Cloud Computing for the Public Sector and its Policy Implications

An introductory paper to the main issues and potential areas for government action

August 2016

AXON

Contents

Exhibits

Contact Information

Axon Partners Group Consulting	
Address	Calle José Ortega y Gasset 25, 1º 28006 Madrid Spain
Contact numbers	Tel. 0034 91 310 2895 Fax 0034 91 141 2811
Contact person	Dimitri Kallinis Partner
Email address	dimitri.kallinis@axonpartnersgroup.com

1. Introduction

Governments across the world are increasingly considering cloud computing solutions in an effort to provide less costly, more efficient and better public services through a variety of cloud deployment models (private, public, hybrid and/or community clouds).

A typical policy objective driving cloud adoption in the public sector is the improvement and streamlining of 'e-Government services' to meet the demands and expectations of citizens who are increasingly embracing digital and mobile technologies.[1] Better e-Government brings about immediately observable benefits to citizens, but poses the challenge of escalating computing and data processing needs. Cloud computing allows government services to address such increased computing and data processing needs in a flexible, scalable, integrated, cost-efficient and secure manner.

Despite this trend, however, it is still generally common for government agencies to be somewhat reluctant to shift their data to the cloud: erring on the side of caution is a typical and understandable public sector initial response to the challenges of a new, paradigm-shifting, technology.

This paper has been prepared by Axon Partners Group Consulting (Axon Consulting) to provide an introduction to the main policy-related issues concerning cloud computing and to the potential areas for government action. In the next sections this paper, first, provides a high level overview of cloud deployment models. This is followed by a discussion of some common concerns behind cloud adoption. The paper then recommends certain policy and regulatory steps that can help counter these concerns, in light of international experience.

Axon Consulting, established in 2006 in Madrid (Spain), provides consulting services to an international client base in the broad technology (TMT) sector. Within the TMT industry, Axon Consulting offers relevant expertise in the fields of strategy

[1] 'e-Government' is an internationally accepted term denoting public services offered to citizens and businesses over the Internet through the intensive use of information and communication technologies (ICT). Examples of such services include the digitization of government records, automation of tax services, procedures for receiving feedback from the citizens electronically, and various forms of data collection and processing.

& innovation, regulation & public policy, and profitability & operations. This includes supporting institutional bodies in developing and setting regulatory policy and competition monitoring; and advising telecom operators on business, commercial and regulatory strategy.

Axon Consulting thanks Microsoft Corporation for providing input and financial support in the preparation of this paper.

2. Basic guide to cloud services and deployment models

The concept of 'Cloud Computing' has emerged over recent years as a genuine paradigm shift for the use of IT services by organizations and individuals. Solid demand and supply have converged on an understanding of what Cloud Computing implies, which is now widely accepted in the industry. The commonly sought characteristics that distinguish a cloud computing service from traditional IT services consist of the following service provision aspects:

- availability of services on-demand
- access to services over a network connection
- utilization of pooled computer resources
- rapid provisioning (and de-provisioning) times with minimal interaction (i.e. formalities)
- metering capability to automatically control and optimize resource usage.

The above characteristics, commonly referred to as the five essential characteristics of cloud computing, are further explained below:

1. On-demand self-service: Consumers have the ability to provision computing capabilities, such as server time and network storage, as needed without requiring human interaction with the service provider.

2. Broad network access: Capabilities are accessed over the network through standard mechanisms that promote use by different platforms such as mobile phones, tablets, laptops and workstations.

3. Resource pooling: The provider's computing resources are pooled to serve multiple consumers, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand.

4. Rapid elasticity: Capabilities can be elastically provisioned and released, in some cases automatically, to scale rapidly outward and inward, commensurate with demand.

5. Measured service: Cloud systems automatically control and optimize resource use. Resource usage can be monitored, controlled, and reported, providing transparency for both the provider and consumer of the utilized service.

These essential characteristics have changed the way people and organizations view IT, as they emphasize monitoring of on-demand service usage and service access and consequently highlight the costs of IT usage. Terms that were previously considered exclusively by technical IT teams such as *number of instances run on a server* or *types of storage* (temporary or permanent), or *required number of CPUs* to run multi-threaded applications have gained visibility across end-users from multiple functions due to the pay-as-you-go nature of cloud services. Most direct users of cloud services today are well acquainted with the concept of a datacenter '**workload**', which encompasses, among others, a combination of the required *instances*, *number of CPUs*, *type of storage*, as well as the *connectivity (e.g. throughput, latency*). Together with the workload (as the commonly accepted measure of cloud services), '**capacity**' of a datacenter to deliver the demands of workloads has also become an important concept for cloud users in assessing the size of a cloud provider. While both terms have long existed in the IT literature, due to the paradigm shift presented by cloud computing they have gained new meanings specific to the cloud.

2.1. The Cloud Deployment models (transformation from Traditional IT to Cloud IT)

2.1.1. The Traditional IT model

Prior to the cloud computing service model, centralized datacenters hosted on-site offered the main IT deployment model for the more sophisticated needs of private and public organizations. In this traditional IT model, every organization owns and manages a datacenter, which is designed with a capacity and computing power deemed to respond to the organization's expected maximum computing requirements.

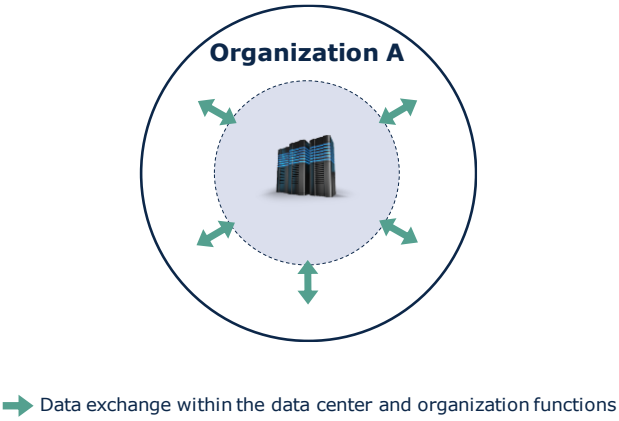

Exhibit 2:1: Visual representation of the traditional IT model [Source: Axon Consulting]

In the traditional IT model, organizations increase their IT capabilities through IT equipment and software procurement cycles, which typically tend to be lengthy and may lag behind the actual computing needs within the organization or, conversely, overestimate them.

