR as a footnote reference. The GDPR focuses on the *protection* of all personal data: data that have the potential to identify a natural person. Good data protection practices ensure the privacy of these natural persons. National legislations must reflect the practices in the GDPR.

Examples of aspects of the GDPR with impact on cloud computing are:

- Key definitions; (direct, indirect, sensitive, pseudonymized, anonymized) personal data, data subject, controller, processor, and so forth
- Data subject rights
- Processing of personal data (includes all types of storage)
- Lawfulness of processing (includes storage)
- Data breaches (includes security procedures)
- Organization of data protection (legal requirements, compliance)
- Data protection by design and by default
- Data location

Full text of the GDPR can be found on https://eur-lex.europa.eu/eli/reg/2016/679/oj

The GDPR only allows processing of data in third countries, which includes transfer of data between the EEA and countries outside the EEA, under an adequacy decision. Currently, the adequacy decision regarding data transferred under the US Privacy Shield is revoked[16]. It is imperative that at the moment of processing an adequacy decision is in place.

[15] You can find the full text of the GDPR in the language that is most familiar to you here: https://eur-lex.europa.eu/. EXIN's Privacy & Data Protection program is designed to inform you about the GDPR and its implications on a few different levels. You can find more information here: https://www.exin.com/qualification-program/exin-privacy-and-data-protection.

[16] You can find the full press release regarding the Privacy Shield on http://bit.ly/PDPF_privacy_shield.

ePrivacy Directive
The *Directive on privacy and electronic communications 2002/58/EC* is also called the ePrivacy Directive. This directive applies to data protection in the digital age. It covers the use of:

- cookies
- spam
- data traffic
- confidentiality of data

This directive deals with all matters that are not covered in the GDPR. In contrast to the GDPR, EU national legislation may differ from the practices in the ePrivacy Directive.

4.3.2.2 United States of America

CLOUD Act
The Clarifying Lawful Overseas Use of Data Act (CLOUD Act) was signed into law in March 2018.

> This bill amends the federal criminal code to specify that an electronic communication service (ECS) or remote computing service (RCS) provider must comply with existing requirements to preserve, backup, or disclose the contents of an electronic communication or noncontent records or information pertaining to a customer or subscriber, regardless of whether the communication or record is located within or outside the United States.

Other legislation
Other US privacy and data protection legislation is based on various sources, such as:

- the Constitution
  - First Amendment: right to free assembly
  - Fourth Amendment: right to be free of unwarranted search or seizure
  - Fourteenth Amendment: due process right
- the concept of 'invasion of privacy' (common law based on jurisprudence)
- the Children's Online Privacy Protection Act (COPPA)
- specific state privacy laws

4.3.3 Clash of the Acts

The last sentence of the general description of the US CLOUD Act is:

> …regardless of whether the communication or record is located within or outside the United States.

The above in conflict with the GDPR, when it applies to 'personal data' (as defined in the GDPR), or any personally identifiable information (PII) (as is the more common name in the USA).

According to the GDPR, a person's personal data is protected *regardless of the nationality of the data subject* if

- the data subject is located within the EEA
- the personal data is physically located within the EEA

According to the CLOUD Act, the US government reserves the right to claim personal data of any person stored by a US service provider (in other words most major cloud service providers), even if the servers are situated within the EEA.

We cannot solve this conflict here. Contact your legal staff or advisors if you are in doubt as to what applies when confronted by this issue

Exam preparation: chapter 4

'Get it' questions

1 / 5
What is the most common security risk to cloud computing?

2 / 5
Which measures can be taken to mitigate that security risk?

3 / 5
Explain what Authentication means within AAA.

4 / 5
Explain the SSO principle.

5 / 5
Is there a difference between personal data and PII?

Give **4** examples of personal data.

Exam terms

- security
- compliance
- mitigating measures
- federation
- identity management
- privacy
- authentication
- authorization
- accounting
- AAA / Triple A
- single sign-on (SSO)
- safeguards
- personal data / PII

Answers to 'get it' questions

1 / 5
The single most common security risk is a data breach, which may be caused by any of the other security risks mentioned in the section Cloud security alliance top threats and their risk mitigation measures.

2 / 5
Data breaches can be prevented by addressing all other top threats. Data security should be a priority for any business. Setting up solid security, both for data-in-rest and for data-in-transit, is important. Access management is key to protecting data.

Please refer to risk mitigation for all other top threats to the section Cloud security alliance top threats and their risk mitigation measures.

3 / 5
AAA, or Triple-A, stands for authentication, authorization, and accounting. These elements are the security cornerstones of IP based network management and policy administration. The accounting element answers the question: What have you done? It tracks how users use their resources and devices, for example to provide an audit trail.

