

CHAPTER 8. Principles for regulating the cloud (1); conclusions from the ontology of cloud computing networks

a. Introduction – scope of this chapter

Having gone through the norms, prevailing schools of thought and currently applicable regulatory approaches regarding cloud computing or the IT applications more closely related to it in Europe and the US, it is now time to look into the principles and best practices that could be derived from each of the two jurisdictions and could serve as guidelines for regulating the cloud, as a unique technology and the backbone of the IT environment of today and tomorrow.

The following parts of this study will be organized in a manner that will have a twofold aim:

- To explain why we need rules specifically for cloud computing besides those already governing the numerous applications based on it
- To formulate these rules not in the strict form of a draft law but as generic regulatory concepts that each jurisdiction can then adopt and adapt to the particularities of its own legal conventions being sure, however, that, if the laws developed have these concepts at heart, the overall governance of the cloud will be more efficient on a cross-border scale.

The proposed principles will be grouped in three chapters, in particular:

- Those stemming from the architecture of the cloud computing network itself and respond to issues related to the way cloud infrastructure is compiled together (Chapter 8)
- Those stemming from the different actors participating across the cloud cycle, i.e. across the workflows developed and facilitated by cloud computing networks, and respond to the way cloud services, businesses and applications are organized and executed (Chapter 9)
- Those responding to the need to build a governance scheme for cloud computing that will differentiate between regulatory challenges on the local and the global level allowing for the concretization of minimum shared standards among different regulations that will permit a more

b. *Constructing the ontology of the cloud; is the cloud one and only thing after all?*

unified tackling of regulatory issues related to the cloud on a cross-jurisdictional basis (Chapter 10).

In comprehensively presenting the regulatory proposals organized into these three groupings, the following methodological tools will primarily be applied:

- Interdisciplinarity, primarily with regard to the principles falling under Chapters 8 and 9
- Legal pluralism, primarily for the principles falling under Chapters 9 and 10
- Harmonization of norms, primarily for the principles under Chapter 10.

b. *Constructing the ontology of the cloud; is the cloud one and only thing after all?*

One of the most common misconceptions regarding cloud computing is that, in laymen as well as in the regulator's eyes, it is usually seen as a concept with just one meaning, that of the means or the medium for the transfer, storage or processing of personal data. Actually, the term 'cloud computing' is much more multi-layered and complex than that and, before getting down to talk about it as a term signifying a whole range of applications serving the above purposes, it is crucial to realize that the cloud has various different facets on a hardware/architectural level⁸³⁷. Particularities in the nature of these facets already lead to the first regulatory principles necessary for an efficient governance of cloud computing.

In computer science, describing and documenting all variations of a technology or the hardware implementations that make it possible is a process called (IT) ontology⁸³⁸. In detail, in computer science and information science, an ontology is an official, analytical naming and mapping of the types, properties, and interrelations of the entities that exist for a particular domain of discourse, i.e. a particular domain of the overall sec-

837 Deepak Puthal, B.P.S. Sahoo, Sambit Mishra & Satyabrata Swain, *Cloud Computing Features, Issues, and Challenges: A Big Picture*, in 2015 International Conference on Computational Intelligence & Networks (CINE), 116–123 (KIIT University ed.)

838 Lamia Youseff, Maria Butrico & Dilma Da Silva, *Toward a Unified Ontology of Cloud Computing*, in 2008 Grid Computing Environments Workshop, 1–10.

tor⁸³⁹. In other words, ontology in IT is a practical application of philosophical ontology, with a taxonomy⁸⁴⁰. An ontology task, in essence, compartmentalizes the variables needed for specific types of computations and, additionally, establishes the relationships between them⁸⁴¹.

Ontology as a tool and practice is increasingly common in several fields of the wider IT sector⁸⁴². To name a few, the fields of artificial intelligence, the Semantic Web, systems engineering, software engineering, biomedical informatics, library science, enterprise bookmarking, and information architecture all resort to ontologies to limit complexity and organize information about and within them⁸⁴³. These ontologies can then be applied to problem solving⁸⁴⁴. The same practice is suggested as a key tool in our effort to analytically comprehend, systematize and, ultimately, regulate cloud computing.

There are several methodologies with which it is possible to map down the ontology of an IT field⁸⁴⁵. The one mostly proposed in relevant literature as the most suitable to grasp and successfully organize all relevant

839 John F. Sowa, *Top-level ontological categories*, 43 *International Journal of Human-Computer Studies* 669–685 (1995.)

840 *Id.*

841 Lamia Youseff, Maria Butrico & Dilma Da Silva (note 838).

842 Ling Liu & M. Tamer Özsu, *Encyclopedia of database systems* (2009.)

843 *Id.*

844 Xiaolong Jin & Jiming Liu, *From Individual Based Modeling to Autonomy Oriented Computation*, in *International Workshop on Computational Autonomy*, 151–169 (2003.)

845 The two fundamental genres of ontology are domain and upper ontology. Domain ontologies (or domain-specific ontologies) represent concepts which belong to part of the world. Particular meanings of terms applied to that domain (i.e. the world) are provided by domain ontology. For instance, the word card has several meanings. An ontology about the domain of poker would model the "playing card" meaning of the word, while an ontology about the domain of computer hardware would model the "sound card" and "video card" meanings. A main feature of domain ontologies is that they represent concepts in very specific, even eclectic ways, becoming often incompatible. As systems that rely on domain ontologies expand, they often need to merge domain ontologies into a more general representation. At the same time, different ontologies in the same domain arise due to different languages, different intended use of the ontologies, and different perceptions of the domain (based on cultural background, education, ideology, etc.).

Another major type is upper ontology (or foundation ontology), i.e. a model of common objects that are generally applicable across a wide range of domain on-

knowledge regarding the cloud is composability⁸⁴⁶. Composability, as an ontology typification method, is inspired by composability as a system design principle⁸⁴⁷; the latter heavily deals with the inter-relationships of components of a system; in this case, of the cloud, as a field of IT⁸⁴⁸.

For reasons of clarity and simplicity, the ontology of the cloud that is endeavored here should be conventionally pictured as a stack of layers. Then, each layer shall encompass one or more cloud services. In addition, cloud services sharing comparable levels of abstraction will be classified as belonging to the same layer, while abstraction will be measured as per which type of users each service is targeted at⁸⁴⁹. For instance, all cloud software environments (i.e. cloud platforms) target programmers, while cloud applications target end users. Therefore, cloud software environments would be all classified in the same but in a different layer than cloud applications, which would, however, also fall all under the same layer.