2.1.2. **Cloud deployment models**

As a paradigm shift away from the traditional IT model, cloud computing offers five deployment models, commonly referred to as **private**, **public**, **hybrid**, **community** and **community hybrid** cloud. We describe each of these models illustratively in the next paragraphs.

Private clouds

The private cloud deployment model can be considered as an incremental change over the traditional IT model. It preserves the concept of hosting an on-site datacenter within an organization. However, the datacenter capabilities are provided to the organization's other departments through the five essential characteristics of cloud computing mentioned previously.[2]

[2] i.e. on-demand self-service, broad network access, resource pooling, rapid elasticity and measured service.

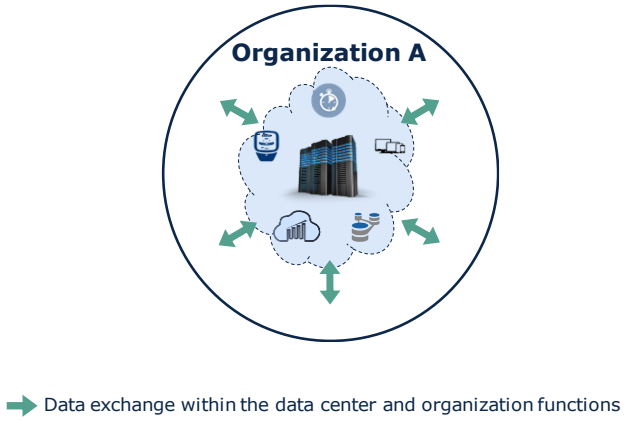

Exhibit 2:2: Visual representation of private cloud deployment [Source: Axon Consulting]

Through the introduction of the concepts of resource pooling and rapid provisioning, the IT department within an organization basically shifts from being a cost center incurring capital expenditures in infrastructure to a profit center that bills the use of IT resources. Such billing takes place through a metered, pay-per-use scheme. The organization's different departments have access to its IT resources through a user-friendly and easy interface, which allows them to activate and de-activate the computing or software resources (e.g. analytical computing) they require.

Public (or hyperscale) clouds

A public cloud comprises one or more datacenters owned and operated by a third party, i.e. a cloud service provider (CSP). In this case, an organization uses the computing and storage resources made available by the CSP, which also offer rapid access to affordable computing resources to other organizations or individuals.

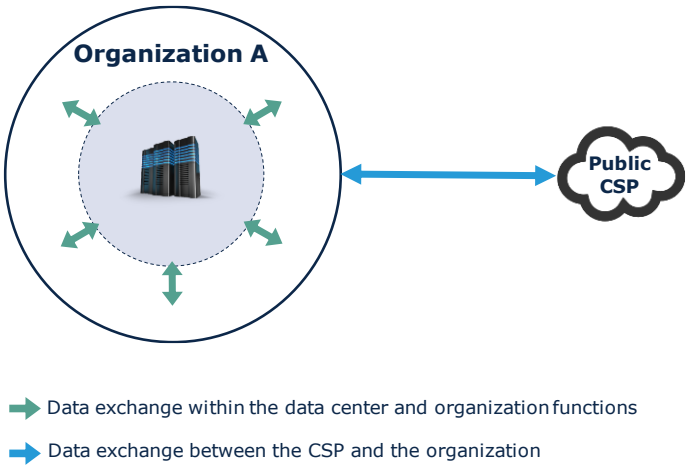

Exhibit 2:3: Visual representation of public cloud deployment [Source: Axon Consulting]

With the shift to a public cloud, relevant processes that were previously covered via traditional services and capabilities of the IT department are relocated to an off-site data center, which is professionally managed. The IT department serves as a reduced interface between the organization and the public CSP.

As public cloud deployments by organizations increase in terms of computing capacity to meet substantially large amounts of workloads, the concept of **hyperscale cloud** deployment has been emerging over the recent years. The exhibit below demonstrates this phenomenon through actual and forecast statistics by Cisco on the global number of workloads demanded from public CSPs.

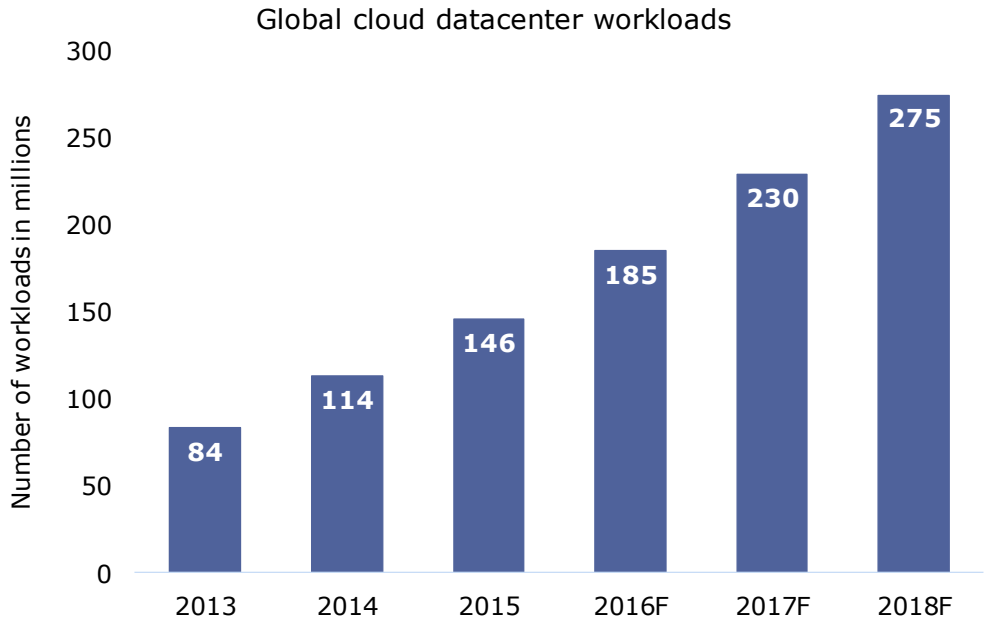

Exhibit 2:4: Past and forecasted values of global cloud computing datacenter workloads in millions [Source: Cisco Global Cloud Index 2015]

The observed increase in global demand for workloads can be attributed to new ways of using computing resources, such as big data analytics, instantaneous data sharing and processing between connected equipment (i.e. Internet of Things), and the overall increase in the amount of data that needs to be stored. All these take public clouds to the next level of hyperscale cloud deployment.