See also section Authentication, Authorization, and Accounting.

4 / 5
Web-based security infrastructure is distributed. Security features and algorithms are spread all over a domain. A solution for this problem is offered by the single sign-on (SSO) principle. All distributed security elements are consolidated on one SSO-server. As a result, a user only must sign on once. SSO architecture uses the so-called SOAP protocol, a protocol for the exchange of information in the implementation of Web Services in the cloud or any other network.

See also section Single sign-on (SSO) for web services.

5 / 5
Both personally identifiable information (PII) and personal data are terms for any information that can be used to directly or indirectly uniquely identify, contact, or locate a natural person. The USA CLOUD Act uses the term PII. The European GDPR uses the term *personal data*.

Possible examples:
- Types of identification: SSN, passport, fingerprints, iris-scans
- Occupational: job title, company name
- Financial: bank numbers, credit records, PIN-number
- Health care: insurance, genetic
- Online activity: logins, IP-address, MAC address
- Demographic: ethnicity
- Contact: phone, e-mail address, social media accounts, street address

See also section High-impact legislation (EU and USA).

5 Evaluation of cloud computing

5.1 The business case for cloud computing

The most important reasons for moving to the cloud in favor of traditional IT infrastructure are:

- Cost savings
- Time to market
- Scalability and elasticity
- Productivity and mobility
- Asset utilization

Making this move means less capital expenditure (CAPEX) in IT infrastructure (including datacenter facilities and network services), and has immediate financial impact. Cloud services provide much more flexibility and better time to market (TTM) of business solutions. Getting a new business application up and running is faster in a cloud environment.

Another reason to move to cloud is that it results in an environmentally responsible business. Cloud computing is more sustainable than traditional IT. Non-profit organizations may find this a compelling reason.

These reasons for moving into the cloud can be evaluated through fundamental financial metrics used to make investment decisions:

- Total Cost of Ownership (TCO)
- Return on investment (RoI)

The TCO and RoI perspectives for the cloud business case are discussed in later paragraphs[17].

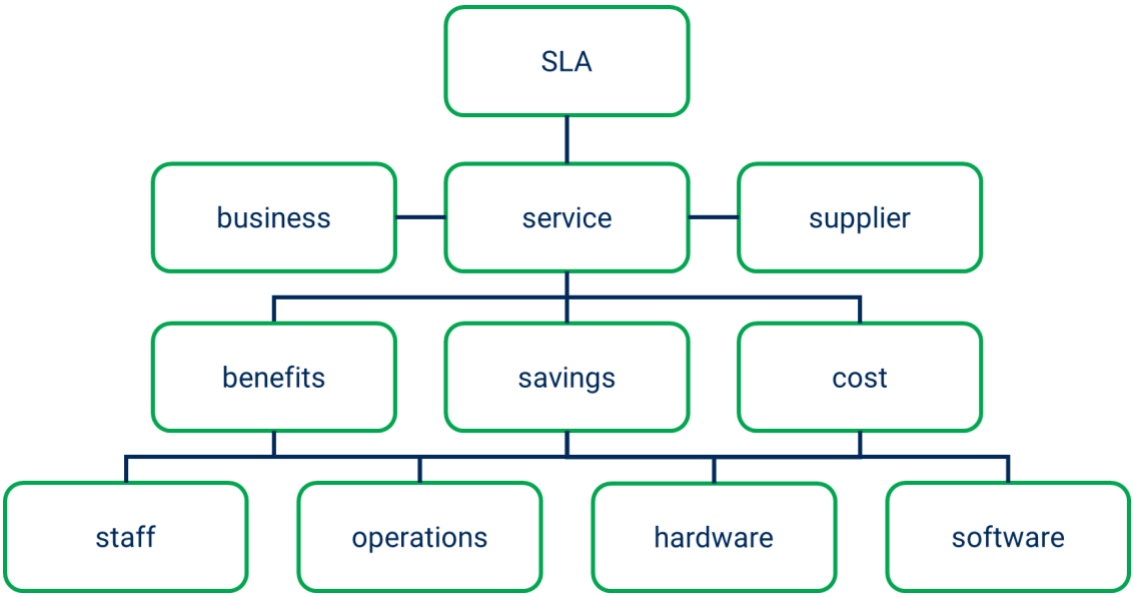

Figure 30 Relations between SLA and the services delivered

[17] An example of an online article that evaluates TCO and ROI can be found on https://technologyfinancepartners.com/2019/07/tco-vs-roi/.

