Capgemini
CONSULTING.TECHNOLOGY.OUTSOURCING

SOGETI

Security Assurance in Cloud Adoption

With a cybersecurity approach that's right for their business, organisations can adopt cloud with confidence

Executive summary

The benefits of cloud computing are now widely understood, yet adoption at enterprise level is still slower than we would expect. A major reason is concern about security. Cloud is often perceived as being inherently less secure than conventional infrastructure, and revelations about the activities of security agencies coupled with growing awareness of the cost of cybercrime have increased would-be adopters' hesitation.

However, with a pragmatic approach based on a solid understanding of the realities of cloud security, cloud can be safe. Organisations can take advantage of cloud-based services while managing their risk responsibly and fulfilling their business, legal, regulatory and compliance obligations.

Whether the requirement is for Software as a Service, Platform as a Service, or Infrastructure as a Service, it is important to choose the right deployment option. Private, community, public and hybrid cloud offer varying degrees of control over security, privacy and assurance, as well as different levels of risk.

The most important dimension of risk to consider (other than in private cloud) relates to multi-tenancy: how the assets of users who share infrastructure are segregated from one another, and who has access to what. Compliance risk is also an important consideration, particularly where physical audits or knowledge of the physical location of data are mandated. Other commonly cited risks relating to nation state access, supplier lock-in and availability are probably no greater in cloud than with on-premise solutions. Failure to address risk can lead to breaches such as data leakage, data loss, or account or service traffic hijacking, which can sometimes have more serious consequences in the cloud, though cloud also has advantages such as easier disaster recovery.

We believe these considerations should not put organisations off using cloud. However, it is important to identify the risks of cloud before deployment, and ensure that they are appropriately managed.

To do so, organisations need to put in place an enterprise architecture framework that ensures that security is built into the infrastructure – and that includes selecting the right cloud deployment option, from a supplier who can offer the right security measures. (It is preferable to put this framework in place at the time of implementation, but it can also be adopted by existing cloud users, subject to any limitations imposed by existing contractual arrangements.) The framework needs to address security concerns with respect to four topics:

- Data storage – covering issues raised by multi-tenancy, such as how different clients' assets are segregated, and what assurances about separation can be provided to the data owners.
- Location of data – including compliance with corporate policies and government requirements regarding issues like data residency and sovereignty.
- Data security – including the approach taken to encryption.
- Incident response planning – how the inevitable attacks will be handled.

The advantages of cloud are so great that organisations can't afford to be held back any longer by security concerns. What they need is a migration strategy with security at its core, enabled by an architectural framework of the kind outlined above. The strategy must identify and quantify any risks that are unacceptable in view of the organisation's appetite for, and tolerance of, risk.

With this strategy and framework in place, it becomes much more straightforward to choose the right cloud deployment options and establish the necessary security controls. The enterprise can then use cloud securely and confidently to enable it to evolve and grow.

Table of contents

Introduction

The cloud computing model has established itself as a viable delivery mechanism for a wide variety of enterprise purposes, from messaging and collaboration to hosting of business-critical services. It is no longer the case (if it ever was) that cloud is only suitable for hosting development and test environments; increasingly, the public cloud is now the default option for production workloads. Indeed, the UK Government has gone further and adopted a principle of "cloud first" for procuring new ICT services.

The benefits of cloud are now well known, but can be summarised as follows:

- Cost-effectiveness of IT operations
- Better total cost of ownership (TCO) control
- Scalability
- Optimised resource utilisation
- Agility
- Flexibility
- Better IT resource management and business focus
- High reliability/availability
- Rapid development, deployment and change management
- Green IT

Despite these attractions, we do not yet see cloud services being adopted as widely as might have been expected, given the advantages. A major inhibitor is that many see the cloud model as inherently less secure than more traditional on-premise IT delivery models. According to one recent survey, 90% of organisations are "very concerned" or "moderately concerned" about security in public cloud[1].

Edward Snowden's revelations about the activities of national security agencies such as the US's National Security Agency (NSA) and UK's Government Communications Headquarters (GCHQ) have undoubtedly increased the unease of would-be cloud adopters. So have other high-profile and well-publicised security incidents, coupled with an increasing awareness of the costs of cybercrime generally (the average yearly cost of cybercrime for each large organisation participating in a recent study was US$7.6m)[2].

Although these are genuine concerns, they need not be an obstacle to cloud adoption provided the right approach is taken. In this paper, we will explain how a pragmatic approach to implementation can reduce the risks associated with cloud to an acceptable level so that adoption can go ahead with confidence.

[1] Information Security Community on LinkedIn. Cloud Security Spotlight Report. http://pages.cloudpassage.com/2015-cloud-security-survey-report-linkedin.html

[2] Ponemon Institute, October 2014. 2014 Global Report on the Cost of Cyber Crime. http://www8.hp.com/us/en/software-solutions/ponemon-cyber-security-report/

Cloud computing – service and deployment models

Before embarking on our discussion of cloud security, we will briefly review the basics of service models (what you can use cloud for) and deployment models (how you can use it).

For the purposes of this paper we will follow the definition of cloud put forward by the National Institute of Standards and Technology (NIST), which suggests that cloud computing has five essential characteristics[3]. These can be understood in terms of an analogy with rental cars:

> **On-demand self-service** (It is easy to rent a car, you can book a reservation by phone or online.)
>
> **Broad network access** (There is a broad network of rental car agencies around the world to give you access to a car rental.)
>
> **Resource pooling** (The rental car companies manage a pool of cars in any given city to meet demand. You don't have to worry about it. If one agency is out of cars they will often refer you to another to help you find a car.)
>
> **Rapid elasticity** (Rental car companies move cars into a particular location when there is a large event and they know demand will be high. They scale up and down to meet the demand.)
>
> **Measured service** (You pay only for the time you used the car. Once you turn it back in you are done. No maintenance, insurance, fuel, tires, etc.)
>
> *Source: Microsoft*[4]

[3] Mell, P, and Grance, T., National Institute of Standards and Technology, 2011. The NIST Definition of Cloud Computing. http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf

[4] Merrihew, A., Microsoft, 2014. Cloud Computing: How to explain it to others in your organization. http://www.microsoft.com/en-us/government/blogs/cloud-computing-how-to-explain-it-to-others-in-your-organization/default.aspx

Cloud service models

There are three basic ways to use cloud.

Software as a Service (SaaS). An application is hosted offsite for multiple clients who access it via the internet on a pay-per-use basis. Applications might include customer relationship management (CRM), enterprise resource planning (ERP), email, web hosting, etc. The cloud service provider (CSP) keeps the infrastructure running and does all the patching and upgrades. The primary benefit is reduction of upfront costs because the application is not procured and licensed in the conventional way. Examples of SaaS are Google Apps and Microsoft Office 365.

Platform as a Service (PaaS). This provides all the resources needed to build services and applications in the cloud without having to download or install the software or applications. It is primarily used for application design, development, testing, support, etc. The main benefit is low development cost. Examples of PaaS are Microsoft Azure and Amazon SimpleDB.

Infrastructure as a Service (IaaS). Here, the CSP rents out servers, storage and network equipment and other infrastructure elements. The infrastructure can be dynamically scaled up or down based on the demands of multiple tenants who are using the resources. Examples of IaaS are Amazon Elastic Compute Cloud (EC2) and CloudSwitch.