Hybrid clouds

The term hybrid cloud refers to the use of a private cloud in combination with a public cloud. Many companies with private clouds gradually end up distributing

workloads across data centers, and private clouds and public clouds thus end up creating hybrid clouds.

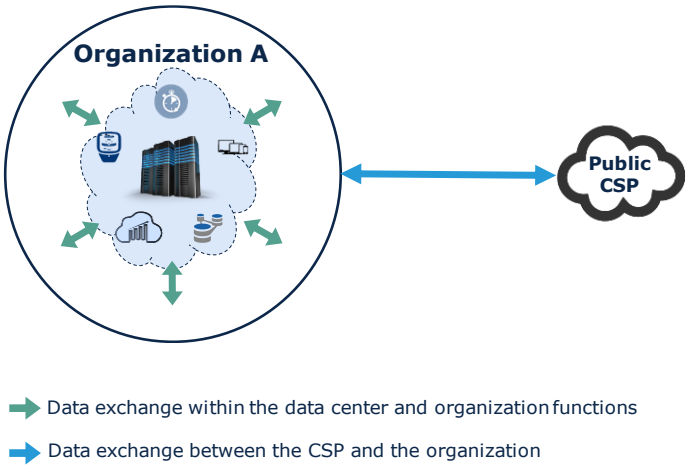

Exhibit 2:5: Visual representation of hybrid cloud deployment [Source: Axon Consulting]

A hybrid cloud combines the benefits associated with a private and a public cloud. Some IT capabilities remain within the organization while others are obtained from an off-site data center. The main challenge associated with this deployment model is to assess the types of data or applications suitable for each of the private and public clouds that comprise the hybrid cloud.

Community clouds

Community cloud deployment models are generally encountered within clusters of public (i.e. state) institutions that aim to benefit from cloud computing, such as a group of academic or research centers. As such, a community cloud is a private cloud installed at a public institution which becomes a central cloud computing facility for several other public institutions. In essence, the organization deploying the private cloud acts as an internal CSP to the other organizations.

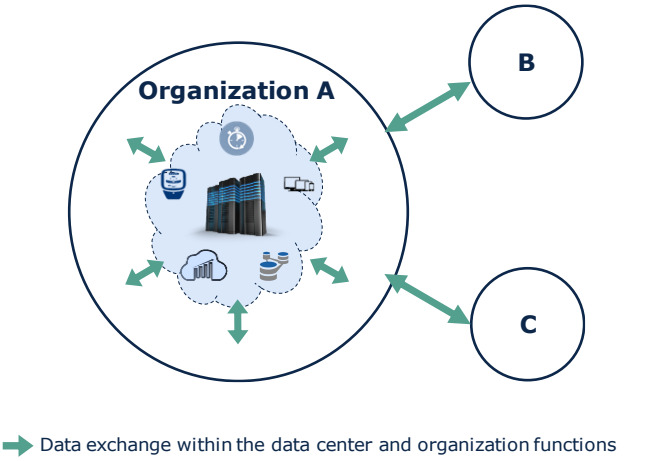

Exhibit 2:6: Visual representation of community cloud deployment [Source: Axon Consulting]

This deployment model builds on the private cloud described earlier, with the central IT department billing the use of IT resources to individual departments through a metered, pay-per-use scheme across several public organizations.

Community hybrid clouds

A community hybrid cloud – typically deployed within clusters of public institutions – merges the community cloud with the public cloud. Once again, an organization acts as an internal CSP among the public institutions and in addition to this cloud capability, all of the public institutions have the possibility to also leverage the services offered by an external CSP.

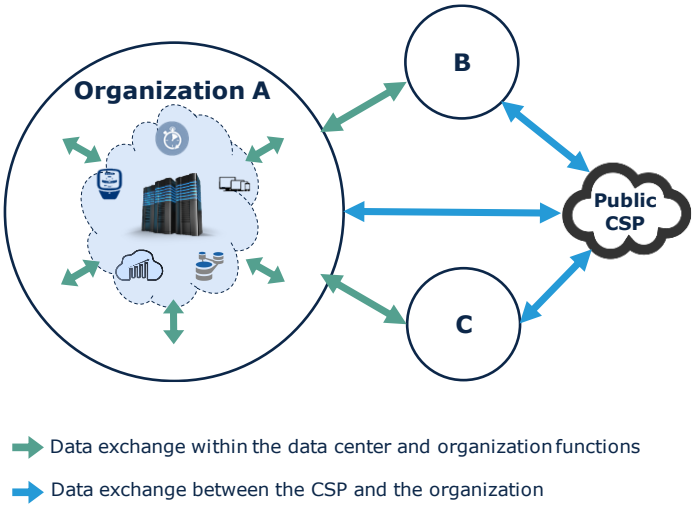

Exhibit 2:7: Visual representation of community hybrid cloud deployment [Source: Axon Consulting]

3. Main public sector concerns regarding a move to the cloud

3.1. Information security and data protection are the top concerns

Cloud computing is a market and technology-driven, rapidly evolving, international industry. Such characteristics may be particularly appealing to the more technology-friendly and agile segments of the private sector, but are less likely to inspire adoption by the traditionally more cautious government and administrative services.[3] And yet, government services are one of the main potential beneficiaries of cloud services, especially in countries with a high level of state involvement in the economy.

Public sector concerns regarding cloud solutions may vary somewhat across countries, but our research suggests that the two most commonly quoted sources of concern are **information security** and **data protection**.[4]

These are two distinct concepts, despite some common overlaps in their interpretation:

- *Information security* refers to the protection of information systems against **unauthorized access**, use, disclosure, disruption, modification or destruction, primarily **by third parties**. In other words, it relates to a cloud service provider's **obligation to protect** its cloud system and its data generally.
- *Data protection* refers to the **protection** of, first and foremost**, the personal data** or other information **of individuals**. Put simply, it can be interpreted as the cloud service provider's **obligation to not share** user data with others.

[3] For example, in a February 2015 report on a Security Framework for Government Clouds, ENISA (the European Union Agency for Network and Information Security) concluded that "*the state of deployment of Governmental Cloud computing is in general at a very early stage*."

[4] This is also confirmed by the ENISA report, which concludes that "*security and privacy issues are considered as key factors to take into account for migration, and at the same time are the main barriers for adoption*".

Detractors of cloud computing invoking information security and data protection risks can sometimes point to real-life major security incidents, such as a security breach experienced in May 2015 by the US Internal Revenue Services (IRS) or the leak of the entire voter database of the Philippine Elections Commission, in March 2016.