All services delivered to the business are managed by a service level agreement (SLA). An SLA puts your requirements, expectations, and key performance indicators (KPIs) in writing, regarding:

- performance
- security
- scalability
- availability
- support
- compliance

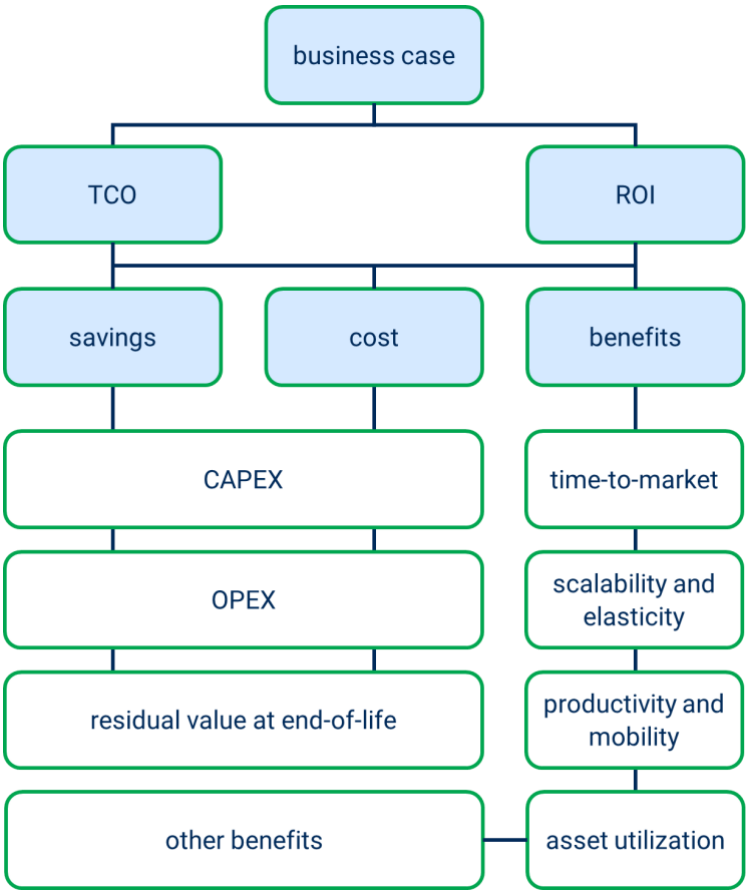

Figure 31 The business case for cloud computing

5.1.1 The total cost of ownership (TCO)

When moving to the cloud, we expect the total cost of ownership (TCO) of IT to be lower. However, this is not necessarily the case. According to Aggarwal & McCabe (2009) “cloud computing shifts the TCO discussion”. Individual components may have low costs, but when all costs are added together, TCO may actually be the same or higher than before moving to the cloud.

The TCO of cloud computing is very dependent on domain or enterprise. The benefit you get from using cloud-based platforms, both in public cloud models and in private cloud models, is related directly to the type of organization you are, and the business processes you support.
Other factors must be considered as well, such as

- existing staff skills
- existing investment in hardware, software, and facilities
- existing laws and regulatory obligations

Using TCO models that do not consider the type of organization and many other factors, are unlikely to result in a truly holistic view of the value of cloud computing for your business.

The main components in TCO calculations are

- hardware costs
- software costs
- support costs

For the finance department these will be further broken down into smaller components like:

- staff costs
- depreciation
- utility costs
- maintenance costs

Hurwitz et al. coined the term total cost of application ownership (TCAO) in their 2009 publication. They calculate the total annual cost of ownership including hardware support plus some amortization costs. The total cost of application ownership has the components:

- server costs
- storage costs
- network costs
- backup and archive costs
- disaster recovery costs
- data center infrastructure costs
- platform costs
- software maintenance costs
    - package software
    - in-house
- help desk support costs
- operational support personnel costs
- infrastructure software costs

There are some gray areas within cloud computing costs that are difficult to understand and prove, but may still need to be part of the TCO calculations. An example is capital costs or capital expenditures (CAPEX). Capital expenditures will be eliminated to a high extent by moving to the cloud. If you do not have any capital, it will not cost you anything.

However, subscription costs and pay-per-use costs may fully replace capital expenditures. In addition, when migrating your application servers and application management to a SaaS service, you may want to buy external support or consultancy. These are operational expenditures (OPEX) that replace former capital expenditures. What you gain in CAPEX savings, return in the form of costs you never had before.

5.1.2 The return on investment (RoI) perspective

Perhaps cost savings should not be the primary reason for moving into the cloud. The real benefit of cloud computing is not in cost savings. The greatest benefit is that IT can react much faster and more effectively to changes in the business. Businesses become more agile by moving IT to cloud services. For a startup company the cloud offers both low CAPEX in IT infrastructure and a much shorter implementation time for the infrastructure. Small Internet-based businesses can be up and running in a very short time.