To show the various cloud service models in action, let's look at the needs of a large airport such as Heathrow:

- The runway and airport operations that form the underlying backbone of the computing infrastructure (servers, network and storage) can be procured from an IaaS provider.
- The plane/coach/transit train operations and associated applications can be customised and developed using a PaaS provider.
- The airline staff/company applications can be hosted as a SaaS service.

Other analogies such as "pizza as a service"[5] may also shed light on the various models.

Deployment models

Whether the requirement is for SaaS, PaaS or IaaS, there is a choice of four main deployment models.

Private cloud. This model offers a distinct and secure environment that is accessible only by a single, specified organisation. Levels of network security, control and privacy are all relatively high. Resources are drawn from a ring-fenced pool of physical computers, hosted internally or externally, and accessed across private leased lines or secure encrypted connections via public networks. As an example, transport company Arriva UK Bus has opted to replace dedicated servers with private cloud to handle spikes in traffic and cope with future demand.

Community cloud. This multi-tenant model allows several organisations with similar needs and concerns – such as compliance and regulatory measures – to work on the same platform. Government, healthcare and some regulated private industries value the added security features of a community cloud, which may allow them to comply with regulatory requirements that rule out public or hybrid cloud. Examples include the UK government's G-Cloud for public sector organisations and ATI Cloud, SITA's dedicated service for the airline industry.

Public cloud. Here a CSP makes resources, such as applications and storage, available to the general public or a large industry group. Examples include Facebook and Gmail.

Hybrid cloud. In this model, two or more clouds (private, community or public) are bound together by standardised or proprietary technology that enables data and application portability and load balancing (for example, "cloud bursting" which uses private or community cloud most of the time but adds capacity from the public cloud at times of peak demand). In some instances, the private element may be a traditional data centre rather than a private cloud. For example, Action for Children, a UK-based charity, adopted cloud for its public web hosting but kept sensitive personal data in-house.

[5] Barron, A., 2014. Pizza as a Service. https://www.linkedin.com/pulse/20140730172610-9679881-pizza-as-a-service

Figure 1 shows the relationship between various cloud deployment models, and their associated strengths and weaknesses.

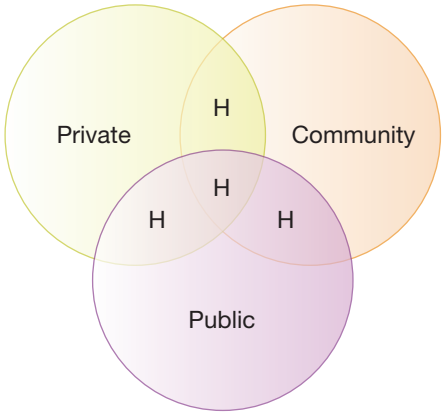

Model	Strengths	Weaknesses
Public	Agile; cost-effective; "illusion of infinite resource"	Multi-tenant; data residency; assurance; standard contracts
Private	Dedicated use; assurance; scope to negotiate SLAs etc	Expensive compared with public model; no "illusion of infinite resource"
Community	Designed for a specific, shared set of security requirements	Difficult to govern; need to manage all stakeholders
Hybrid (shown as H in diagram)	"Best of breed" suppliers can be switched in and out	"Weakest link"; must cater for security issues across ALL suppliers

Figure 1: Deployment models (source: Capgemini based on NIST information)

Capgemini expects the hybrid deployment model to become increasingly dominant in future as it can ease the cloud consumer's transition to cloud services while addressing issues such as payment card industry (PCI) compliance and data protection laws restricting international export of personal data.

Cloud service brokers

Cloud service brokers provide an intermediate layer between multiple cloud vendors and users while offering services such as selection, aggregation, integration, performance management and security. They should be able to unify legacy services and new multi-sourced cloud-based offers into a common management platform, and provision preconfigured applications as part of a service integration solution.

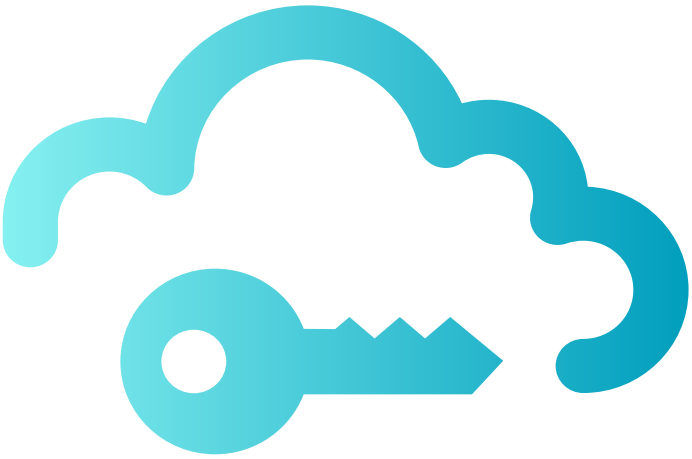

Risks associated with cloud computing

Having reviewed the service and deployment models, we will now consider the risks that are commonly believed to be associated with cloud. These will influence the choice of deployment model, as well as decisions about whether to use cloud at all.

There are undoubtedly risks associated with the use of cloud-based services, just as there are risks associated with other delivery models. The most talked-about risks regarding cloud computing lie in five areas: multi-tenancy, compliance, nation state access, lock-in and availability. We assess each in detail below, and give examples of real-life security breaches.

Multi-tenancy

Multi-tenancy is a core component of many cloud services (other than private cloud) and can entail sharing of storage, compute and/or application resources. This is probably the most important dimension of risk to consider. Provided sufficient thought is given to how the assets of users who share infrastructure are segregated from one another, and who has access to what, it need not be an obstacle to cloud adoption.

In the cloud environment, organisations must place their confidence in the security barriers operated by the CSP. The main difference between this and placing trust in third parties for traditional service delivery or data hosting lies in the extent to which resources are shared. Whereas in a traditional outsourced environment services will generally run on dedicated hardware, in a cloud environment the physical infrastructure will typically be shared.

Clearly, this sharing represents an increased level of risk compared to the use of dedicated hardware, but we believe the risk can be managed, for example through the use of cryptographic controls to separate each tenant's security domains from the rest, and through the use of security monitoring technologies to detect unusual or suspicious activities indicative of a compromise.

As with any business decision, cloud-derived benefits such as flexibility and speed to operation need to be weighed against the elevated security risks associated with the introduction of new attack surfaces due to multi-tenancy. The risk management function of the business (and not IT) must ultimately adjudicate in each case.

Often the benefits will outweigh the risks, especially as the increased abstraction, logical segmentation and automation often used to support multi-tenancy, for instance via API-driven image provisioning, can provide additional security benefits.

Compliance

This area includes:

1. Legislative and regulatory compliance, such as data protection legislation including the UK Data Protection Act 1998.
2. Other compliance requirements such as those associated with the Payment Card Industry Data Security Standard (PCI-DSS).

In our view, the difficulties of complying with data protection legislation in the cloud environment are often overplayed. Various mechanisms make it possible to export personal data overseas, for example:

- The Safe Harbor Agreement between the US Department of Commerce and the European Union, which enables the personal data of EU citizens to be exported to US organisations that abide by it.
- Binding corporate rules (BCRs). These allow global organisations to make intra-company transfers of personal data in compliance with the relevant EU Directive.
- Model contract clauses issued by the European Commission, which ensure compliance with EU requirements relating to the export of personal data (Principle 8 of the UK Data Protection Act 1998).