A first remark concerning such incidents is that they remain rare and are by no means specific to the cloud: they demonstrate risks associated with all types of data processing, be it in traditional (in-house) IT systems or in the cloud. Indeed, there is a potentially misleading perception that cloud computing may pose an increased risk due to the absence of direct supervision by the public entity that owns the data. This reaction confuses the level of data security with the identity of the data controller: the level of security will depend, among other, on the security standards applied by the data controller, whoever this may be. There is nothing suggesting that control by the original creator or owner of data will be more secure than control by a CSP whose whole commercial proposition actually consists in providing the most secure data processing environment possible to third parties. In a recent commentary, the Federal Chief Information Officer of the United States also confirmed this point by likening the CSPs to banks in the way that CSPs possess the skills, the abilities and the motivation to ensure the provision of best peace-of-mind practices in the field of cloud services compared to any one company or organization[5].

In fact, it is safe to assume that huge volumes of data are lost, destroyed, corrupted or stolen, on a daily basis, albeit on a smaller scale, at the level of private corporate IT systems, including those run internally by various public authorities. However, such breaches, if and when they occur, are not immediately (if at all) found and addressed, and are rarely reported, unless the user concerned has a regulatory obligation to do so.

3.2. **Data localization**

A step taken by governments reportedly to protect government and other sensitive data is to impose **data localization** requirements, i.e., an obligation to keep data uploaded to cloud servers within the national territory and thus under the direct jurisdiction of the local government authorities. To mention a few examples:

[5] http://www.cio.com/article/2996268/cloud-computing/us-cio-tells-it-leaders-to-trust-the-cloud.html

- Recently, PayPal's license to operate in Turkey was not renewed, because the datacenters it uses to store its customers' financial information are located outside Turkey and thus represent a data protection risk, as well as a violation of Turkey's banking laws.
- Since July 2015, public authorities in Germany concluding an agreement with a private CSP are required to ensure it holds an ISO 27001 or equivalent security certification and does not host their data outside the country.
- New Zealand's Ministry of Health is developing its Health Information and Governance Framework to restrict reliance on any cloud service providers that store the data outside the country until after they are in compliance with a very stringent set of rules, to address information security and data protection concerns.
- Russia's privacy legislation requires data localization for the data of Russian citizens.
- China's proposed cybersecurity provisions will reportedly impose security reviews on any operators of crucial information infrastructure facilities who wish to store data overseas.

Such data localization requirements are typically justified as an information security and data protection measure. However, information security and data protection are not necessarily a function of a data center's physical location. If data are accessible through the Internet, they are accessible from anywhere, and security and privacy breaches may also originate from the other end of the world. Moreover, policy-makers often ignore the cost of regulation. In this case, obligatory uses of local cloud services due to regulatory requirements are almost certain (unless subsidized by the government) to cost more than commercially available hyperscale cloud services. For certain workloads, such as the hosting of public-facing websites of government agencies, a relaxation of the regulatory requirements on data location would definitely bring advantages in terms of the cost and the reliability of the provided hosting services.

From a user's perspective, the real factors ensuring adequate information security and data protection are a function of three sets of parameters:

- The security measures in place by the CSP, regardless of the system's physical location(s).
- The jurisdiction(s) governing privacy and security issues in relation to the data concerned, and the public authorities' powers to access such data. Again, when it comes to the cloud, this is not necessarily determined by the physical location of the data (which can vary or be multiple, in any event). The US authorities'

recent attempts to assert extraterritorial jurisdiction over a Hotmail account hosted in Ireland, which was contested by Microsoft before the US courts, is a prominent example of this type of problem. More recently, the 2nd Circuit Court of Appeals has concluded in Microsoft's favor on the grounds that the US Stored Communications Act (SCA) does not mention extraterritorial application, i.e. application for non-US cases. However, the court also remarked that if there was any evidence of the account holder being a US citizen or a US resident, the warrant could have been valid. The opinion also calls for the modernization of the SCA, which could potentially be revised to allow the assertion of extraterritorial jurisdiction by US authorities in the future.

- The legal rights and responsibilities of CSPs and users alike, and the applicable rules on liability, responsibility and dispute resolution in the event of a security or data breach.

In the remainder of this document, we shall examine some of the regulatory tools available to governments to address these issues.

4. Potential cloud policy approaches for governments

4.1. The pros and cons of regulation

Regulation is not an end in itself; but legal certainty is. Legal certainty can be partly or fully based on legislation, state regulation or other statutory rules, but it can also rely on commercial contracts, if these are sufficiently clear, comprehensive, enforceable and not overridden by law.

Cloud computing has grown rapidly without cloud-specific laws but a mosaic of existing rules on data protection, privacy, information security, laws of contract, intellectual property etc. whose provisions can be interpreted, where necessary, in a cloud context. The result varies across jurisdictions: those with a significant body of established, transparent rules, case-law and experience in ICT matters may be better off without new, cloud-specific rules, but those without, e.g., an existing regime on data protection or contract rules for cloud-type services, will need to find ways to fill these gaps to establish legal certainty.

Hence governments considering any form of regulatory action on cloud computing need to strike a careful balance between:

- ensuring that legacy regulations are brought forward in a way that makes sense,
- avoiding a potentially heavy-handed and possibly obsolete regulation that is specific to the cloud, and
- filling in existing real legal gaps in a way that will help CSPs and cloud users (also including government services), and address, in particular, information security and data privacy risks associated with cloud computing.

Similarly, from a policy viewpoint, state measures on cloud computing can signify either an unnecessarily interventionist approach with a chilling effect on the local take up of cloud solutions or a welcome public commitment of the state to the cloud as the way forward – with a positive ripple effect across both the public and the private sectors. One dynamic that should be avoided is requiring more rigorous security and privacy requirements for cloud services than what has been, and is, required for traditional on premise solutions. Ideally, policy defines relevant

security and privacy requirements for ICT solutions generally – whether offered from the cloud or a traditional server.

Clearly, therefore, even if there is no one-size-fits-all solution for the regulation of the cloud and its use by government services, there is at least a list of key issues associated with information security and data privacy that can benefit greatly from legal certainty. We propose to elaborate on these issues and the recommended regulatory or legal solutions, with examples from international practice.