Under composability, one cloud layer is classified as being higher in the cloud stack, when its services can be composed from the services of the underlying layer⁸⁵⁰. For example, when it comes to the cloud application layer, since cloud applications are made possible, i.e. are developed, using cloud software environments, it can be said that cloud applications are composable from cloud software environments, and, consequently, the cloud application layer is higher in the cloud stack⁸⁵¹. Following this logic, the cloud stack is composed from bottom up of the following layers:

- The Firmware/hardware layer (HaaS)
- The Software Kernel layer
- The Cloud Software Infrastructure layer, which is further broken down to Computational Resources (IaaS), Storage (DaaS), and Communications (CaaS)

tologies. It usually employs a core glossary that contains the terms and associated object descriptions as they are used in various relevant domain sets.

Lastly, a hybrid is an ontology incorporating elements from both the domain and upper model.

846 Lamia Youseff, Maria Butrico & Dilma Da Silva (note 838).

847 John F. Sowa (note 839).

848 Lamia Youseff, Maria Butrico & Dilma Da Silva (note 838).

849 *Id.*

850 John F. Sowa (note 839).

851 Lamia Youseff, Maria Butrico & Dilma Da Silva (note 838).

- The Cloud Software Environment layer and,
- The Cloud Application layer (SaaS)⁸⁵².

An analytical presentation of these layers will permit us, afterwards to pinpoint some essential regulatory guidelines for the cloud.

i. The Firmware/Hardware layer

By application of the tools described above, the ontology scheme of the cloud has the firmware/hardware layer at its foundations. It comprises the actual physical hardware and infrastructure that form the backbone of the cloud, as technology and as network⁸⁵³. On this layer, the main users are big enterprises with voluminous IT requirements which, most commonly, are in need of a service constituting of subleasing hardware which they will then use for their own computational needs or purposes (HaaS)⁸⁵⁴. As a rule, the entities acting as HaaS providers at this level have the tasks of operating, managing and upgrading the said hardware on behalf of their consumers, for as long as the sub-lease contracts they have entered into with customers remain in force. One of the classic examples of HaaS are the contracts banking service providers enter into with big data storage providers in order to cover their computational needs⁸⁵⁵. At this layer, users have predefined workloads with characteristics that impose strict performance requirements.

ii. The Software Kernel layer

On this cloud layer are to be allocated all pieces of basic software management for the physical servers composing the cloud. Software kernels⁸⁵⁶ at

852 IEEE INFOCOM 2010 – IEEE Conference on Computer Communications.

853 Mike P. Papazoglou & Willem-Jan van den Heuvel, *Service oriented architectures. Approaches, technologies and research issues*, 16 The VLDB Journal 389–415 (2007.)

854 *Id.*

855 Morgan Stanley's sublease contract with IBM in 2004.

856 In computer science, the kernel (also named the nucleus) is a computer program that constitutes the core of a computer's (or computer network's) operating system. The kernel has complete control over everything that occurs in the system. As such, it is the first program loaded on system startup, and it then manages the

this level are implemented as an OS kernel⁸⁵⁷, hypervisor⁸⁵⁸, virtual machine monitor⁸⁵⁹ and/or clustering middleware⁸⁶⁰. Traditionally, grid computing applications were deployed to run on this layer on several interconnected clusters of machines⁸⁶¹. However, due to the absence of the virtualization element in grid computing, those tasks were closely tied to the actual hardware infrastructure; consequently, providing migration, checkpointing and load balancing to the applications at this level used to be a complicated task⁸⁶². In the meantime, a considerable body of research in grid computing has led to several grid-developed concepts being realized today in cloud computing⁸⁶³.

remainder of the startup process, as well as input/output requests from software, by translating them into data processing instructions for the central processing unit. It is also responsible for managing memory, and for communicating with computing peripherals, like printers, speakers, etc. The kernel is a fundamental part of a modern computer's operating system. Mutatis mutandis, in the context of a cloud computing network the kernel is its most basic software, the one managing its most fundamental and elementary functions and processes, which are basically dedicated in making sure that the network itself will run properly.

- 857 The OS kernel as a term essentially is synonymous to the term 'software kernel'.
- 858 A hypervisor or virtual machine monitor (VMM) is a piece of computer software (there are firmware or hardware typifications of hypervisors but they call outside the scope of this study) that creates and runs virtual machines. The hypervisor presents the guest operating systems with a virtual operating platform and manages the execution of the guest operating systems. The term hypervisor is a variant of supervisor, a traditional term for the kernel of an operating system: the hypervisor is the supervisor of the supervisor, with hyper- used as a stronger variant of super-.
- 859 A virtual machine is a software computer that, like a physical computer, runs an operating system and applications. The virtual machine is comprised of a set of specification and configuration files and is backed by the physical resources of a host.
- 860 In the context of a computing cluster, the activities of computing nodes are orchestrated by "clustering middleware", a software layer that sits atop the nodes and allows the users to treat the cluster as by and large one cohesive computing unit, e.g. via a single system image concept.
- 861 Stephanos Androutsellis-Theotokis & Diomidis Spinellis, *A survey of peer-to-peer content distribution technologies*, 36 ACM Comput. Surv. 335–371 (2004.)
- 862 2008 Grid Computing Environments Workshop.
- 863 Ian Foster, Yong Zhao, Ioan Raicu & Shiyong Lu (note 92).

iii. The Cloud Software Infrastructure layer

The cloud software infrastructure layer hosts fundamental resources which are essential so that other higher-level layers can be used to construct new cloud software environments or cloud applications. The main reason why resources allocated on this layer are set apart from the two highest levels in the cloud stack is that the latter can bypass the cloud infrastructure layer in directly building their system⁸⁶⁴. Often this bypass can enhance the efficiency of the system, yet it comes at the cost of simplicity and minimum development efforts necessary⁸⁶⁵. The services allocated on this layer are further divided into: computational resources, data storage, and communications.