Legal advice should be sought to ensure that these standard mechanisms are sufficient for specific organisational needs.

It can nonetheless be problematic to achieve end-to-end PCI-DSS certification for a service delivered via the cloud. The PCI Security Standards Council released specific guidance on the use of cloud computing in 2013. In essence, those seeking PCI-DSS compliance must ensure that they have a complete understanding of the scope of any claimed PCI-certified services offered by their CSPs. The consumer of the service can then build its own PCI-compliant services on top of the existing certifications.

Defining the split in responsibilities between consumer and CSP to a level sufficient to achieve PCI-DSS compliance is a non-trivial task – but this is possible using either private, community or hybrid cloud providers.

This risk should not deter organisations from pursuing a cloud strategy. Compliance can be a problematic area even with

traditional delivery models, and it is always the responsibility of the end-user organisation. With cloud, the lower down the stack the CSP stops, the greater the cloud consumer's responsibility for ensuring compliance – i.e. the consumer has most responsibility in IaaS and least in SaaS.

Nation state access

Greater insight into the activities of the NSA, GCHQ and other nation state agents provided by Edward Snowden and others have heightened concerns about the security of data when hosted in cloud services located in foreign nations. Even before Snowden, many organisations were worried about the data access provisions incorporated into the US Patriot Act, passed following the attacks of 11 September 2001.

Organisations worried about the ability of national security services to access data hosted in a cloud service should ask themselves how confident they can be that their on-premise security systems are adequate to keep out such well-resourced threats.

For instance, the Patriot Act is as applicable to standard models as it is to the cloud model, a point illustrated by the US Government's seizure of the Belgium-based SWIFT organisation's European payments data, which had been mirrored to a US-based operating centre. The Act is not even limited to data held within the US; its applicability is determined by the location of the holding company rather than that of the data. Data can therefore be seized under traditional outsourcing arrangements with US-based suppliers (or suppliers with a US-based subsidiary that cannot demonstrate adequate independence).

Given the technical capabilities of state security actors and the non-technical mechanisms available to them in the form of bribery or extortion, it is highly likely that such agencies will compromise your data if they want it highly enough – whether your data is held in the cloud or on-premise. Hosting your data in a trusted physical location will not help you unless you have complete control of the entirety of your supply chain, operational management and all communications.

The bottom line is that businesses can lose out on the benefits of cloud adoption, or spend so much money as to negate the benefits, in order to combat a risk that is in reality minimal (as well as equally applicable to non-cloud hosting). Once again, this risk should not be seen as an insurmountable barrier to cloud adoption.

Lock-in

Some prospective users are concerned about getting locked in to using a particular supplier. It's true that, at present, it is not straightforward to move between cloud suppliers, or to bring a cloud activity back on-premise.

SaaS. To migrate between suppliers, the data must be exported from the current environment and transformed into a format compatible with the new one. Providing the volume of data is manageable, this may be relatively straightforward as most providers support the export of data as .csv files, or in another commonly parsable format. The data elements may still require some transformation to support a different data model, and the user organisation will need to be able to export and then re-implement the access controls governing the data and functionality available to its users.

PaaS. This can be the most difficult model to move. The organisation must export not only the data but also the code. This can be a significant undertaking, particularly if there is a mismatch between the functionality (e.g. software versions or scripting support) supported by the PaaS providers.

IaaS. Here, migration is usually relatively straightforward, unless operating system images are saved in proprietary formats that preclude running on similar hypervisor environments elsewhere. Once again, data must be exported. If an organisation is planning to bring a service on-premise from a CSP, it must ensure that it has all of the appropriate resources available to host the service.

However, migration is likely to become easier in future. Just as it has become easier to switch between UK electricity suppliers as competition and standardisation have increased, the cloud market is likely to see an emergence of common methodologies and competition that will simplify the process of switching CSPs. Work continues on many different standards in the cloud space, via the Open Group, the Distributed Management Task Force (DMTF), the IEEE and elsewhere, which helps to promote standardisation and encourage the development of a flexible and dynamic market.

In addition, cloud service brokers will increasingly shoulder the burden of migration. These brokers integrate and manage cloud offerings from numerous providers to present unified atomic business services to one or more consumers. They will encourage CSPs to provide similar and compatible features.

So, although it may look as if lock-in is a major risk, that risk can be expected to reduce over time. In addition, it should be remembered that migrating services in the on-premise world can also be a painful process.

Availability

Critics of cloud provisioning ask organisations to consider the potential damage should they move business-critical services such as email to the cloud and then lose their connection to the cloud or experience an outage in the cloud service itself. These concerns are not baseless as there have been well-publicised outages in leading services including Microsoft Azure.

Again, though, these concerns need to be put in perspective. If an organisation has already outsourced its back-office functions to a more traditional outsourcer then they face these risks now – even though the cloud outages attract more publicity. Moreover, the cloud service is likely to offer more resilience at a much more competitive price thanks to its distributed nature, and cloud systems can provide affordable disaster recovery (DR) options.

To evaluate this risk, organisations should compare availability statistics for their current services with equivalent cloud services, also taking cost into account. Cloud usually does quite well in these comparisons.

Evaluating the risks

Of the risks outlined above, we consider only two categories to be significantly more applicable to cloud computing than to traditional delivery models. These are multi-tenancy, due to the new attack surfaces not present in single-tenant models, and compliance, particularly where compliance requirements mandate the use of physical audits or knowledge of the physical location of data assets.

On the other hand, cloud computing can offer significant security benefits to an organisation. These include, inter alia, improved data centre security, increased resilience, introduction of new security capabilities through increased automation and abstraction, reduced reliance on portable media, and greater concentration of skilled security resources.

To get a realistic perspective on the risks associated with cloud, organisations need to be far more clear-eyed about those inherent in their current delivery model, including risks associated with outdated data centres or legacy applications, which tend to be underestimated. Similarly, thought needs to be given to how data is managed, secured and transported within or between organisations; for example how much sensitive data is still being transported on unencrypted devices such as memory sticks or smartphones?

Viewing cloud risks in this light makes them look more acceptable, and we believe it will encourage more users to go ahead with cloud deployment. Before doing so, however, it is vital to establish an enterprise architecture framework that ensures security is built into the infrastructure – and that includes selecting the right cloud deployment option, from a supplier who can offer the right security measures. In the next chapter we discuss the best way of tackling this task.

The real threat of security breaches

Some recent security breaches illustrate the risks noted above and their implications. Breaches can be divided into several groups[6], of which the most important arguably are:

Data breach

Researchers have described[7] how a virtual machine could use side channel timing information to extract private cryptographic keys being used in other virtual machines on the same physical server. However, in many cases an attacker wouldn't even need to go to such lengths. If a multi-tenant cloud service database is not properly designed, a flaw in one client's application could allow an attacker access not only to that client's data, but to every other client's data as well.

While data loss and data leakage are serious threats in cloud computing, the measures put in place to mitigate one of these threats can exacerbate the other. You may be able to encrypt your data to reduce the impact of a data breach, but if you lose your encryption key, you'll lose your data as well. Conversely, you may keep offline backups of your data to reduce the impact of a catastrophic data loss, but this increases your exposure to data breaches.

Data loss

Data stored in the cloud can be lost through various factors including malicious hackers[8], accidental deletion by the CSP or physical catastrophe such as a fire or earthquake. The CSP should take adequate measures to back up data, but the customer needs to take responsibility as well. Many compliance policies require organisations to retain audit records or other documentation. If an organisation stores this data in the cloud, loss of that data does not take away the organisation's responsibility regarding data protection.