4.2. Regulatory measures for better information security

4.2.1. Classification of data and information based on their origin, content and level of sensitivity

Establishing how cloud computing is going to work across various government services requires, first of all, rules identifying which data and applications are "cloud-eligible" and the level of protective measures they require. Therefore, a first critical regulatory task for better information security is **data classification**, through which government agencies can determine and assign different levels of security to different categories of data.

For example, Australia, United Kingdom and the US[6] have recently gone through separate processes of data classification. While each government has decided on an independent data classification regime, the results are strikingly similar. The common aspects of all three regimes, illustrated in Exhibit 4:1 are:

- Three levels of data according to sensitivity: i) public/non-sensitive, ii) sensitive and iii) secret/critical data.
- Most of the data by volume belongs to the public/non-sensitive data category.

[6] Note that, in the Exhibit that follows, FISMA stands for Federal Information Security Management Act.

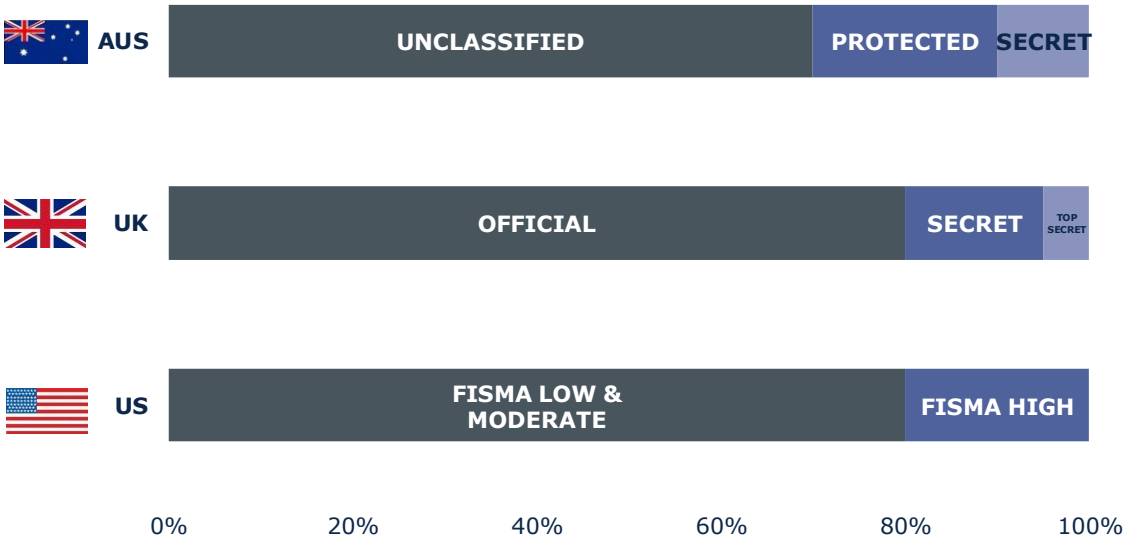

Exhibit 4:1: Illustration of international data classification schemes [Source: Axon Consulting and Microsoft]

Most data classification regimes are designed to be simple, so as to clearly sort out what needs protection and what not, and define the appropriate level of protection by reference to a deployment model and class-specific standards of security. In many cases, most of the available data do not classify as sensitive, and are therefore suitable for processing in a public cloud. The table below provides some examples of such services with possible appropriate cloud deployment models, which have been previously described in Section 2:

CATEGORY	SERVICE EXAMPLES	POSSIBLE CLOUD DEPLOYMENT MODEL
Collaboration	▶ Agency-wide email to the cloud ▶ Office productivity tools	**Public Cloud**
Workflow	▶ Employment verification ▶ Grants management ▶ Claims processing ▶ Customer relationship management	**Hybrid Cloud**
Infrastructure	▶ Public-facing websites hosted in the cloud ▶ Application development/testing ▶ Content delivery ▶ Virtualized data centers	**Public Cloud**
Business Intelligence	▶ Data analytics ▶ Performance management	**Hybrid Cloud**
Information Security	▶ Identity management ▶ Security management services	**Private Cloud**

Exhibit 4:2: Candidate areas for cloud assessment and their corresponding possible deployment models [Source: Axon Consulting]

The differences between these various deployment models have been discussed in Section 2, and the right choice will obviously depend on a number of factors. As a rough rule of thumb, government agencies and public authorities in general seem to have a preference for a hybrid deployment model, as it allows them to take advantage of several levels of information security and data protection. These range from the more secure private cloud for the agency's or authority's internal needs to the more flexible, ubiquitous and cost-effective public cloud, with high processing power suitable for the large amount of non-sensitive data processed by cloud service providers on behalf of government services.

In all cases, government services wishing to make use of the cloud should remain free to define different classes of data for information security purposes and classify their data accordingly, based on their own policy and regulatory constraints. This is and should remain a largely discretionary and internal process, and regulators can, at most, provide some guidance on the type of classification that seems appropriate for different types of government data.

One area that policy-makers should watch out for in this regard are the incentives for individual agencies or entities to over-classify their data – because decision-makers tend to err on the side of caution but also, in some cases, because of the importance or value that they place on their work and corresponding data. In any event, the same data classification systems cannot easily apply across all public services. For example, it is clear that the level of information security required,

e.g., by the ministries of defense or health, will normally be very different from that required by the ministry of culture.

It may, therefore, be relatively easy for a government to adopt a "top down" data classification for information security and cloud computing purposes at the level of one or more ministries; it will be more difficult– but certainly not uncommon - to do so for the whole government sector; and a combination between general ("umbrella") rules for all state agencies with separate, more detailed or agency-specific information security rules for one or more individual agencies is also possible.

As an example of the last, mixed, model, in the US, the Federal Information Security Management Act (FISMA) requires agencies to maintain an information security program commensurate with their risk profile. The FISMA accreditation of agency systems has three primary objectives, namely confidentiality, integrity and availability, and it distinguishes between low, moderate and high security standards that need to be met, depending on the data an agency is processing.

Although it is also covered by FISMA annual compliance reports among other agencies, the US Department of Defense has adopted its own cloud security strategy, given its particular situation. This is set out in a "Security Requirements Guide" with a Cloud Security Model that initially defined six information "Impact Levels" ranging from "Unclassified Information approved for Public release" (Level 1) to "Classified Information up to SECRET" (Level 6). In order to simplify the selection process, the number of levels was then reduced from 6 to 4.[7]

4.2.2. **Definition of authorization levels and applicable security standards for cloud service providers**

Once different categories of data sensitivity have been defined for cloud computing purposes, the next logical step for information security is to lay down corresponding certification or authorization levels that establish the specific categories of data that may be handled by each particular CSP. Doing so ensures that CSPs that handle data with higher sensitivity are subject to more rigorous information security standards, designed for their authorization level.