When business agility is the most important reason to move to the cloud, TCO models do not work well to prove business value. This can be better presented from a return on investment (RoI) perspective. The RoI reflects the received value, set against the investment over a given period of time.[18]

Return on investment is calculated by the formula:

$$RoI = \frac{\text{Final value of the investment} - \text{Initial value of the investment}}{\text{Total cost of the investment}}$$

5.1.2.1 Time to market

Regardless of the cloud deployment model or cloud service model used, a key benefit is the speed of change that supports your IT-enabled product deployments. Emerging technologies, increased capacity for new products, or additional computing capacity can be implemented in a very short time. This will shorten the time to market for any of your new products or services.

A short time to market means short product deployment timelines. Cloud customers can evaluate the monetary value of this benefit in terms of:

- accelerated margin or market share
- fast worldwide coverage
- market leadership if you can get to the market early

5.1.2.2 Scalability and elasticity

Seasonal business peaks and promo-actions increase IT resource demands. When IT performance is not up to the task, this seriously affects customer experience. Customers expect websites to respond instantly. When they notice response times, it is usually only if they are too long. A slow website will result in customers abandoning the check-out process. IT capacity has a direct influence on business performance. Cloud characteristics like scalability and elasticity easily support fluctuating demands with real-time capacity scale-ups.

Businesses can evaluate the monetary value of scalability and elasticity by comparing the expected extra revenue that can be generated with an enhanced customer experience to the increased costs of the used cloud solution. Especially businesses with a non-cloud IT solution will benefit, because they do not need to keep expensive redundant resources just for peak performance times.

Even businesses without seasonal dynamics experience periods during which extra workload is required from the IT systems. Systems overloads cause expensive overtime. Businesses may even face penalties when performance failures create situations that lead to compliance failures.
The cloud's capacity for elasticity organically adapts the cloud solution to any business change. Cloud solutions easily support re-allocation of IT capacity from discontinued services to new solutions and new lines of business. Cloud natively scales IT capacity to correspond with business and workforce changes. These business changes are more frequently needed in the current market conditions. Cloud solutions scale almost in real-time and under the pay-per-use concept. In contrast to traditional IT solutions that have fixed one-year budgets, with cloud solutions you pay for the capacity you actually need at the time you need it.

Evaluating the benefits of scalability and elasticity will help build a business case for cloud models with inbuilt scalability.

[18] Hardware tends to be useful for around 5 years in a business. Therefore, 5 years is generally an acceptable time period for ROI and TCO calculations.

5.1.2.3 Modern productivity and mobility

Modern cloud services have out-of-the-box capabilities that can seriously improve your workforce outcomes by enabling them to use their 'core devices' like personal smartphones. These capabilities allow the employees to work anytime and anywhere, and support bring-your-own-device (BYOD) concepts.

Smartphones tend to be always on and within the employee's reach. This can seriously drive collaboration and business engagement. Increased communication drives innovation and gives a new flexibility to the working environment. In turn, this may improve the employee experience, which may have a positive impact on their wellbeing, and on business productivity and efficiency. Cloud solutions that match employees' experience in daily life can help retain talents from a generation that expects their experience at work to be great. Keeping high-performing employees onboard is important to all successful businesses.

To determine the value of cloud solutions that foster 21st century communication, you need an overview of:

- current workforce costs
- the (expected) revenue per employee of the solution
- the identified (expected) productivity gains
- the number of potential users of the cloud solution

You should also consider whether the proposed cloud solutions will help you generate new business in line with your targets.

5.1.2.4 Asset utilization

Usually, when the total cost of ownership (TCO) of cloud solutions are evaluated, cost savings are not as high as expected. However, that does not mean that there is no business case for cloud purely from a cost perspective. Since cloud elasticity usually works under a pay-per-use concept, this impacts asset utilization.

Since traditional IT struggles with scalability, uniformity, and long procurement cycles, many businesses have more capacity than needed. The overcapacity is used as a risk buffer for fluctuating capacity demands. Often less than 50% of the assets, both IT assets and human resources, are used on a regular basis. If IT budgets are restricted to traditional, yearly formats, future capacity demands may outpace them. It may also happen that IT budgets are simply not managed at all, because they are cumbersome to deal with.

With proper cloud governance, automated cloud usage policies, suitable IT budget management flexibility and financial management for IT services, cloud services can manage to increase asset utilization to 90% or more!

By avoiding direct costs associated with traditional IT, cloud benefits should make you seriously rethink your TCO-driven picture. Once you have reliable figures for your current asset utilization, you can set realistic targets for asset utilization in a cloud environment. To assess the financial gain this brings, the RoI of investing in a cloud solution can be computed. This may completely change the financial impact of IT over time.