There are additional reasons to take this threat seriously: under the new EU data protection rules, data destruction and corruption of personal data are considered forms of data breach and would require appropriate notifications. Moreover, loss of personal data in the cloud could incur hefty fines.

Managing this threat is not straightforward. As noted above, if a customer encrypts data before uploading it to the cloud, but loses the encryption key, the data will be lost as well.

Account or service traffic hijacking

Even well-regarded CSPs such as Amazon and WordPress have experienced theft of credentials enabling hijacking, as well as fraudulent use of their services leading to theft or unexpected usage charges for their customers.

Attack methods such as phishing, fraud and exploitation of software vulnerabilities are not new but still achieve results. End-users' reuse of credentials and passwords between systems increases the impact of such attacks. With stolen credentials, attackers can often access critical areas of deployed cloud computing services, allowing them to compromise the confidentiality, integrity and availability of those services, and cause reputational damage to companies whose accounts or service instances they use to launch further attacks.

Organisations should be aware of these threats and the available strategies that can contain the damage (and possible litigation) resulting from a breach. They should look to prohibit the sharing of account credentials between users and services, leverage strong two-factor authentication techniques where possible, and secure connections (e.g. SSH and SSL).

---

[6] Cloud Security Alliance, 2013. The Notorious Nine: Cloud Computing Top Threats in 2013. https://cloudsecurityalliance.org/download the-notorious-nine-cloud-computing-top-threats-in-2013/

[7] Zhang, Y., Juels=A., Reiter, M.K., Ristenpart, T., 2012. Cross-VM Side Channels and Their Use to Extract Private Keys. https://www.cs.unc.edu/~reiter/papers/2012/CCS.pdf

[8] Honan, M., 2012. How Apple and Amazon Security Flaws Led to My Epic Hacking http://www.wired.com/2012/08/apple-amazon-mat-honan-hacking

Security architecture for cloud services

Organisations need to ensure that security is built into their cloud infrastructure – which includes selecting the right cloud deployment option, from a supplier who can offer the right security measures.

Putting in place the right enterprise architecture framework makes it easier to do this: the framework should contain the processes, products, tools and techniques needed to create a complete IT system architecture for all infrastructure – not only cloud. Capgemini's Integrated Architecture Framework (IAF) (illustrated in figure 2) is an example of such a framework. Other examples include SABSA and TOGAF.

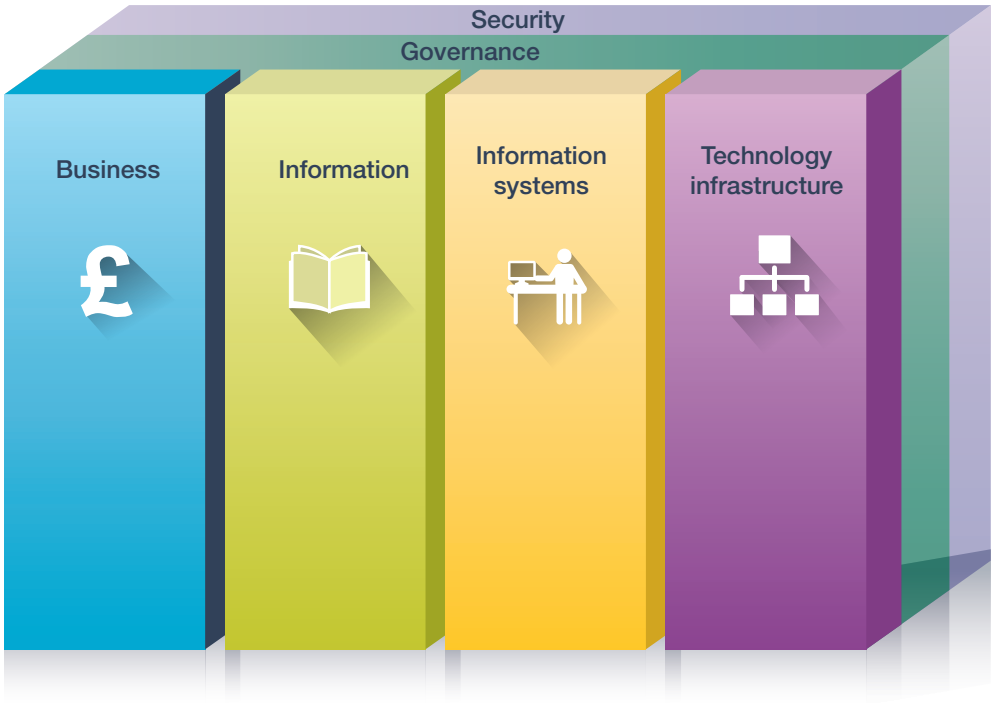

Figure 2: Capgemini's Integrated Architecture Framework – an example of a framework incorporating security and governance[9]

The framework helps to ensure that security and governance are an integral part of the overall architecture, positioning the organisation to create the right controls (including monitoring). It also helps to limit cost: it is almost always cheaper to build security in from the start than to retrofit it. Control of non-financial costs such as those associated with damage to legislative compliance, reputation and customer confidence is another driver for adopting a framework.

Following an architecture framework has another important advantage: it makes it possible to trace characteristics of the technical implication back to either requirements or risks – something that is often vital for business governance and legal compliance.

[9] Capgemini. Integrated Architecture Framework (IAF), version 4

Security architecture design for cloud

Whichever architecture framework is selected, the security architecture design for cloud-based services or systems will follow the approach shown in figure 3.