[7] For more information, see http://iasecontent.disa.mil/cloud/Downloads/Cloud_Computing_SRG_v1r2.pdf.

Typically, data classification and the definition of specific security standards per data category go hand-in-hand, set out in one and the same measure, with security standards then evolving over time.

An example of the link between different information security levels and corresponding security standards is the "Multi-Tier Cloud Security (MTCS) Singapore Standard (SS)", presented as "the world's first cloud security standard that covers multiple tiers", with three levels of security for both the private and the public sector. The main purpose of this standard is to define three distinct and comprehensive sets of security standards, each suitable for a different level of security. CSPs certified under this scheme can specify the levels of security they can offer to their users. In parallel, potential users wishing to rely on such CSP services can use the MTCS SS to understand and assess the cloud security standards and procedures that they require.

A potential risk of any such top-down government data classification is that it may effectively result in a protectionist regime, which complicates or even excludes the provision of cloud services by non-domestic CSPs. This may be legally incompatible with a country's international obligations but also constitute a wrong policy move, especially if the local pool of CSPs is not sufficiently experienced and credible, and cannot meet international high standards of quality and security. Put simply, if a government is ready to rely on foreign suppliers of state-of-the-art weaponry to protect its country's national sovereignty and security, it should be also ready to rely on the best state-of-the-art cloud solutions to minimize security and privacy risks, regardless of the CSP's nationality. Similarly, an overly-rigorous, custom approach may favor large multi-national CSPs who can invest in complying with such requirements while excluding smaller, local providers.

A way to avoid an unnecessarily protectionist approach and enhance competition is to rely on common or reputable international information security cloud standards and authentication procedures, at least as an alternative to domestic standards, if any. Examples include the ISO/IEC 27001 standard for information security controls, the Cloud Security Alliance (CSA) Cloud Controls Matrix (CCM) or the Payment Card Industry Data Security Standards Level 1 version 3.1 certification standard for organizations that accept most payment cards and process cardholder data.

These standards are generally not mandatory; but a government, its regulators and public agencies are free to require them for specific categories of government data or cloud services, be it as a binding requirement, one among several presumptions

of compliance with the required security standards or as an optional add-on that weighs in favor of a CSP's selection.

4.2.3. Liabilities and notification obligations regarding security breaches or information leaks

In the context of information security, a **security breach** refers to the intentional overriding or disabling of any security measures implemented in a datacenter or other ICT system that aims to protect information from unauthorized access. An **information leak** consists of revealing private information inadvertently, as a result of poor or lacking implementation of security measures.

While exposure to these two risks can never be excluded entirely, defining the CSPs' civil, administrative and criminal liabilities to their clients in such incidents and their obligations, if any, of notifying the affected parties and the supervisory authorities can help limit the associated risks to a commercially acceptable level.

The required provisions can be part of a country's data protection rules, which, although relating to personal data only and not to any type of user content, can perhaps be relied on to protect users against security breaches or information leaks of other types of user content stored in the cloud. For greater certainty, however, government authorities may need to set out special rules either in their procurement of cloud services or at a more general, legislative level (potentially also expanding the benefit of these rules to private users of cloud services).

4.3. Regulatory measures for data protection

4.3.1. Establishing quality standards on the collection, security, fair use and processing of collected data

Most countries have some form of sector-specific or general data protection rules, governing the collection, use, or other processing of personal data and personal information.

The European Union (EU) has arguably the world's strictest data protection regime, introduced through a Directive in 1995 (which still allows for slightly different national rules), which will be replaced by an even more comprehensive Regulation by 2018. The new regime will largely harmonize national provisions and expand their substantive and territorial reach, with implications outside the EU. Although

references to cloud computing are conspicuously absent from the Regulation's text, it is clear that its provisions will govern the whole range of data protection issues arising in a cloud content, and may well have an inherently controversial spillover effect outside Europe.

The merits of this new data protection regime still need to be tested in practice and since the level of the protection it affords to personal data is very high and involves complex enforcement provisions governments outside the EU are unlikely to emulate it entirely – nor would it seem a good idea to do so in radically different jurisdictional and market environments. They may, however, use the Regulation as a "maximalist check list" from which they can cherry-pick specific provisions, such as those dealing with the required standards for the collection, security, fair use and processing of collected data.

However, there may be other, less intrusive and more cloud-specific, ways for governments across the world to draw inspiration from the provisions of the EU Regulation or, possibly, other developed data protection regimes. A good example is the ISO/IEC 27018 code of practice "for protection of personally identifiable information (PII) in public clouds acting as PII processors", an addendum to ISO/IEC 27001, which, although inspired by EU data protection rules, is more international and less controversial and coercive that the full set of EU rules. It should be therefore relatively easier for government agencies to request that CSPs wishing to offer them services demonstrate their compliance with this code of practice.

4.3.2. **Jurisdiction in data protection matters and cross-border data transfers**

As a ubiquitous platform, the cloud raises novel issues of territoriality and jurisdiction. From a data protection perspective, these have led to conflicts in more than one ways:

- As previously mentioned, a US court has attempted to assert jurisdiction over user data stored outside the US territory. However, disclosure of such data to the US authorities may be inconsistent with the hosting country's data protection rules. The most high-profile example is the EU, which protects personal data regardless of the data owner's nationality and sets very strict conditions for their transfer outside the EU. CSPs caught in the middle of such a conflict of laws have no easy way out: they cannot keep both sides happy. A possible way around this conundrum, at least so far, is the solution of a "Data

Trustee", as agreed between Microsoft and Deutsche Telekom in Germany, in November 2015. Under this arrangement, Microsoft offers Azure, Office 365 and Dynamics CRM Online from two datacenters in Germany. Customer data are stored in Germany only, and a Deutsche Telekom unit acts as the "Data Trustee" for the data of Microsoft's customers, thus bringing such data outside the territorial reach of US or other government authorities other than those of Germany.