As an example, suppose that your current asset utilization is 40%. Imagine that you have 100 servers running, but you are only using 40 of them. That is a 40% asset utilization. If you scale down to owning 50 servers, your utilization is 80%, you are still using 40 servers, but now out of the available 50. Cloud solutions allow you to reduce the number of servers you pay for from 100 to 50, without losing elasticity for sudden performance demands. Since it is based on a pay-per-use model, you are not spending your capital on 100 servers that must be written off over time, but instead on 50 servers that are replaced by the service provider when necessary. The bottom line is that we can reduce the TCO in this specific example by 50%.

5.1.2.5 Other benefits

Some other real benefits from cloud solutions lie with operations and staffing. Operational tasks such as implementation, maintenance and support will be performed by the cloud service provider. This allows the business to do with fewer staff and lowers training demands for your remaining staff.

Examples of operational benefits:

- managed services
- self-service (unmanaged services)
- instant server deployment
- software licensing without impact on CAPEX
- uptimes are guaranteed
- back-ups as a Service (always off-site)

Examples of staffing benefits:

- lower IT staff costs
- decreased need for hard-to-find IT experts
- lower recruitment, HR and training costs

The impact of these benefits depends on the type of business you are in and the scale you operate on. It also depends on the scale, dynamics, type, and current IT Infrastructure.

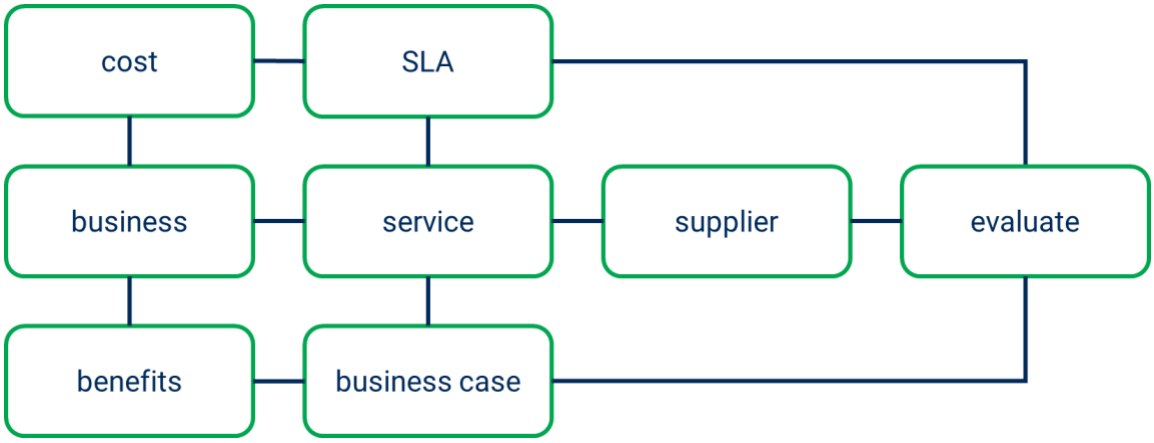

Figure 32 Relation between benefits and business case

5.2 Evaluation of cloud computing implementations

5.2.1 Evaluating performance factors, management requirements, and satisfaction factors

When outsourcing IT to cloud service providers, support and maintenance are no longer in-house and may take longer. Control over security measures will also be outsourced when using cloud services.

To ensure that a move to a cloud solution, or a change of cloud service providers is successful, the following questions should be evaluated:

- Can the cloud services support the business?
  - What are the costs, savings, and benefits of the cloud solution?
  - Is there a solid business case for moving to a cloud solution?
  - Is there a cloud solution that fulfills the requirements of the business?
  - Are multiple cloud services necessary?
- How robust is the security of the cloud solution?
  - How does cloud security level compare to your own?
  - Does the level of security measure up to what the business expects and needs?
- Does the cloud provider's SLA measure up to SLAs your business has?
  - Which expectations does the business have of the cloud provider's SLAs?
  - What are the service level targets you need the cloud provider to keep?
- What is the system performance of the cloud solution?
  - Is it higher than maintaining your own data center, private network, and systems?
  - What connection and transaction speeds do the business require?
- What are the cloud service billing mechanism and units?
  - Are there any special discounts and price altering options?
  - Are there volume discounts or capacity trade-in practices?

It makes sense to do a comparative study of several providers before you sign a contract.

5.2.2 Evaluating service providers and their services

During the process of choosing a cloud service provider, the business will have compared service providers and chosen services that fit the business needs. As a part of *due diligence*, the business must approve of security practices and review data protection measures taken by the cloud service providers. The service level agreements (SLAs) clearly show the performance you may expect. But the business cannot stop evaluating.

When outsourcing IT, key IT processes are no longer under full control of the business. This creates the need for a proper governance framework. The cloud service must be evaluated on performance, financial performance, and compliance.