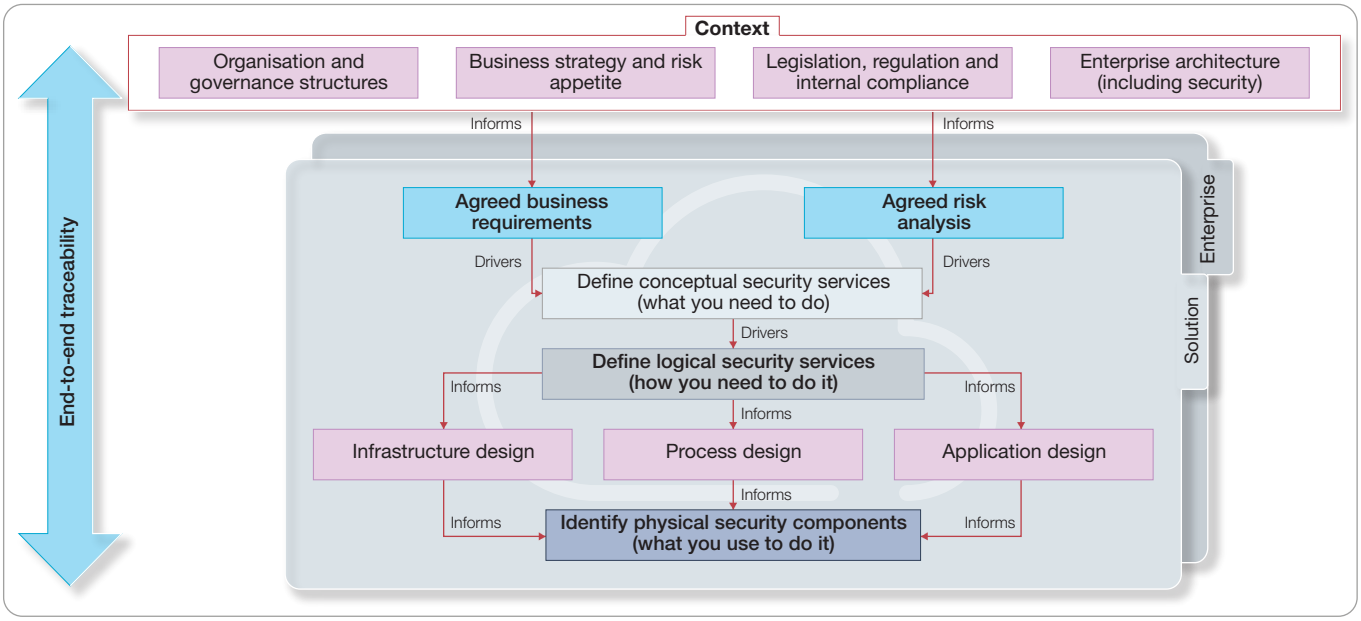

Figure 3: Approach to security architecture design

The design entails four views: the context, the conceptual security architecture, the logical security architecture and the physical security components.

Context. The context consists of:

- High-level business requirements
- Business needs
- Regulations and legislation
- High-level risk assessments
- The division of responsibility between the consumer, CSP and third parties

Conceptual security architecture. Based on the context, the architect defines the business context, security domains and the interconnections between the domains, which together constitute the conceptual security architecture. This architecture must be agreed by all the relevant stakeholders before proceeding with the logical controls selection.

Logical security architecture. The logical security architecture for cloud is then derived from the conceptual security architecture. The architect selects controls from applicable standards, compliance and legal regulations such as:

- Cloud Security Alliance (CSA) Cloud Controls Matrix (CCM) – the Security, Trust and Assurance Registry (STAR) program provides varying assurance requirements and maturity levels of CSPs and consumers, and is used globally by customers, providers, industries and governments.
- ISO/IEC 27001 – the requirements for information security management systems (ISMS).
- ISO/IEC 27018:2014 – an auditable standard for CSPs who process personal data. Includes around 70 controls from different international data protection laws.
- Cyber Essentials – a UK government-backed, industry-supported scheme to help organisations protect themselves against common cyber attacks. It provides a clear statement of the basic controls all organisations should implement to mitigate the risk from common internet-based threats, within the context of the UK government's *10 Steps to Cyber Security*[10].
- PCI-DSS – standard for protecting credit card data.

Physical security components. The next step is the selection of the physical components that will be used to implement the logical security controls just identified. In-depth knowledge of existing services and offers, and their underlying security controls, is important here, especially during the bid or proposal stage.

[10] https://www.gov.uk/government/publications/cyber-risk-management-a-board-level-responsibility

Using a security reference model (SRM)

A security reference model (SRM) can be helpful during security architecture design. Figure 4 shows a generic SRM, which illustrates a series of conceptual security services grouped into high-level areas such as hosting, security governance and access management.

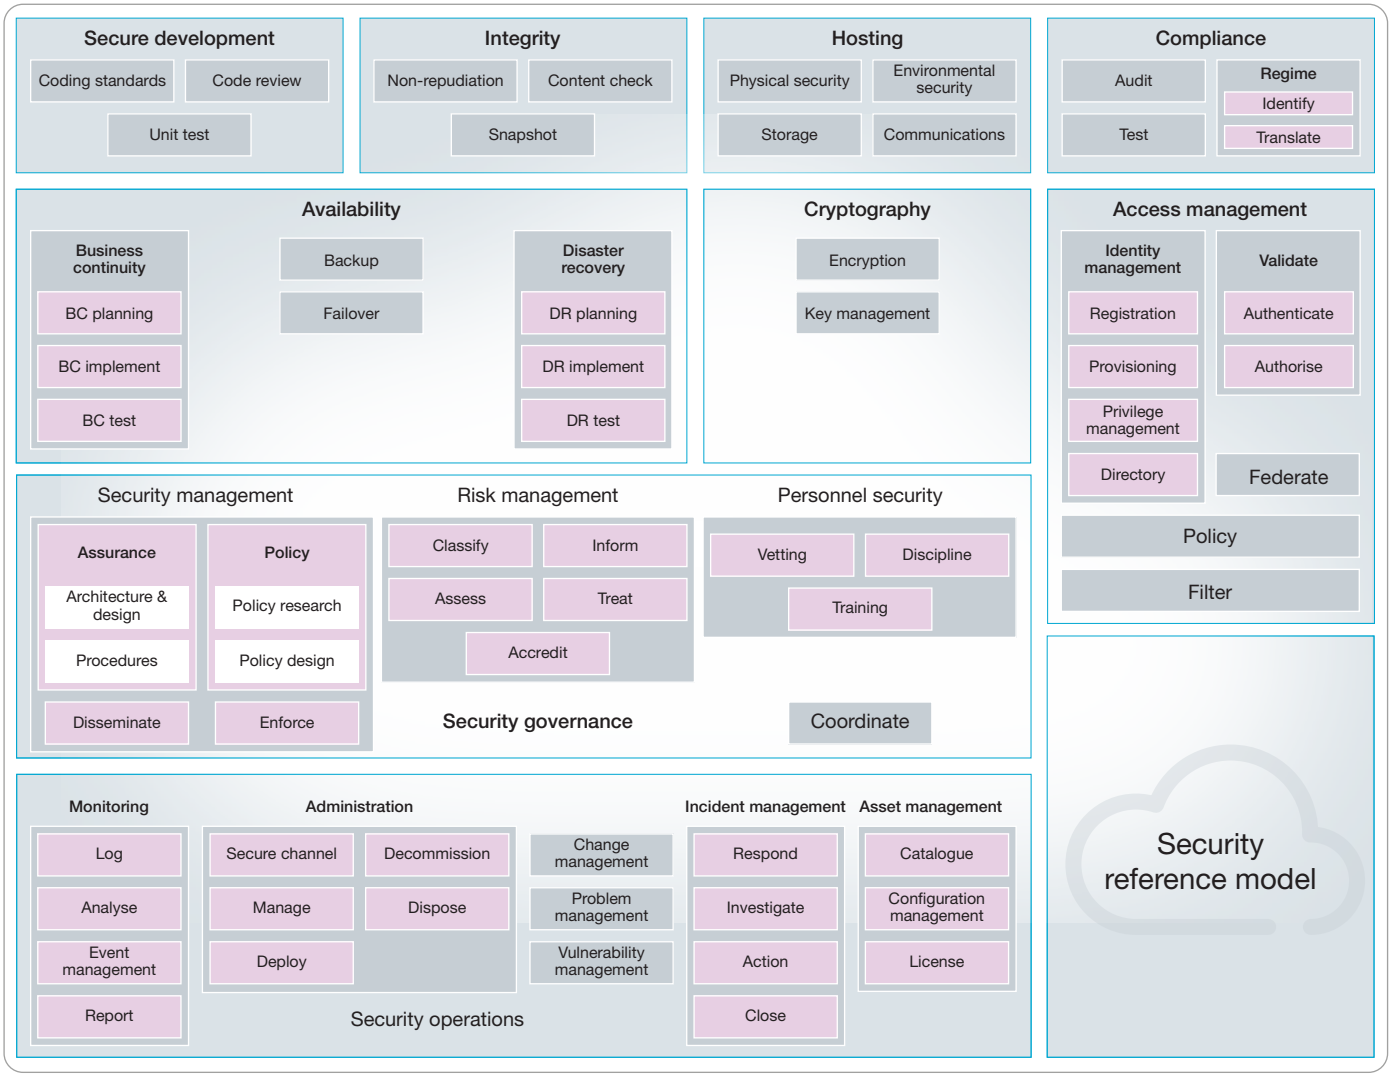

Figure 4: Generic security reference model (SRM)

This generic SRM can be used for identifying services that are in scope for various security areas. It can show the division of responsibilities for different cloud service models. It can also be used during procurement to conduct due diligence of the security controls claimed by CSPs.