- computational resources: Virtual machines (VMs) are the most common form for providing computational resources to cloud users at this layer which they can subsequently use to customize the software stack for performance and efficiency⁸⁶⁶. Conventionally, such services are dubbed Infrastructure as a Service (IaaS)⁸⁶⁷. Virtualization is the enabling technology which offers unprecedented flexibility to users in configuring their settings while protecting the physical infrastructure of the provider's data center⁸⁶⁸. However, since VMs can by nature co-exist on the same data storage hardware facility, the lack of a strict performance isolation between them while sharing the same physical node can at any time result in the inability of cloud providers to give strong guarantees for performance to their clients⁸⁶⁹. Such weak guarantees, unfortunately, can inject themselves up the layers of the cloud stack⁸⁷⁰.
- data storage: The second infrastructure resource is data storage, which constitutes what cloud computing is probably most widely known for: allowing users to store their data at remote storage facilities and access them anytime from anywhere⁸⁷¹. This service is commonly quoted as Data-Storage as a Service (DaaS), and it permits cloud applications to

864 Deepak Puthal, B.P.S. Sahoo, Sambit Mishra & Satyabrata Swain (note 837).

865 Mike P. Papazoglou & Willem-Jan van den Heuvel (note 853).

866 Refer also to Chapter 2.

867 Refer also to Chapter 2.

868 Refer also to Chapters 2 and 6.

869 Dimitrios Zissis & Dimitrios Lekkas, *Addressing cloud computing security issues*, 28 *Future Generation Computer Systems* 583–592 (2012.)

870 Lamia Youseff, Maria Butrico & Dilma Da Silva (note 838).

871 See Chapter 2.

scale beyond their limited servers. Data storage systems are as a standard expected to meet several rigorous requirements for maintaining users' data and information, including high availability, reliability, performance, replication and data consistency⁸⁷²; however, precisely because of the conflicting nature of all these requirements, no system can implement all of them together⁸⁷³. For instance, availability, scalability and data consistency are regarded as three conflicting goals from a technical point of view. Given that those features are hard to be simultaneously achieved with general data storage systems, DaaS-providers implement their system to favor one feature over the others, while indicating their choice through their SLA. However, there is no legal warranty at the moment regarding the minimum that needs to be achieved for any one of the most common performance requirements causing considerable irregularities and, thus, insecurities throughout the cloud market.

- communication: As cloud systems evolve and become more and more popular and the means for developing a wide range of IT services for the general public, so does the need for guaranteed quality of service for network communication, with communication becoming a vital component of the cloud infrastructure. As a result of this demand, cloud systems have focused on developing features enhancing communication capability in a service-oriented, configurable, schedulable, predictable, and reliable manner⁸⁷⁴. Towards this end, the concept of Communication as a Service (CaaS) emerged. Although at the beginning this model was the least discussed and adopted in commercial cloud systems, it is gaining more and more in popularity over the last years⁸⁷⁵. Inter alia, systems that belong to CaaS are VoIP telephone systems, audio and video conferencing as well as instant messaging apps are cloud applications that are already or are expected to be based on CaaS⁸⁷⁶.

872 See Chapter 2.

873 *Id.*

874 Ozalp Babaoglu, M. Jelasity, Anne Marie Kermarrec, Alberto Montresor & Maarten van Steen, *Operating Systems Review (ACM)*, available at: <http://dl.acm.org/citation.cfm?doid=1151374.1151379>.

875 Mike P. Papazoglou & Willem-Jan van den Heuvel (note 853).

876 Lamia Youseff, Maria Butrico & Dilma Da Silva (note 838).

The three sublayers composing the infrastructure layer also share common challenges besides the ones particular to each of them. Among others, security of the services, availability and quality are the most commonly addressed concerns for all cloud infrastructure components⁸⁷⁷.

iv. The Cloud Software Environment layer

The following layer in cloud ontology is the cloud software environment layer (or, simply, the software platform layer). Users of this layer are cloud applications' developers, who implement their applications for and deploy them on the cloud⁸⁷⁸. Providers of this layer, on the other hand, supply developers with a programming-language-level environment aimed at facilitating interaction between programming environments and cloud applications, as well as at accelerating deployment and supporting scalability necessary for those cloud applications⁸⁷⁹. Services provided by cloud systems in this layer are commonly referred to as Platform as a Service (PaaS)⁸⁸⁰. A classic example of systems in this layer is Google's App Engine, which provides a python runtime environment and APIs for applications to interact with Google's cloud runtime environment or Salesforce's Apex language permitting developers of cloud applications to design the page layout, workflow or customer reports according to the logic of their applications⁸⁸¹. In a nutshell, cloud software environments facilitate the process of the development of cloud applications⁸⁸².

v. The Cloud Application layer (SaaS)

The cloud application layer is the one closest to the end-users of the cloud. It basically corresponds to the very cloud-based applications we all know and use in daily life, from our email service, to Dropbox or similar file storage and management services etc. This model has exponentially

877 Deepak Puthal, B.P.S. Sahoo, Sambit Mishra & Satyabrata Swain (note 837).

878 Lamia Youseff, Maria Butrico & Dilma Da Silva (note 838).

879 *Id.*

880 See Chapter 2.

881 Xiaolong Jin & Jiming Liu (note 844).

882 Lamia Youseff, Maria Butrico & Dilma Da Silva (note 838).

gained popularity for all the reasons explained in earlier parts of this study⁸⁸³.

c. Different uses but the same ontology: what does this mean for cloud computing regulatory principles?

From the analysis in the previous section in combination with the technical overview of cloud computing in Chapter 2 of this study we can draw the conclusion that there is a clear dichotomy between cloud computing as a technical arrangement, as a technology and infrastructure, on the one side, and the cloud as the applications through which we have the possibility to use in various forms the capacity of this infrastructure, on the other.

Viewing the above observation through the basic reasoning proposed by the doctrine of law and knowledge⁸⁸⁴, which is rapidly gaining popularity particularly in public law, it can be argued that this dichotomy has caused a fussy picture, at least on the front of end-consumers and on the regulatory front, due to the fact that the infrastructural nature of the cloud is not, at most times, immediately visible and, therefore, comprehensible to non-technically-savvy actors. It is of course, undeniable that there are lots of different ways to deploy the same kind of infrastructure and this means that the (regulatory) challenges coming with one type of cloud environment will not necessarily be the same with those of another. For instance, a great deal of issues regarding privacy raised by public clouds are non-existent or they are satisfactorily tackled when the same resources are utilized to set up a private cloud computing network⁸⁸⁵. However, the technical expertise, the mechanical skills and the very materials (i.e. pieces of

883 See Chapter 2.

884 Hans-Heinrich Trute (note 432). For further details on the doctrine of law and knowledge and the broader issue of how knowledge converts into or affects the law, refer to: Hans Christian Röhl, *Wissen, zur kognitiven Dimension des Rechts*, vol. 9 (2010); Gunnar Folke Schuppert & Andreas Vosskuhle, *Governance von und durch Wissen*, Bd. 12 (2008); Mariana Valverde, *Law's Dream of a Common Knowledge* (2009); Brett M. Frischmann, Michael J. Madison & Katherine Jo Strandburg, *Governing knowledge commons* (2014); Friedrich A. von Hayek, *The Use of Knowledge in Society*, 35 *The American Economic Review* 519–530 (1945); Adrian Vermeule ed., *Local and Global Knowledge in the Administrative State* (2013.)