5.2.2.1 Performance evaluation

Technical performance is usually monitored, based on:

- monthly technical performance reports
- monthly exception reports
- quarterly management reviews

5.2.2.2 Financial performance evaluation

Financial performance is usually monitored based on:

- monthly capacity utilization reports
- initial usage polices, built by review of existing (customer's) policies and procedures
- continuous improvements for the usage policies itself and for their cloud usage automation[19]
- continuous improvements
- policies automation

5.2.2.3 Compliance

Compliance is usually evaluated yearly, based on:

- yearly IT audits
  - businesses should ensure their service providers have yearly audits
  - ideally the audit requirement is a section in the SLA
- a *Statement on Auditing Standards No. 70 (SAS70*), which is the current standard for Reporting on Controls at a Service Organization
- third-party assurance for service organizations such as:
  - an *SSAE 16*, which is an enhancement to, the SAS70 report
  - an International Standards for Assurance Engagements No. 3402 (ISAE3402 - Assurance Reports on Controls at a Service Organization) report
- if your provider is ISO certified, the third-party certification statements
  - most likely at least for ISO/IEC 20000 and ISO 27001
- proof of compliance with the GDPR, CLOUD Act and other data protection legislation

[19] For example, cloud usage policy-driven brokering for the different clouds with different pricing for the hybrid cloud, dynamic scale-up and scale down by the performance-driven or utilization monitoring-driven rules or business demands, forecasts, promo actions.

Exam preparation: chapter 5

'Get it' questions

1 / 5
Does Capital Expenditure (CAPEX) in cloud computing always result in savings or reduction of costs?

2 / 5
What types of benefit are listed below, operational or staffing?

Benefit:	Operational	Staffing
Instant server deployment		
Less training		
Backups as a Service		
Managed services		

3 / 5
Before you sign a contract with a cloud provider, it is wise to prepare yourself and to do a comparative study of several providers.

Give an example of a question you should ask.

4 / 5
Part of proper cloud governance is making sure the cloud service provider's performance matches the SLAs.

What should you do to evaluate the technical performance of the cloud service provider?

5 / 5
What key benefit areas of the RoI study should be taken into account?

Exam terms

- evaluation
- costs and savings
- operational benefits
- performance factors
- management requirements
- satisfaction with service providers
- evaluation of service providers
- capital expenditure (CAPEX)
- operational expenditure (OPEX)

Answers to 'get it' questions

1 / 5
No. Capital expenditure may be replaced by subscription costs and pay-per-use costs, and operational expenditures.

See also section The total cost of ownership (TCO).

2 / 5

Benefit:	Operational	Staffing
Instant server deployment	✓	
Less training		✓
Backups as a Service	✓	
Managed services	✓	

See also section Other benefits.

3 / 5
Your answer should be one of these main questions:

- Can the cloud services support the business?
- How robust is the security of the cloud solution?
- Does the cloud provider's SLA measure up to SLAs your business has?
- What is the system performance of the cloud solution?
- What are the cloud service billing mechanism and units?

See also section Evaluating performance factors, management requirements, and satisfaction factors.

4 / 5
You should check the monthly technical performance reports, exception reports, and quarterly management reviews.

Please also see the section Evaluating service providers and their services.

5 / 5
The key benefit areas of the RoI study that should be taken into account are:

- TTM
- scalability
- modern productivity
- assets utilization

Please also see the section The return on investment (RoI) perspective.

List of basic concepts

Terms are listed in alphabetical order.

AAA / Triple A (authentication, authorization, accounting)
application
application hosting
audit
availability
back-up
back-up service
bandwidth
blog
bps (bits per second)
Bps (Bytes per second)
business logic

capital expenditure (CAPEX)
cell phone
CIFS (common internet file system)
claim-based solution
client
client-server
cloud access architecture
cloud presence
cloud technology
common carrier
compliance
confidentiality
cost
CRM tool (customer relation management tool)
customer
datacenter
database
datacenter architecture
denial-of-service attack (DoS)
deployability
digital identity
distributed denial-of-service attack (DDoS)
distributed management taskforce (DMTF)
Dropbox
e-commerce
economic benefit
e-mail
encrypted federation
extranet
failover
federation
frame relay network
GDPR (General Data Protection Regulation)
green IT

guest operating system
hardware
HTML (hypertext markup language)
hybrid cloud
hypervisor
IaaS (infrastructure as a service)
identity
identity management
IM (instant messaging)
IMPS (instant messaging and presence service)
Institute for Electrical and Electronics Engineers (IEEE)
integrity
Internet protocol security (IPSec)
interoperability
intranet
ISO (International Standards Organization)
IT infrastructure
IT service
JavaScript
JSON (JavaScript Object Notation)
LAN (local area network)
latency
location independent
loosely coupled (architecture)
mainframe
man-in-the-middle attack
memory
messaging protocol
microcomputer
middleware
migration
minicomputer
MMS (multimedia message service)
mobile device
mobility
multiprocessing
multi-programming
multiprotocol label switching (MPLS)
multi-purpose architecture
multi-sided platform (MSP)
multi-user
National Security Agency (NSA)
network
network attached storage (NAS)

network infrastructure
network protocol
online games
Open Cloud Consortium (OCC)
open systems interconnection (OSI)
open virtualization format (OVF)

open-ID
operating system
operational benefit
operational expenditure (OPEX)