- The scope of the EU Directive on data protection has an extraterritorial element, in that its rules can potentially also apply to companies processing data outside the EU, with only a tenuous territorial link to the EU, as was confirmed by the EU Court of Justice in the *Google Spain* case. The EU is not alone in pushing for such jurisdictional overreach in the data protection sector. For example:
  - Under the Privacy Amendment Act of 2012, an organization "carrying on business" in Australia must comply with the Australian Privacy Principles, even if it has no physical presence in Australia.
  - In Canada, the Privacy Commissioner has jurisdiction to investigate foreign entities compliance with the requirements concerning the collection and processing of personal information of Canadian residents.
- As mentioned earlier, the transfer of personal data outside the EU is subject to strict conditions. In principle, it is allowed only if the non-EU country ensures an "adequate level of protection". This is subject to certain derogations (such as the data subject's unambiguous consent to the proposed transfer). Under the 1995 Directive, the European Commission may find that a third country does ensure such an adequate level of protection and in 2000, the Commission adopted the so-called "Safe Harbor Decision" finding that the US met this adequate level of protection standard, albeit only with regard to those US companies that chose to adhere to the "Safe Harbor" scheme. More than 4000 US companies relied on this decision to operate in Europe, moving data back and forth. However, in October 2015, the EU Court of Justice declared the Safe Harbor Decision invalid. It has been replaced by a new arrangement (the "EU-U.S. Privacy Shield"), as of 12 July 2016, which many CSPs have accepted in order to comply with EU rules.

The takeaway from these examples is that governments wishing to make use of cloud services should be aware of the risk of such jurisdictional conflicts and take appropriate measures to avoid them. At a minimum, they should insist on clear choice of law and jurisdiction rules. They should also bear in mind that, even though jurisdictional rules in the cloud space can never be absolutely airtight, most

governments will be reluctant to step directly, through legal means, into each other's "jurisdictional cyberspace", when it comes to government data.

Further, government users of cloud services should have a clear view of the personal data they expect to store on the cloud, their level of sensitivity and a realistic assessment of the implications of their transfer to a physical location outside the country – for the purposes of an informed response rather than a flat-out rejection. Finally, from a user's perspective, the simultaneous application of more than one set of national data protection rules is not necessarily a problem: on the contrary, it may mean that the user can pick and choose those friendliest to the data owner.

4.4. **Implementing Data Classification**

As previously mentioned in Section 4.2.1, data classification is one of the steps needed to ensure better information security. With this in mind, the UK, US, Australia and many other countries have strived to differentiate data based on their required level of protection and security standards.

For policymakers interested in implementing a similar classification regime it may be tempting to classify most public data as 'top secret', 'secret' or at least 'classified'. However, there are practical and economic downsides to such an over-classification, which should be considered upfront. The cloud deployment models described in Section 2.1.2 have different associated costs that are a function of the complexity of deployment, any applicable security measures in place and whether or not economies of scale of a large public datacenter are being achieved. In the exhibit below we provide an illustration of costs associated with each deployment model:

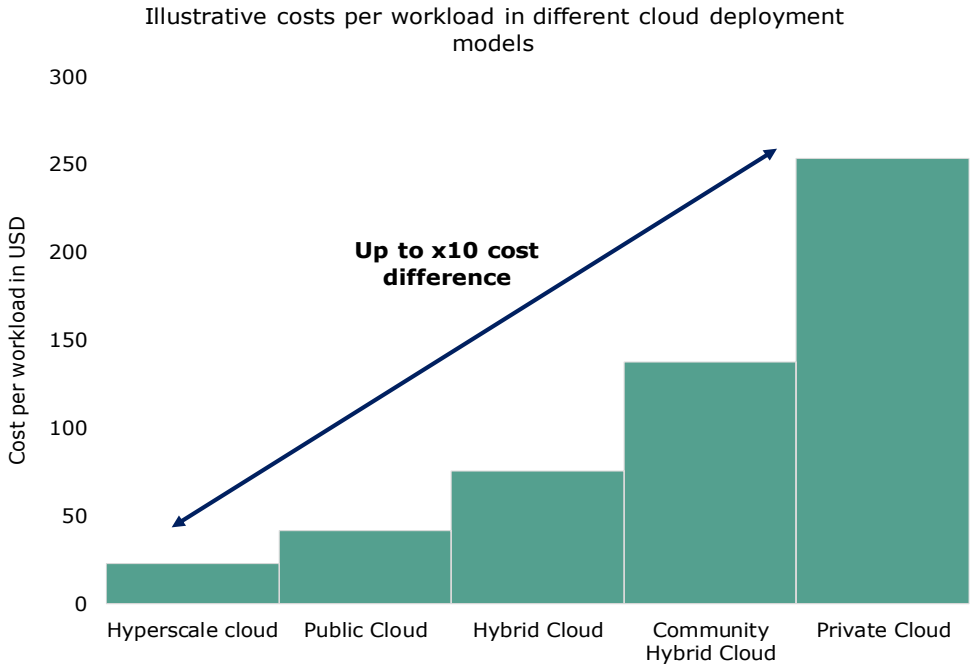

Exhibit 4:3: Illustrative costs of different cloud deployment models [Source: Axon Consulting]

As can be observed above, the difference in cost per workload could be as much as 10x depending on the cloud deployment model. Keeping such material cost differences in mind, policymakers may want to benefit from less costly public cloud deployment models as much as possible by limiting reliance on a private cloud to highly classified data. In such a balancing act, hybrid cloud models often represent a good compromise between cost and security, ensuring that the most cost-effective cloud service is used for any given workload.

To be able to utilize a hybrid cloud deployment model more effectively, it is important to classify data upfront in a way that makes clear which classification is suitable for which data category. Over-classification of data is a potential risk that should be avoided as much as practicable, as it creates an unnecessary burden on the decision making process and generates dispensable costs.

5. The challenges of designing a regulatory regime for the cloud

5.1. Cloud rules vs. general rules

At least on its face, designing a regulatory regime to specifically address cloud computing seems like the most direct way to address public sector concerns associated with moving to the cloud. Such an approach has the advantage of being tailor-made for the cloud computing industry and thus the most obvious way to remove any ambiguity on the applicable rules.

On the flip side, however, such an approach has hardly any international precedents to rely on and no guarantee of success. Although cloud-specific rules have certain merits, cloud computing does not sit well within a system of detailed and inflexible national legislative and mandatory rules. Therefore, a basic principle guiding any attempt to regulate the cloud should be that the more detailed the rules, the less mandatory they should be.