885 See Chapter 2 for the difference between public and private cloud networks.

hardware) that are necessary in order to build up either a public (with just the standard protection features) or a private (with as advanced protection features as possible) cloud ecosystem are, in essence, the same. In both cases, and in every other in between, one will need pieces of the same kind of infrastructure, the same kind of information science and IT engineering knowledge that will permit one to put those pieces of hardware into meaningful working arrangements and, of course, even the features that will differentiate them and make them stand apart from each other will be based on the same technical principles and scientific intel that makes the overall concept of cloud computing technology possible. Consequently, it becomes evident that, despite the great variety in which cloud services and networks appear on the market and the substantial differences which might exist between all these variations of cloud environments, there is a common underlying connecting tissue that binds them all, and that is the knowledge (of informatics, computing engineering and other disciplines) related to them which is one and the same.

To put it more illustratively, let's take the example of two data hosting and sharing facilities, such as Dropbox, one public and commercially available and the other private and customized to be accessible by a specific circle of users only, probably also cut out in a manner that will provide answers to their very particular needs. It is true that a great deal of elements of the two applications might look totally different from each other, from the interfaces to the layers and tools each of them uses to ensure privacy and security for its users. But no matter how different the two applications may look, the basic principles and knowledge behind them are the same; as a result, from each of these two manifestations of the cloud there are minimum common expectations which call for minimum shared regulatory principles that would settle them in a unanimous manner. This unanimity could and should be not just within the boundaries of one jurisdiction but on a cross-jurisdictional basis. This does not in any case necessitate some kind of unification of different jurisdictions into one or the introduction of one extra supranational legal order just for the sake of IT regulation. Jurisdictional particularities and traditions of every legal order could very well be upheld and respected in the field of IT law as it is done in any other legal sector. What we need to make sure is that these commonly shared principles will advance the achievement of the same goals from every jurisdiction on each and every matter of cloud computing regulation.

In other words, the challenge is not to homogenize IT laws or pulverize jurisdictional particularities. It rather is to set common goals and establish rules that will contribute to their achievement. The path towards achieving these goals can and will expectedly be different, both because cloud computing manifests itself through various different arrangements and because two or more identical cloud networks in different environments will naturally be treated in differentiated manners according to the legal culture in each environment. However, as long as the same purposes are pursued and, ultimately, materialize, the path and the means need not be identical.

With this in mind, the ontology of the cloud as it was previously analyzed allows us to define a first set of regulatory principles for cloud computing based on the knowledge that makes the cloud possible.

d. Mapping the life cycle of data on cloud computing networks: risks, security and privacy issues as indicators for the nature of cloud computing regulation rules

Having analyzed what cloud computing as technology and technological arrangement actually consists of via the tool of ontology, it is worth also mapping down the life cycle data follows while circulating through the various layers presented above. Presenting the blueprint of the path of data through the cloud will also allow us to pinpoint the risks they are exposed to from a technical perspective. This knowledge, which, as it has been argued in the case of cloud ontology already, is universal and applies to all different kinds of cloud networks no matter whether they host public services or others available only to a limited circle of users, can then lead us to the concretization of the regulatory principles stemming from the ontology of the cloud.

For starters, in the context of the analysis following below, the term ‘data life cycle’ should be interpreted as referring to the entire process from generation to destruction of any kind of digital data⁸⁸⁶. This path consists of seven distinct stages, the essence, features and main risks of which are summarized as follows. It needs to be noted that, in keeping with the dynamic relations between the different layers of the cloud ontol-

886 Deyan Chen & Hong Zhao, *Data Security and Privacy Protection Issues in Cloud Computing*, in 2012 International Conference on Computer Science and Electronics Engineering, 647–651 (2012.)

ogy as they were previously analyzed, the stages of the data life cycle on the cloud can occur equally dynamic. It is only for descriptive easiness that they are hereunder separately analyzed and their sequence of presence does not imply at all that each stage is immune to or sealed from the others.

i. Data generation

Data generation describes the moment when data is actually created for the first time, regardless of whether it is original or data resulting from processing of preexisting data sets⁸⁸⁷. At this stage of data generation, the ownership status of data is also determined⁸⁸⁸. In pre-cloud IT environments users of whichever size, i.e. from individual users to large scale organizations, used to own and manage the data they were the creators of⁸⁸⁹. However, in an IT environment where data increasingly, if not by default, migrate to the cloud immediately after their creation or they are even created directly there, the issue of ownership cannot be answered so self-evidently. In other words, regulatory principles are need which will either allow the question of data ownership to be answered at all times during the circulation of data on a cloud network or, if so preferred, will provide enough safeguards to data owners regarding what extend of their personal private information is being collected by other actors on the cloud network. Last but not least, principles that will determine under which conditions data owners may put a stop to collection and use of personal information regardless of the layer within the cloud network where such practice occurs are also necessary. However, it needs to be made sure, at the same time, that these rules need be realistic and promise realistic levels of protection to data owners, unlike what seems to happen with the respective provisions of the GDPR⁸⁹⁰.

887 Michael Backes & Peng Ning eds., *Computer security – ESORICS 2009*. 14th European Symposium on Research in Computer Security, Saint-Malo, France, September 21-23, 2009 : proceedings, vol. 5789 (2009.)

888 Elen Stokes, *Regulating Technologies: Legal Futures, Regulatory Frames and Technological Fixes* by Roger Brownsword and Karen Yeung (eds), 73 *The Modern Law Review* 682–689 (2010.)