PaaS (platform as a service)
pay-as-you-go model

performance factors
permissive federation
personal identifiable information (PII)
portability
Pretty Good Privacy (PGP)
privacy
privacy notice
private cloud
processing
protocol analyzer
public cloud
recovery
redundancy
remote data center
replication
risk
RoI (return on investment)
SaaS (software as a service)
satisfaction factors
scalability
scripting language
security
server
service level
service level agreement (SLA)
service-oriented architecture (SOA)
single sign-on (SSO)
slide share
smartphone

SMS (short message service)
social media
software
staffing benefit
stakeholder
storage
storage management initiative-specification (SMI-S)
subcontracted supplier
supplier contract
support
system management architecture for system hardware (SMASH)
TCO (total costs of ownership)
TCP/IP (transmission control protocol / Internet protocol)
thin client
throughput
tiered architecture
time-to-market
time-to-value
traceability
track
user
utility
verified federation
video telecommunication
virtual machine (VM)
virtualization
virtualization management initiative (VMAN)
virtualized environment
virus (infection)
VoIP (voice-over-Internet protocol)
VPN (virtual private network)
web browser
web frontend
web service management (WS-MAN)
web-based enterprise management (WBEM)
webmail
website
Wiki
Wikispace
workload
XML (extensible markup language)
XMPP (extensible messaging and presence protocol)

References

Aggarwal, S. & McCabe, L. (2009). *The Compelling TCO Case for Cloud Computing in SMB and Mid-Market Enterprises* Hurwitz & Associates. http://www.netsuite.com/portal/pdf/wp-hurwitztco-study-dynamics.pdf

Aitoro, J. R. (2008). *The Basics: A Glossary of Federal Technology: Identity Management* https://nextgov.com

Baker, J. (2010). *How to Evaluate cloud computing Providers* Data Center Knowledge. https://www.datacenterknowledge.com/archives/2010/06/01/how-to-evaluate-cloud-computing-providers

Bernstein, P.A. (1996). Middleware: a model for distributed system services. *Communications of the ACM*, 39(2), 86-98. https://dl.acm.org/doi/pdf/10.1145/230798.230809

Carr, N. (2009). *The big switch: Rewiring the world, from Edison to Google.* WW Norton & Company. 2008. ISBN 9780393062281.

Cloud Industry Forum (CIF) (2012). *Cloud UK - Paper five: Cloud Definitions, Deployment Considerations & Diversity* https://issuu.com/cloudexperts/docs/cif-white-paper-5-2012-cloud-definitions-deploymen

Cloud Security Alliance (2019). *Top Threats to Cloud Computing: Egregious Eleven* https://cloudsecurityalliance.org/artifacts/top-threats-to-cloud-computing-egregious-eleven/

Encyclopaedia Britannica (2015). https://britannica.com

Geva, P. (2010). *Evaluating cloud Cloud computing Computing Services: Criteria to Consider*, https://searchcloudcomputing.techtarget.com/feature/Evaluating-cloud-computing-services-Criteria-to-consider http://searchcloudcomputing.techtarget.com.

Greene, T. (2011). *Researchers Find "Massive" Security Flaws in Cloud Architectures* Computer World. https://www.computerworld.com/article/2499518/researchers-find--massive--security-flaws-in-cloud-architectures.html

Harding, C. (2011). *Cloud Computing for Business: The Open Group Guide* Van Haren Publishing. ISBN: 9789087536572.