The conceptual security architecture can be used as the basis for a logical security architecture. An architect can use the conceptual and logical security architecture (together with the wider business and non-functional requirements) to identify appropriate physical components and then implement appropriately secure solutions.

Issues to consider when designing cloud services architecture

- **Functional and non-functional requirements.** Gather the requirements and group similar ones together. You may find that groups of requirements point towards obvious conceptual security services that you will need.
- **Risk assessments.** Identify your assets, your threats and the business impacts of compromise, and again the need for certain conceptual security services may become obvious. Risk assessment methodologies and tools[11] that you may decide to adopt include:
  - ISO/IEC 27005:2011 – an international standard providing guidelines for information risk management.
  - Information Security Forum (ISF) – The Standard of Good Practice for Information Security is comprehensive and compatible with other well-known standards. It is particularly geared towards ISF's own Information Risk Analysis Methodology (IRAM) and automated tool, Risk Analyst Workbench (RAW).
  - CESG Information Standard 1/2 – with its supporting documents, this provides a suite of information risk management guidance for use, predominantly, by central government departments, the wider public sector and their suppliers. However, it can be used by any organisation to assess and manage technical risks.
  - US National Institute of Standards and Technology (NIST) Special Publication 800-30 – this is the US government's preferred risk assessment methodology, and is mandated for US government agencies. The methodology is usable by organisations of all sizes in both the private and public sectors. It is designed to be consistent with the ISO standards, and flexible enough to be used with other risk management frameworks.
- **Architecture frameworks.** Look at architecture frameworks like TOGAF 9.1, IAF and SABSA.
- **Interactions between on-premise and on-cloud.** Identify whether services are joint or provider responsibilities. This leads to a better understanding of the interface (management, contractual and security) between consumer and CSP.
- **Physical realisation.** After identifying the logical requirements of the service, start to identify the technologies or processes that best fit your needs – including the best deployment models.
- **Traceability and defensibility.** Identify and address any compliance and governance aspects. Traceability from either requirements or risks through to technical implementation is fundamental here – and convenient when the auditors visit to ask their usual questions.

Security considerations for cloud architecture can best be understood in terms of four fundamental topics:

- Data storage
- Location of data
- Data security
- Incident response planning

All of these four topics are relevant to the choice of deployment model, as well as to the way the chosen model is used. The next chapters consider each of them in turn.

[11] CESG, 2015. Analysis of information risk management methodologies. https://www.gov.uk/analysis-of-information-risk-management-methodologies


PLATFORM
COMMUNICATIO
WEBSITE
NETWORK
APPLICATION
COLLABORATION
SEARCH
RUNTIME
IDENTITY
COMPUTE
CONTENT
INFRASTRUCTURE

Data storage

Storing data in the cloud gives convenient access to it from anywhere, on demand and on multiple devices. In the cloud data storage model, the data is stored in logical pools, the physical storage spans multiple servers (and often locations), and the physical environment is typically owned and managed by a cloud hosting provider. It is based on highly virtualised infrastructure and has the accessible interfaces, near-instant elasticity and scalability, multi-tenancy, and metered resources associated with cloud generally.

Data storage security challenges in the cloud

When data is distributed, it is stored at more locations, and this increases the risk of unauthorised physical access to the data. Storage devices are generally where an organisation's critical and sensitive data resides. Providing data storage services is the CSP's responsibility for all types of cloud service model (IaaS/PaaS/SaaS); data security is the joint responsibility of the CSP and the customer, except in SaaS where the entire responsibility lies with the CSP.

In private cloud, the data for a given customer is stored on dedicated servers within the CSP's data centre. The customer normally has full visibility of their data. In the hybrid and public cloud deployment scenarios, by contrast, there are multiple security concerns, potentially including all of those discussed in our earlier chapter on "Risks associated with cloud computing".

Responding to data storage security challenges

Customers can tackle these challenges through a careful evaluation of cloud solutions generally and individual CSPs in particular. We recommend the following steps.

Assess data security and privacy. Carefully evaluate how the CSP and associated storage services handle data confidentiality, privacy, integrity and availability. Data location and data security are discussed in earlier sections. A comprehensive risk assessment are discussed in the following sections.

- What access rights does the CSP have over your data? Some CSPs offer encryption of customer data (at rest) and thus have no knowledge of the customer data. Sometimes the CSP may have rights to access stored data.
- Is any zero-knowledge policy offered by the CSP? A CSP with a zero-knowledge policy claims that the data is encrypted and the (private) keys necessary for decryption are only available to the customer.
- What are the CSP's service level agreements, non-disclosure agreements and industry accreditations?
- What type of cooperation with and support for forensic investigation is provided if there is a security breach?

It is the customer's responsibility (required by law for certain data, e.g. personal data) to judge the provider's stance on data security, but there has to be a trade-off between performance, flexibility and security.

Analyse data availability. Understand the CSP's DR capability. The provider should be sufficiently mature and technically competent to service customers from its DR centres with the same level of security, performance and flexibility as is normally provided. The customer must consider the availability factor carefully during the design and implementation phase and thereafter. This is not the responsibility of the CSP.

Evaluate connected applications and devices. Understand the applications and devices that are connected and have access rights to the cloud-based storage. Periodic audit cycles and vulnerability management of devices and applications can prevent intruders/hackers from exploiting application and firmware vulnerabilities.

Look for virtualisation. Evaluate whether the CSP supports **virtual storage area network (VSAN) and virtual volumes (VVOLs)** within its storage capabilities. VSAN and VVOLs provide data services at the virtual machine level and put an end to impractical logical unit number (LUN) storage management. If virtualisation is provided, consider whether it is sufficient to address any segregation requirements.

Data storage considerations for specific service models

Data storage security in the SaaS service model is the sole responsibility of the CSP, whereas in the IaaS and PaaS model it is shared. Capgemini recommends that IaaS customers adopt on-premise tokenisation/encryption technology to protect personal and business-critical data. The encryption should be either at OS level or application layer for IaaS, and at application layer for PaaS. Tokenisation and encryption are discussed in more detail in the "Data security" chapter, along with key management.

Location of data

Many companies face challenging concerns over data residency and data sovereignty when contemplating a move to the cloud. The degree to which an organisation moves to the cloud, and its choice of deployment models, will depend partly on corporate policies and on laws and regulations. These policies and laws may require the company to know where the data will be stored, how it will be managed and who will have access to it.

Even if the organisation is not required to know these things, lack of knowledge may be costly, for example if the data is stored in a country where data protection legislation is inadequate. It therefore pays to address this concern.