889 Xiaolong Jin & Jiming Liu (note 844).

890 See Chapter 4.

ii. Transfer

Another fundamental block of the overall life cycle of data on the cloud is the transfer of them. As a rule, in the pre-cloud enterprise status quo data transmission did not require encryption or, at the most, only simple data encryption measures would suffice⁸⁹¹. However, in the new enterprise environment facilitated by cloud computing, where it is far from given that the network of one enterprise does not overlap with that of another, transfers of data are commonplace and both data confidentiality and integrity should be ensured in order to prevent tapping and tampering by unauthorized users⁸⁹². From a technical point of view, this cannot be guaranteed by data encryption alone nor with technical measures only⁸⁹³. For optimal data integrity in the cloud confidentiality is also crucial⁸⁹⁴ and it can only be achieved if trustworthy transfer protocols are legally necessitated. In fact, these should be maintained not only throughout the length of a single cloud network but also during circulations of data from one network to the other. In other words, the relevant rules providing for integral transfer mechanisms should be developed having in mind both the horizontal and the vertical data transfers which are possible in cloud environments.

iii. Use

While being used on the resources of a network digital static data appear in either of the following formats: as static data being used on a simple storage service (such as most of storage services addressed to end-users, like Amazon S3 or Dropbox) where data encryption is already feasible⁸⁹⁵. However, static data on the cloud can also be used by cloud-based applications on the PaaS or SaaS layer and, in those cases, data encryption is not always feasible⁸⁹⁶. In fact, on layers prior to the end-applications level data encryption is very likely to lead to problems of indexing and query, ab-

891 S. Subashini & V. Kavitha (note 119).

892 Nir Kshetri, *Privacy and security issues in cloud computing. The role of institutions and institutional evolution*, 37 Telecommunications Policy 372–386 (2013.)

893 Deyan Chen & Hong Zhao (note 886).

894 Deyan Chen & Hong Zhao (note 886); David W. Opderbeck (note 628).

895 Ozalp Babaoglu, M. Jelasity, Anne Marie Kermarrec, Alberto Montresor & Maarten van Steen (note 874).

896 *Id.*

normalities which would undermine the smooth functioning of the overall cloud network⁸⁹⁷. It turns out that, contrary to what may be commonly perceived as a result of the perception simple end cloud services try to cultivate on customers⁸⁹⁸, in cloud just as in traditional IT environments, the data being treated is almost not encrypted for any program that deals with it on a layer prior to the end applications level⁸⁹⁹. Moreover, due to the multi-tenancy feature⁹⁰⁰ of cloud computing models, the data being processed by cloud based applications is in many instances stored together with the data of other users at least when they are used by applications and actors other than the end users. Given that this technical arrangement is technically utopic that it will cease to exist, it becomes evident that regulatory principles defining codes of conduct for any actor using data at any point during their life cycle and on any layer of a cloud network are necessary.

iv. Sharing

Data sharing, which is a function continuously performed by data owners and several different types of actors that have access to data stored on a cloud network, is an action expanding the use range of the data thus rendering data permissions more complex⁹⁰¹. This is of course a very known issue about cloud computing, which existing legislation is already striving to cope with, at least with regard to the specific cloud-enabled applications for which there is regulation in place. However, given that data owners

897 Deyan Chen & Hong Zhao (note 886).

898 Huaiqing Wang, Matthew K. O. Lee & Chen Wang (note 12).

899 Deyan Chen & Hong Zhao (note 886).

900 "Software multitenancy", which is largely considered as one of the cornerstone features of cloud computing, refers to a software architecture in which a single instance of software runs on a server and serves multiple tenants. The term "tenant" denotes a group of users who share common access with specific privileges to the software instance. Under multitenant architecture, a software application is designed to provide every tenant a dedicated share of the instance – including its data, configuration, user management, tenant individual functionality and non-functional properties. Multitenancy contrasts with multi-instance architectures, where separate software instances operate on behalf of different tenants. For more, refer to: Krebs, R., Momm, C., & Kounev, S. (2012). Architectural Concerns in Multi-tenant SaaS Applications. *Closer*, 12, 426-431.

901 (note 852).

can authorize data access for one party, which can further share the data with another party without the consent of the original data owner and taking into account that this chain of sharing occurrences can go on and on extending to users that are far from the jurisdiction of the data owner, we can never realistically expect that simply by devising new methods for extending law applicability universal legal safety cannot be achieved. Therefore, only the endorsement of common regulatory principles on the cloud and the use as foundations of cloud governing laws by as many jurisdictions as possible can be expected to provide trustworthy answers to the issues discussed.

v. Storage

Possibly the most common activity with regard to data on the cloud is storage. In fact, data is stored on cloud networks in two distinct contexts, i.e. in IaaS environments, such as those of any standard cloud storage service, and in PaaS or SaaS environment, where data related to the core code of cloud based applications are stored⁹⁰².

In computer science, data stored in cloud storages is treated in the same manner as data stored in any other kind of facility, pre-existing or concurrent to cloud computing⁹⁰³. With that in mind, computer science literature applied to data stored on the cloud the classic three criteria in order to assess how securely they are stored⁹⁰⁴: confidentiality, integrity and availability.

As far as data confidentiality is concerned, the solution advanced so far from a technical perspective is data encryption⁹⁰⁵. The particularities of cloud environments, involving large amounts of data transmission, storage and handling, as well as processing speed and computational efficiency of encrypting large amounts of data, make the use of symmetric encryption

902 Deyan Chen & Hong Zhao (note 886).

903 Ozalp Babaoğlu, M. Jelasity, Anne Marie Kerमारrec, Alberto Montesor & Maarten van Steen (note 874).

904 Nir Kshetri (note 892).

905 Robert Gellman (note 696).

algorithms⁹⁰⁶ more suitable than asymmetric ones. Moreover, another key question coming immediately after the choice of the most suitable encryption pattern is key management⁹⁰⁷ and who is responsible for it. An ideal answer would be that it is the data owners but, for the time being and in the foreseeable future, average users do not possess enough expertise to manage these keys and, as a standard, they entrust key management with the cloud providers. Consequently, the enormous range of tasks for the latter means their key management responsibilities are way more complex and difficult to cope with but, in any case, imperative. Switching focus to data integrity⁹⁰⁸, the essential question is how users, who put several gigabytes or more of data into the cloud, can check the integrity of it. This turns out to be not an easy question to answer given that rapid elasticity as an elementary feature of cloud computing resources makes it impossible for the average end user to know where their data is being stored at all times⁹⁰⁹. As data is dynamic in cloud storage environments, traditional technologies to ensure data integrity may not be effective⁹¹⁰. Last but not least, in a traditional IT environment the main threat to data availability comes from external attacks⁹¹¹. In the cloud, however, in addition to external attacks, there are several other factors that may put data availability under threat⁹¹², namely the availability of cloud computing services; whether cloud providers have committed themselves to continue to operate in the future or what safeguards they have undertaken in case their op-