Hurwitz, J., Bloor, R., Kaufman, M. and Halper, F. (2009). *How to Calculate the Cost of Applications in a Cloud Computing Data Center* Dummies.com. https://www.dummies.com/programming/networking/how-to-calculate-the-cost-of-applications-in-a-cloud-computing-data-center/

International Standards Organization (2023). *ISO/IEC 22123 Information Technology —Cloud Computing — Part1: Vocabulary, Part 2: Concepts, and Part 3: Reference architecture* https://www.iso.org/home.html

International Standards Organization (2018). *ISO/IEC 20000 Information Technology — Service Management – Part 1: Service management system requirements, and Part 2: Guidance on the application of service management systems* https://www.iso.org/home.html

International Standards Organization (2020). *ISO/IEC 19944 Cloud computing and distributed platforms —Data flow, data categories and data use – Part 1: Fundamentals* https://www.iso.org/home.html

International Standards Organization (2015). *ISO/IEC 27017 Information technology — Security techniques — Code of practice for information security controls based on ISO/IEC 27002 for cloud services* https://www.iso.org/standard/43757.html

International Standards Organization (2022). *ISO/IEC 27002 Information security, cybersecurity and privacy protection — Information security controls https://www.iso.org/standard/75652.html*

International Standards Organization (2018). *ISO/IEC 27000 Information technology — Security techniques — Information security management systems — Overview and vocabulary* https://www.iso.org/obp/ui/#iso:std:iso-iec:27000:ed-5:v1:en

International Standards Organization (2022). *Popular Standards: ISO/IEC 27001 Information Security Management* https://www.iso.org/isoiec-27001-information-security.html

International Standards Organization (2019). *Information Technology (IT) in General* https://www.iso.org/ics/35.020/x/

International Standards Organization (2019). *ISO/IEC 27018 Information technology — Security techniques — Code of practice for protection of personally identifiable information (PII) in public clouds acting as PII processors* https://www.iso.org/standard/76559.html

International Standards Organization (2015). *ISO/IEC 33001 Information technology — Process assessment — Concepts and terminology https://www.iso.org/standard/54175.html*

International Standards Organization (2024). *ISO/IEC 38500 Information technology — Governance of IT for the organization https://www.iso.org/standard/81684.html*

International Standards Organization (2016). *ISO/IEC 27036 Information technology — Security techniques — Information security for supplier relationships – Part 4: Guidelines for security of cloud services* https://www.iso.org/standard/59689.html

International Standards Organization (2019). *ISO/IEC 19086 Information technology – Cloud computing — Service level agreement (SLA) framework – Part 1: Overview and concepts, Part 2: Metric model, Part 3: Core conformance requirements, and Part 4: Components of security and of protection of PII* https://www.iso.org/home.html

International Standards Organization (2017). *ISO/IEC 19941 Information technology — Cloud computing — Interoperability and portability https://www.iso.org/standard/66639.html*

International Standards Organization (1994). *ISO/IEC 7498-1 Information technology — Open Systems Interconnection — Basic Reference Model: The Basic Model https://www.iso.org/standard/20269.html*

Krill, P. (2009). *The Cloud-SOA Connection: IT Groups that Understand SOA May Be Able to Take Better Advantage of the Cloud* https://www.infoworld.com/article/2676125/the-cloud-soa-connection.html

Kuznetzky, D. (2011). *Virtualization: A Manager's Guide* O'Reilly Media. ISBN: 9781449306458.
Lemos, R. (2011). *Recent Breaches Spur New Thinking on Cloud Security* https://www.darkreading.com/attacks-and-breaches/recent-breaches-spur-new-thinking-on-cloud-security/d/d-id/1097493?piddl_msgorder=asc www.darkreading.com
Livingstone, R. (2011). *How Low Cost is That Low-Cost Cloud?* CFO.com https://www.cfo.com/the-cloud/2011/09/how-low-cost-is-that-low-cost-cloud/
MacVittie, L. (2009). *Load Balancing is Key to Successful Cloud-based (Dynamic) Architectures*, dev/central. https://devcentral.f5.com/s/articles/load-balancing-is-key-to-successful-cloud-based-dynamic-architectures, www.Devcentral.f5.com
McKay, D. (2010). *Evaluating Cloud Solutions: What Type of Cloud is Right for Me?* Security Week. https://www.securityweek.com/evaluating-cloud-solutions-what-type-cloud-right-mewww.securityweek.com
Mell, P. and Grance, T. (2011). *The NIST Definition of Cloud Computing* NIST Special publication Publication 800-145. https://csrc.nist.gov/publications/detail/sp/800-145/final
Parkhill, Douglas D. (1966). *The Challenge of the Computer Utility*, Addison-Wesley. ISBN 0201057204.
Rymer, J.R., Staten, J. Burris, P., Mines, C., and Whittaker D. (2014). *The Forrester Wave™: Enterprise Public Cloud Platforms, Q4 2014* Forrester. https://www.forrester.com/report/The+Forrester+Wave+Enterprise+public+cloud+Platforms+Q4+2014/-/E-RES118381
Srinivasan, S. R. (2011). *Multi-Tenancy Misconceptions in Cloud Computing* SYS-CON Media, Inc.
Tetz, E. (2011). *Cisco Networking All-in-One for Dummies* Wiley.

Certified for what's next

Contact EXIN

www.exin.com