Before attempting to regulate the cloud, governments should have a clear understanding and consensus on the objectives and needs they want to address:

- Is there a real need for specific measures and if so, why? What are the main gaps or ambiguities under the existing legal regime that need to be addressed?
- Which are the priority legal areas that need to be covered? Other than information security and data protection, are the existing rules on law of contract, consumer protection, intellectual property, unlawful content and standards sufficiently clear, applicable to the cloud and appropriate on substance to address possibly novel cloud-specific issues?
- Is the policy objective the creation of a more open cloud market or does the government want to protect local CSPs to the possible detriment of quality and take-up of cloud services?
- Can the public authorities afford the time, budget and effort involved in developing national standards and authentication procedures?
- What are the tax or other incentives that can be offered to CSPs to attract their investment in locally hosted data centers? etc.

As mentioned at the outset, information security and data protection seem to top the list of areas calling for legal certainty and, where this does not already exist, new, cloud-specific, regulation. An example of assessing the need for regulation in these and other areas is shown in Exhibit 5:1, below:

	Are there relevant laws?	Aligned with Cloud computing needs?	Need for alignment
Data protection	***Yes***	***No***	
Information security	***No***	***-***	
Intellectual property infringment	***Yes***	***No***	
Unlawful contents	***No***	***-***	
Certification and standardization	***No***	***-***	
Consumer protection	***Yes***	***No***	

Exhibit 5:1: Example assessment to determine the necessary regulatory regime for the cloud [Source: Axon Consulting]

5.2. **Regulatory responsibility for the cloud**

Cloud computing raises a very mixed range of legal issues and its regulation may therefore involve a number of different government agencies. A typical example of the overlapping competence (and hence required coordination) involved in regulating the cloud is shown in the Exhibit below:

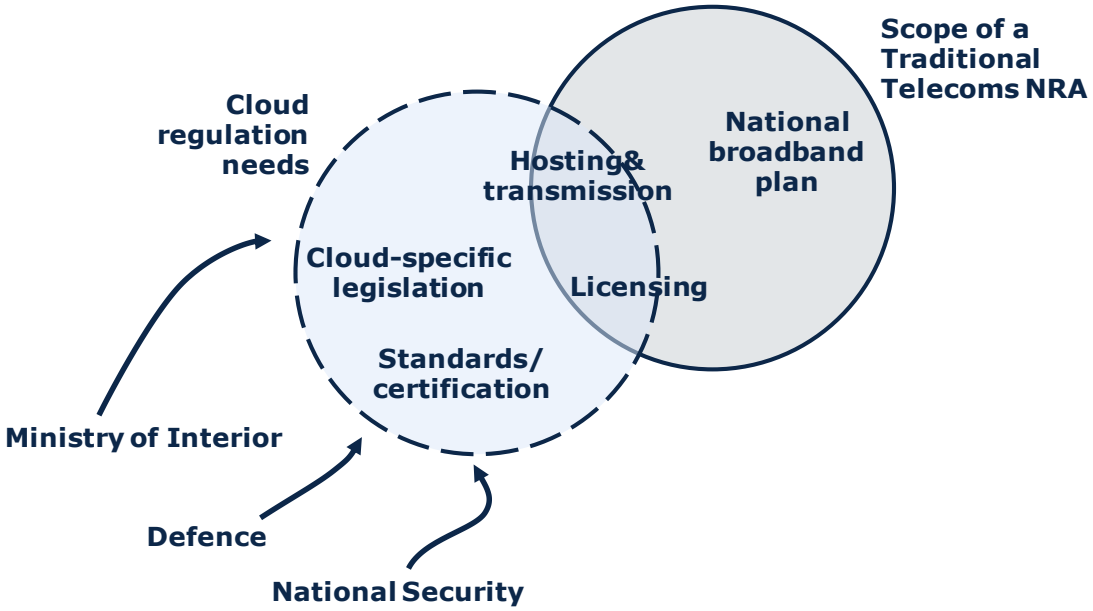

Exhibit 5:2: Cloud computing regulation requires regulatory oversight from various public agencies [Source: Axon Consulting]

Effective coordination between government agencies, without time-consuming and counterproductive overlaps and "turf-wars", is one of the possible pitfalls than need to be addressed early on by any government wishing to regulate the cloud in a timely and effective manner. This dynamic is particularly challenging in verticals such as healthcare, which is often regulated by multiple entities. Ideally, overall responsibility should be entrusted to a single ministry (e.g., the one in charge of telecoms or technology) or regulatory agency (e.g., the national telecoms or ICT regulator) to provide high-level guidance across government and key industries. Parallel, overlapping and even conflicting initiatives by several agencies and authorities, simply because the cloud is a hot issue, will muddy the waters rather than ensure an effective regulatory framework for cloud services in the country concerned.

6. Recommendations

- The main public sector concerns regarding a move to the cloud are information security and data protection (or privacy). Hence any new, cloud-specific rules should aim to address, as a first step, potential gaps or uncertainties concerning the application of these rules to the cloud.

- Data localization rules may not be the right way to address information security and data protection concerns. If generalized, they can have an adverse effect on the local take-up of cloud services by public and private clients.

- It is important for governments to rely on a classification of data and information based on their level of sensitivity, with different security and accreditation rules for each level. Such a classification need not be uniform throughout the whole state mechanism, but its basic principles should be coherent and transparent.

- Data protection and information security rules largely overlap or complement each other, but they are not substitutes, nor do they address the same objectives. Cloud-specific rules should aim to reconcile any conflicts between the two.

- Overall, the EU data protection regime offers the strictest set of such rules in the world, and although it does not necessarily provide a model it is at least a helpful checklist. Other, less burdensome and more cloud-specific models can be found the ISO/IEC 27018 code of practice. Having said that, the sheer weight of the EU economy, combined with its restrictive regime on the transfer of data to third countries and certain extraterritorial provisions may well have a spillover effect outside the EU and hence exercise some influence on future data protection rules across the world. Non-EU public authorities dealing with such rules should keep this in mind.

- Cloud computing raises novel issues of jurisdiction. Governments wishing to make use of cloud services should anticipate any associated risks bearing in mind that these may well be disassociated from the physical location of the data.

- Cloud regulation should not be seen as an end in itself, but should address clear existing needs or gaps, and can consist of a mix of mandatory and non-binding

rules. As a rough rule of thumb, the more detailed the rules, the less binding they should be.

- As a policy matter, it may be difficult to justify cloud-specific legislation. Any government embarking on such a task should therefore use this opportunity to address as comprehensive a set of cloud-specific issues as possible under one and the same legislative initiative, and try to ensure that any associated regulatory and supervisory functions are entrusted to a single ministry or authority.