906 Symmetric-key algorithms (applied in symmetric encryption) are algorithms for cryptography that use the same cryptographic keys for both encryption of plaintext and decryption of ciphertext. The keys may be identical or there may be a simple transformation to go between them. In the actual practice of data cryptography, the keys represent a shared secret between two or more parties that can be used to maintain a private information link. This requirement that both parties have access to the secret key is one of the main drawbacks of symmetric key encryption, in comparison to public-key encryption (which is what is known as asymmetric key encryption). For more details, refer to: Hans Delfs & Helmut Knebl, *Introduction to cryptography. Principles and applications*, 2007: 1 (2007); Christof Paar & Jan Pelzl, *Understanding cryptography. A textbook for students and practitioners* (2010.)

907 Hans Delfs & Helmut Knebl (note 906).

908 Deepak Puthal, B.P.S. Sahoo, Sambit Mishra & Satyabrata Swain (note 837).

909 See also Chapter 2.

910 Deyan Chen & Hong Zhao (note 886).

911 S. Subashini & V. Kavitha (note 119).

912 Michael Backes & Peng Ning eds. (note 887).

eration is suspended; or whether the cloud storage services provide trustworthy backup functionalities.

As it becomes evident from the analysis above, there is not a single nor an obvious answer as to which actor throughout the layers of a cloud network is responsible for making sure storage of data standards live up to the expectations that have been described as essential. In other words, having in place rules that put the burden of such issues to specific entities relevant to specific types of cloud-enabled business will never be a regulatory strategy generic enough to provide us with answers to these challenges in every instance, even in situations which may not be market applicable at this point but are technically feasible in any case. As a result, the need for general cloud computing regulatory principles emerges once more.

vi. Archival

The key criteria for evaluating archiving of data from a technical perspective⁹¹³ are the storage media on which archival is done, whether off-site storage is provided or not and, last but not least, how long archival storage lasts. If the media chosen for archival are portable and, at some point, they get out of control, the archived data are exposed to the risk of leakage. On the other hand, if cloud service providers do not provide off-site archiving, availability of data is put under question. In addition, archival services are not adequate if they are not made to last over a certain minimum amount of time; otherwise, they may result in availability or privacy threats. These issues occurring with reference to archiving as a fundamental function of cloud services should also be answered in the framework of a set of generic regulatory principles for the cloud.

vii. Destruction

When a given set of data is no longer required, it needs to be destroyed⁹¹⁴. The physical dimension of cloud computing storage facilities as well as

913 Deyan Chen & Hong Zhao (note 886).; Dimitrios Zissis & Dimitrios Lekkas (note 869).

914 Deyan Chen & Hong Zhao (note 886).

the archiving capabilities cloud services are supposed to offer in order to increase integrity and availability of data pose questions as to after what point data can be regarded to have been effectively deleted without any possibility of being restored⁹¹⁵. Given the different variations of data and the different instances throughout the cloud ontology when they might be created or be rendered useless, it becomes again evident that generic principles for governing the cloud are highly advisable.

e. Regulatory principles derived from the ontology of cloud computing

Bearing into account the ontology of the cloud as it was analytically described above we can recognize on each layer certain functions and/or actors that primarily aim at the same goals with their functions no matter whether the cloud network they are part of is a public or private one. Therefore, making use of the teleological perspective, it can be argued that, despite the particularities of each network and its specific features, which may neutralize some challenges or, anyway, make them easier to be tackled by respective actors, we can agree on minimum rules that will need to be observed by the network and its constituent entities so that the ultimate goal of the entire workflow is fulfilled⁹¹⁶:

i. On the hardware/firmware layer

As we have seen, this constitutes the backbone of the cloud network, primary gravity is placed on the issue of security, integrity and (constant) availability of resources⁹¹⁷. Given that, regardless of whether a cloud network's infrastructure is utilized by the network owners themselves or whether it is outsourced to third parties, it has to maintain at all times high levels of security and integrity in order for the data stored or the processes executed on it to be available and run smoothly at all times⁹¹⁸, rules contributing to the achievement of these prerequisites are of vital importance. This trend can already be observed across various examples of resource

915 Dimitrios Zissis & Dimitrios Lekkas (note 869).

916 See also Chapter 5.

917 Stephanos Androutsellis-Theotokis & Diomidis Spinellis (note 861).

918 See also Chapter 5.

outsourcing on this layer, with service level agreements (SLAs) imposing strict rules and obligations to cloud network owners/managers who lease, in whole or in part, their resources to third parties⁹¹⁹. Given that these requirements of integrity and constant availability are already considered as *sine qua non* by all affected actors⁹²⁰, it is high time for them to be incorporated into laws. All the more so if we take into account the fact that the current regime of SLAs, which are the subject matter of negotiation between contracting parties, regularly leads to situations of imbalances where, even within the same market, greatly variable degrees of integrity and trustworthiness are required or expected from cloud infrastructure owners, even though their resources will be utilized for the provision of cloud-based services or the execution of equally sensitive computational tasks⁹²¹. It goes without saying that the relevant rules should legislate on the minimum standards necessary leaving of course room for even more elevated commitments at the discretion of the parties in each and every case. Regulating on the minimum standards and leaving room for more elevated commitments at the discretion of the parties will also contribute to the rules that will be adopted being more harmonized on a cross-jurisdictional level since law subjects of a particular jurisdiction will be able to expand their activities to others simply by adapting the standards on all or part of their infrastructure to those prescribed by the jurisdiction(s) they wish to enter. As it has been demonstrated⁹²², mechanics of the cloud perfectly permit the infrastructure of a cloud network to be treated either unanimously or in a compartmentalized manner. Consequently, if the law states the minimum standards a cloud network needs to uphold at all times, it is then always possible to divide part of the overall resources and adjust it to further elevated standards in order to satisfy requirements of more than one jurisdictions at the same time. The only requirement would be, of course, to have a basic principle of non-confluence between resources utilized for processing tasks falling under rules dictating different standards. This does not imply at all that the principle of ultimate utilization of resources (which it should never be forgotten that it is one of the core features of the cloud⁹²³) should be compromised. On this issue, we

919 Xiaolong Jin & Jiming Liu (note 844).

920 See also Chapter 5.

921 Lamia Youseff, Maria Butrico & Dilma Da Silva (note 838).

922 See Chapter 2.

923 Refer also to Chapter 2.

should once again resort to the engineering flexibilities that characterize cloud technologies and can ensure that resources of a network dedicated to processing tasks falling under the rules of a specific jurisdiction will abide by the minimum standards set by the rules of that jurisdiction and will adapt themselves when switched to tasks governed by different laws.