Location challenges

Passing information across international borders raises many legal complexities. An organisation with data in the cloud has to comply with the laws of the nation hosting the cloud data, as well as its own local laws. The country of origin may require measures that are not permitted by the host nation – for example, China does not allow certain algorithms or key lengths to be stored in its data centres. Clearly there can be conflicts between the requirements of the country of origin and host nation; in addition the host nation's laws may allow official authorities to access cloud data without any consultation or notification with the customer organisation.

Recent research from cloud security provider[12], based on data collected from its internal register of over 7,000 CSPs, concluded that 99% of cloud services do not meet the requirements of the reformed EU directive on Data Protection, including the right to be forgotten/data infidelity and deletion policies; data residency; data breach detection and notification; encryption and secure passwords.

If an enterprise does not deal with storage issues, it could be subject to fines and prosecution. These vary from country to country, but the new EU regulation proposes penalties of up to 5% of a company's annual revenue or €100 million, whichever is the higher.

Responding to location challenges

Define a strategy. Before entering the cloud, an organisation should consider creating a data storage plan or cloud strategy based on legal advice. This should define which data can be stored in the cloud, which host countries it is willing to use, and how it will ensure that storage complies with the laws of both local and host nations.

Determine data security approach. The enterprise should consider what security techniques it will employ (tokenisation or encryption) and what it expects from a CSP (for example, data centre certifications). Access by government agencies may not be avoidable so it is important that the organisation creates a way to deal with such events and to be informed when such activities take place. This will clarify the exact requirements that need to be met when approaching CSPs.

Create a governance framework with the CSP. This should outline which information can and cannot leave the country and any reporting standards the cloud supplier must meet, such as monthly reporting of breaches and notifications of any requests for enterprise data from third parties, including government agencies. It should also outline the steps taken by the CSP to ensure data security, such as understanding all disaster recovery procedures, the state of back-up technologies and locations, and any encryption technologies used (and when these are changed at any point).

Consider encryption. If data is hosted in a different country, not decrypting information in that country can provide added security. Encryption, tokenisation and key management are discussed in the "Data security" chapter.

Choose the right location. Our view is that when storing sensitive information that should not be viewed by anybody outside the enterprise, or storing information that can identify an individual, CSPs hosting within the local nation should always be chosen. Although this can be a more expensive choice, the cost will be offset by the reduced risk of data leakage and reduced effort and cost of ensuring compliance.

Choose the right provider. Particularly when storing information in other countries, the most important factor is the CSP itself. Scrutinise and research the CSP as much as possible, and review the contract with your legal department. Also assess the level of trust you can place in them. If the CSP is transparent from the beginning about all its processes and procedures regarding data hosting, and is able to meet and negotiate security requirements, then it is likely to work with you to help keep data safe and not hide any cost-cutting strategies it has in place. If a CSP has limited flexibility and is unable to meet your exact security requirements, then the lower costs of this service may often be more expensive in the long term.

[12] Help Net Security, August 11, 2014. Only 1 in 100 cloud providers meet proposed EU Data Protection requirements. http://www.net-security.org/secworld.php?id=17238

Data security

More and more sensitive information is hosted in the cloud, and traditional data protection controls like encryption of data at rest are unlikely to be applied once it's there.

According to the 2013 Global Encryption Trends Study from the Ponemon Institute[13], 53% of organisations have transferred sensitive or confidential data to the cloud. Yet only 39% of data in SaaS applications is encrypted, and that number drops to 26% when it comes to platform as a service (PaaS) and infrastructure as a service (IaaS) deployments.

Data security challenges

Data protection controls in a cloud context can be challenging to implement architecturally. Certain portions of the technology stack are opaque to the customer, and are managed by the CSP. So, unless the CSP specifically provides data protection features (and some do), a customer's ability to implement technical data protection controls without additional engineering might be constrained.

From a process standpoint, these controls can be challenging to pull off. There might be legitimate reasons why a CSP requires access to enterprise data – for example, to debug application functionality or to ensure that customers do not abuse resources. This means that the logistics of who will hold the encryption keys requires discussion, planning and well-thought-out procedures established in advance.

Finally, if the customer does not have visibility over how the CSP will be securing the data, it is difficult to demonstrate compliance with existing data protection standards, such as PCI-DSS, the US Health Insurance Portability and Accountability Act (HIPAA), and the UK's Cyber Essentials scheme.

Responding to data security challenges

Several key factors can help with the challenge of securing data in the cloud.

Data sensitivity. First of all, it is important to understand how sensitive the data is. A time-logging application used internally will store less sensitive information than an e-commerce store, which stores users' personal information and payment details. Even in this case, not all the information stored by the application can be deemed sensitive and securing data selectively can help minimise the processing overhead.

Encryption. A common approach to securing data is to encrypt it before it reaches the cloud. In the event of a data breach on the CSP's network, the attacker will need to invest additional time and effort to decrypt the customers' data. Encryption strength has improved as standards such as AES-256 have become widespread among cloud security providers.

Homomorphic encryption (HE). The idea behind this is to allow operations to be performed on data even though it is encrypted. With this recent improvement in encryption, the customer can prevent data from being decrypted in the cloud by doing the decryption on a specialised gateway or – often less practically – within its own network. While improvements have been made to make HE more practical, it is still an theoretical concept and only partial operations are currently supported.

Tokenisation. If the data is so sensitive that it cannot be stored outside a company's network, even if it is strongly encrypted, then an alternative such as tokenisation can be used. Tokenisation (also called pseudonymisation) is a process of substituting sensitive data with a randomly generated non-sensitive equivalent.

The original data is stored inside the organisation's network for maximum control and security, while the tokens can be safely stored elsewhere. Continuing developments in areas such as tokenisation and HE will allow customers to interact with their data in its encrypted state, which increases the privacy and security of data residing in the cloud.

[13] Thales, 2014. Global Encryption Trends Study. https://www.thales-esecurity.com/knowledge-base/analyst-reports/global-encryption-trends-study

Key management. Encryption and tokenisation have been proven to work, and should be important parts of the overall design. In both cases, however, the most important aspect is achieving control of key management. There is no benefit in using strong encryption if encryption keys can be stolen by attackers or handed over to government agencies, and the best way to ensure this does not happen is for the customer to have control over the master keys. However, if the CSP is offering encryption, they may control the encryption key. In the case of a data breach, or an inside attack, a leaked secret key will lead to the data being exposed. A better solution is to use a pair of keys – one is stored by the provider, the other by the customer. Even if the CSP is required to hand over information to a government agency, the data cannot be decrypted unless the customer provides their own key.

The role of the broker

A cloud access security broker (CASB) interjects security controls – such as tokenisation and encryption technologies – between users and the cloud-based services they consume.

CASBs can be useful in a range of situations. For example, they may be called in by organisations where cloud services have been established outside the control of IT; here, the CASB can help IT regain visibility and control.

Incident response planning (IRP)

Incident response is a linchpin of information security, and more so when it deals with cloud: the response process changes because you're likely to be co-mingling systems and resources in ways that dramatically complicate the process.

Even the most diligent planning, implementation and execution of preventive security controls cannot decisively eliminate the possibility of an attack on the IT infrastructure and resources. This means it's imperative to have a carefully thought-out incident response plan in place to ensure that the proper steps are taken to detect, isolate, contain and eradicate such an incident. The basic idea of incident response is to ensure damage control in case of an attack and gather evidence to take proper legal action to ensure recovery and containment of a security incident.