Regarding the security challenges at this level, from a technical point of view at HaaS a developer has better control over security⁹²⁴; nevertheless, the length of this grip should not provoke any security gap in the virtualization element of the cloud network. Similarly, on the other side of the coin, virtual machines have in principle the capacity to address these integrity of virtualization issues, yet in practice there are a lot of security questions that remain unsettled⁹²⁵. The other security element that keeps calling urgently for resolution is the unwavering quality of the information that is put on the cloud supplier's infrastructure. The powerful presence of virtualization across all types of cloud processing and the proximity in which it brings data from different users make both holding a definitive control over information and paying respect to the physical area/resources that hosts it primary responsibilities of the information owner/cloud resources user⁹²⁶. It becomes clear that, in order to achieve most extreme trust and security on the HaaS layer, a few procedures starting from both sides of the provider and user need to be coordinated. Currently, security obligations of both supplier and client incredibly vary from one cloud network to the other due to different cloud administration models which, at the lack of minimum requirements prescribed by laws, are arbitrarily developed by the market⁹²⁷. Undoubtedly, a private cloud is better protected against security threats on the infrastructure level compared to a public cloud. Nonetheless, regardless of the deployment model every single cloud facility has one elementary yet extremely crucial challenge to live up to: protecting the physical infrastructure of data centers⁹²⁸. Relevant basic rules should reflect on damage done by any natural disaster but also any damage done to the facility deliberately. It should not fail our attention that in every case the infrastructure that needs to be protected is not only the hardware where data is processed and stored but also that where it is

924 S. Subashini & V. Kavitha (note 119).

925 *Id.*

926 Lamia Youseff, Maria Butrico & Dilma Da Silva (note 838).

927 *Id.*

928 *Id.*

getting transmitted⁹²⁹. In the cloud reality data transmitted from the source to destination typically may pass through the resources of a large number of third parties⁹³⁰. Consequently, rules need to be established that will also prescribe for the minimum security benchmarks that these third parties will need to guarantee at all times, in a generic manner and without reference only to particular types of data processing. Regardless of the fact that heavy security measures are normally set up in the cloud, still information is transmitted through ordinary internet routes⁹³¹. It goes without saying, then, that there is also the need to establish rules that will necessitate from cloud infrastructure providers to seal their facilities against threats that may intrude to them from the world wide web. There are already several technical options that can help secure transmission of information inside the cloud⁹³². Encryption techniques tackle those needs to a certain degree yet they are not connection oriented⁹³³. Concerns with respect to interruption of the flow of information or even interception of it by outer non-clients of the network through the web need to additionally be considered. In a nutshell, security on the HaaS layer has both an internal and an external dimension and the principles regulating it need to make sure that every cloud environment will be not only internally secure and integral but also sealed and isolated towards the internet to also deter external security threats, such as cyber-criminal attacks.

ii. On the software/kernel layer

As it has already been described, on this layer we find the basic software tools needed for the management of the physical servers that compose the cloud network. The roles and duties appearing on this level are almost identical to those of the hardware layer, only focusing on the software as-

929 Mike P. Papazoglou & Willem-Jan van den Heuvel (note 853).

930 Ling Liu & M. Tamer Özsu (note 842).

931 Lamia Youseff, Maria Butrico & Dilma Da Silva (note 838).

932 Cong Wang, Qian Wang, Kui Ren & Wenjing Lou, *Privacy-Preserving Public Auditing for Data Storage Security in Cloud Computing*, in IEEE INFOCOM 2010 – IEEE Conference on Computer Communications, 1–9; Niels Fallenbeck & Claudia Eckert, *IT-Sicherheit und Cloud Computing*, in Handbuch Industrie 4.0 Bd.4, 137–171 (Birgit Vogel-Heuser, Thomas Bauernhansl & Michael ten Hompel eds., 2017.)

933 Ling Liu & M. Tamer Özsu (note 842).

pect and processes for keeping the cloud at its foundations sealed and integral from external threats⁹³⁴. Therefore, with regard to any regulatory principles stemming from this stack, reference is made to the recommendations analyzed in the previous section, i.e. the hardware layer.

iii. On the cloud software infrastructure layer

As it was analytically presented, fundamental resources to other higher-level layers are provided through this layer, which are then used to construct new cloud software environments or cloud applications. This is the first instance across the cloud network where we observe that the roles of provider and user of the resources of the network are so closely intertwined and enter each other's territory⁹³⁵. To put it more descriptively, a software infrastructure provider is at the same time a user of the network's resources, as he uses part of the network's hardware resources to host his processing activities that make the services it offers to entities of the above layers possible. In addition, a user of the software infrastructure providers' services is, at the same time, a provider of other cloud-based services addressed to end users of the network. This intertwining of roles brings to the forefront the need for cloud computing rules to be based on the teleological principle and follow, as much as possible, generic formulation patterns and depart from the case-based logic⁹³⁶. The blurred lines between roles and functions of actors on the cloud software infrastructure layer reveal the need to establish rules that will delineate duties, rights and obligations for entities across the cloud network without personalizing them or referring to specific arrangements/applications made possible thanks to the uses of the resources of that network.

934 S. Subashini & V. Kavitha (note 119).

935 Ozalp Babaoglu, M. Jelasity, Anne Marie Kermarrec, Alberto Montresor & Maarten van Steen (note 874).

936 William N. Eskridge & Philip P. Frickey, *Statutory Interpretation as Practical Reasoning*, 42 *Stanford Law Review* 321–384 (1990); Henry Prakken, *An exercise in formalising teleological case-based reasoning. Artificial Intelligence and Law*, 10 *Artificial Intelligence and Law* 113–133 (2002.)

iv. On the PaaS and SaaS layers

The last two layers of the ontology, corresponding to the cloud software environment and the cloud application layer, are the ones commonly referred to as PaaS⁹³⁷ and SaaS⁹³⁸. With regard to them, some observations regarding security issues and, respectively, rules that should be established to regulate them merit to be raised. On the PaaS layer, what actually happens on a technical level is that the administration supplier gives partial control to the customer in order for the latter to be able to manufacture applications on top of that layer⁹³⁹. However, for these applications to function properly and without interruptions, it is imperative that no insecurities beneath the software environment level occur. The cloud software environment layer is meant to empower cloud application designers to assemble their own particular applications on top of the platform⁹⁴⁰. Therefore, system durability and trustworthiness in relation to the underlying layers is of primary significance. Until now, this has been reflected on the affirmations suppliers bring on the table when negotiating contract services with potential customers⁹⁴¹. Till now, these clauses have been observed to be of great variety extending even to questionable security gimmicks in an effort, on behalf of suppliers, to enter into a kind of assurances fight towards potential customers, which have been found to extend to technical safeguards of doubtful trustworthiness⁹⁴². Therefore, rules establishing the minimum that should be achieved regarding these standards of safety are necessary. In fact, these rules are technically possible to be based even on objective technical measurements that will survey the viability of each cloud network's application security features at this level⁹⁴³. Some of those measurements with immediate application are defenselessness scores⁹⁴⁴ and patch scope⁹⁴⁵. These indices can show the quality of application coding based on the security features and the way the resources of each cloud network are brought together. One additional reason why secu-

937 See Chapter 2.

938 See Chapter 2.

939 Deepak Puthal, B.P.S. Sahoo, Sambit Mishra & Satyabrata Swain (note 837).

940 Dimitrios Zissis & Dimitrios Lekkas (note 869).

941 Lamia Youseff, Maria Butrico & Dilma Da Silva (note 838).

942 Nicholas Platten (note 42).

943 S. Subashini & V. Kavitha (note 119).

944 Stephanos Androutsellis-Theotokis & Diomidis Spinellis (note 861).

945 *Id.*

rity standards on the cloud software environment layer need to be reaffirmed via clear rules is that, especially when it comes to software-related vulnerabilities, such weaknesses or malicious elements of a cloud network can easily extend as far as the web applications that will be made available to users via the said cloud network, thus endangering or undermining the integrity of the wider web. Therefore, clear precautionary rules that will aim at containing them already on the PaaS layer are strongly advisable and urgently needed.

v. On the SaaS layer in particular

On the upper level of the cloud ontology, the cloud application layer or SaaS, lie the end cloud-based applications that are made available to end users with the only prerequisite that they have access to the internet, even partially, or that they can in any other technically feasible way access the cloud network where the service they make use of is hosted. The quintessence of affairs on this layer is that the client needs to be able to rely on the supplier to feel safety, in a whole range of different aspects⁹⁴⁶. Initially, it is elementary from the part of the supplier that he must actively prove that he can keep his clients from seeing or accessing without authorization one another's information. Simultaneously, it is imperative that the provider guarantees and makes sure that the application will be always accessible, not just because the other way around would put the client's confidence in the application in danger but also because, from a legal point of view, making sure that the application is always on and users can access it anytime they wish is a strong determinant towards the fact that the provider made sure users had unwavering and continuous possibility to exercise the expected functions through the application environment, among which also precautionary safety controls about their data. Thanks to SaaS, it is becoming increasingly possible (and popular) to switch to net program or software applications over 'old-fashioned', (usually) offline ones⁹⁴⁷. Consequently, primary focus is not so much on portability of uses, given that, after all, the new cloud-based apps usually do offer simpler and friendlier interfaces to users to do things. Rather, the focus lies nowadays

946 (note 852).

947 Xiaolong Jin & Jiming Liu (note 844).

on safeguarding or upgrading the security element in comparison to the standards offered in this front by the older applications and achieving effective information relocation and resource management while maintaining the elevated security standards as well⁹⁴⁸. What this model and its set targets practically mean is that in SaaS programming the service provider may host the application on its own private server farm or on a cloud computing facility administering it through a framework provided by an outsider supplier (e.g. Amazon, Google, etc.)⁹⁴⁹. This arrangement, where the involved actors and what each one of them is expected to carry out are so open, in terms of multitude, is one more pointer to the need of establishing rules on cloud regulation that will focus on the teleological principle, i.e. on who is expected to achieve what only ‘who’ should be understood in a generic sense (as a number of actors and not specific entities) and ‘what’ should be understood in the sense of functionality or body of functionalities within all those comprising the network and not as specific manifestations that come out when these functionalities are put to work. The cornerstone of the SaaS model, is that data is stored at the SaaS provider’s data center, along with the data of other users⁹⁵⁰. Even more, if the SaaS provider is depending on a public cloud computing service, users’ data might be stored on the same facilities along with the data of other unrelated SaaS applications. It is also quite a standard practice that the cloud supplier imitates the information at numerous locations across borders for reasons of keeping up the high accessibility prerequisite⁹⁵¹. Consequently, there are several security issues raised such as data security, network security, data locality, data integrity, data segregation, data access, authentication and authorization⁹⁵². Apart from any specialized rules that may establish specific standards or security policies as the necessary minimum, it is essential to take the leap and move from the specific to the broad context: technology and the constant evolution of science related to the cloud will make available more and more tools that will add up to the security levels of cloud networks⁹⁵³. As a result, it is not so imperative to legislate on which specific measures cloud networks should adopt to stand above the

948 S. Subashini & V. Kavitha (note 119).

949 John F. Sowa (note 839).

950 See Chapter 2.

951 See Chapter 2.

952 S. Subashini & V. Kavitha (note 119).

953 Dimitrios Zissis & Dimitrios Lekkas (note 869).

benchmark for network security but, rather, on what should be achieved in terms of security-related milestones. In other words, the SaaS layer of the cloud ontology and the way it is constructed exposes the need to develop a cloud regulation framework that will have at its core, not the ephemeral features of such a rapidly evolving phenomenon. Now that we have decomposed and exhaustively analyzed what the cloud actually is about, it is evident that cloud computing regulation should not be regarded as a new body of law that will replace existing legislation on particular manifestations of cloud computing, because of the latter being insufficient or dysfunctional. On the contrary, the proposed laws will come to serve as the currently missing cohesion element from the field of IT law, the one that will boost the integrity of this corpus of legislation as it will take advantage of its inherent features and will focus on its inherent flaws, at the same time, trying to correct them or, at least, seal cloud networks, as much as possible, against them.