Challenges of incident response in the cloud

The key benefit of opting for cloud adoption is to leverage the capability of rapid and elastic provisioning of infrastructure, and to automatically scale out and scale down. The ability to provide what appears to be unlimited resource comes from resource pooling and virtualisation.

However, virtualisation and resource pooling coupled with huge amounts of data flowing in and out of the cloud pose a nightmarish challenge for incident response and forensics including:

- Identification and isolation of affected instance(s)
- Segregating the malicious traffic in a multi-tenant environment
- Assuring non-contamination of infrastructure resources
- Separation of non-incident related traffic and data

It can be seen that incident response planning in the cloud poses complicated challenges, partly because of cloud's inherent properties (elasticity, multi-tenancy, etc.). Dependence on the CSP adds further complexity. So does division of responsibilities between customer and CSP – though some responsibilities are always clear: for example it will always be the customer's responsibility to make sure an adequate DR plan exists.

Meeting incident response challenges

Organisations using cloud services need to put in place an incident response workflow of the type shown in figure 6. Each of the phases is reviewed below.

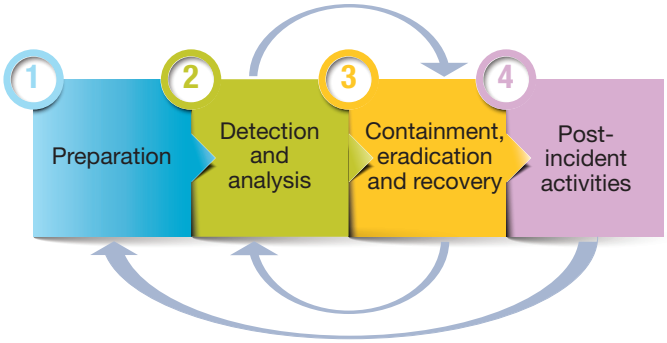

Figure 6: Incident response workflow

Preparation

In this key IRP phase, the cloud setup is assessed and gap analysis of existing security carried out, so that security measures can be adjusted. This exercise should be completed before cloud migration, and then frequently revisited to ensure that there is always an up-to-date understanding of the threat landscape, together with appropriate controls to mitigate the threats.

The first layer of control is the firewall, a basic perimeter-level blockade; it can be coupled with an intrusion detection/prevention system for better results. Anti-virus/anti-malware tools can help contain and eradicate known signature-based threats. A sandbox environment for malware analysis can enhance the threat mitigation capability, as can proxy servers and advanced next-generation firewalls.

Controls need to be established to mitigate distributed denial of service (DDOS) attacks and other risks to business continuity. A security incident and event management (SIEM) solution can give a robust holistic overview of the enterprise security landscape; its centralised log repository and rules-based or statistical correlation engine can help analyse huge volumes of logs from all connected devices to detect anomalies and potential attacks and incidents.

The volume of data and logs generated in the cloud may argue in favour of even more advanced approaches, like Bayesian analysis for faster and more accurate – even pre-emptive – detection of threats and attacks. The preparation stage should also define incident verification steps, classification and prioritisation guidelines, and triaging matrices to help the incident response/forensics team respond better in the event of an attack.

Detection and analysis
The detection and analysis phase of IRP always depends on the capability of the tools and controls in place, coupled with the skill of the incident response/forensics team. Intrusion detection systems (IDS), intrusion prevention systems (IPS) and SIEM play a key role in this phase, where the initial detection of any anomaly takes place.

Detailed log collection is a primary requirement for ensuring that a coherent picture is painted during analysis. Cloud makes all this more challenging as we need to ensure logs from each and every device are added to the log stream (particularly important due to the elastic nature of cloud), ensure the hypervisor logs are collected (proper communication and agreement with the CSP needs to exist in this regard), and acquire from the CSP the audit logs for the network, application, cloud administration roles and accesses, backup and restore activities, maintenance access and change management activity.

It's clear, then, that where cloud is used creating a coherent picture depends on the CSP's willingness and ability to collect and provide data (logs, forensic artefacts, etc.) relating to the provider-owned infrastructure. The customer also depends on the CSP for the analysis of that data, since the customer usually has little knowledge of the CSP-owned infrastructure. The elastic nature of cloud services and the use of resource pooling add to the complication of collection and interpretation of such data. In addition:

- Clock synchronisation across the infrastructure landscape is needed for reliable forensics, and the CSP must provide sufficient information in logs to accurately identify the time zone.
- Incident verification, classification, triaging and prioritisation by the incident response team, and communication between customer and CSP, are important to ensure that indicators of compromise or attack are shared at the right time to mitigate any business loss.
- Both parties need to maintain the integrity and authenticity of information (out-of-band communication channels and proper encryption schemas help).

Containment, eradication and recovery
Close cooperation between all stakeholders is paramount for effective and efficient containment, eradication and recovery after an incident, and for all the legal and privacy implications to be addressed.

In the cloud, given the varied service and deployment models, it is important to segregate responsibility for the containment and eradication strategies from ownership of the technology platform that is the target of the attack. For example, an IaaS consumer is primarily responsible for containment, eradication and recovery from incidents, but it is the provider's responsibility to mitigate perimeter attacks on the network or infrastructure (e.g. DDOS attack).

For SaaS and PaaS service models, it is advisable to assess the containment, eradication and recovery plan before subscribing. In these models, the customer's ability to respond to a security incident is diminished as access to the underlying infrastructure and associated control data is greatly reduced.

Post-incident activities (follow-up)
The organisation must update the incident knowledgebase with the outcome, and must carry out a gap analysis to identify if any additional measure, control or process can prevent such an event in future.

The customer organisation must work with the CSP to fill gaps. They should also request that security vendors update their signatures (for antivirus, anti malware, IDS/IPS, etc.) and issue new use cases for SIEM tools to pre-empt similar threats. Depending on the service model, this could be up to either the customer or the CSP.

Other considerations for IRP in the cloud

Prerequisites for IRP in a cloud context include:

- A thorough understanding of the respective responsibilities of the customer and CSP, so you can determine each party's role in analysing and responding to an incident.
- Synergy and a seamless flow of information between the customer's and CSP's incident response teams.
- Clarity about the CSP's own IRP processes, including the way it classifies incident response triggers.
- An individual contact for every time period in case of an incident.
- Written service level agreements (SLAs) for incident response backed by shared financial implications (e.g. penalties for the supplier) if the SLAs aren't met.
- Clear knowledge of the systems, security controls, data and process deployed by the CSP in order to understand where to respond in the event of an incident.
- A thorough assessment of security controls provisioned in the cloud, and a gap analysis, and any required measure to close the gap.
- A DR plan in case of an outage – this can use the customer's internal infrastructure or another provider. The CSP should collect data that helps the customer to meet the recovery point objective (RPO) and recovery time objective (RTO), and have logistics in place for moving data to an alternative system.

IRP considerations for specific service and deployment models

In addition, certain considerations apply to specific service models. For example, in a SaaS consumption model, the consumer is entirely dependent on the CSP and must have clarity of the incident response process, monitoring controls, and whether that information is accessible to the consumer. All this must be defined in the SaaS escrow agreement, which should clearly define the RPO and RTO.

In an IaaS consumption model, security controls should be in line with any contractual requirements from the CSP.

Private cloud deployment requires a clear understanding of the response and investigative impacts of the underlying technology. Understand how to segregate network traffic for a shared resource pool, how to isolate instances, how to carry out forensics in case of a security incident, how to create a centralised logging mechanism and patching mechanism, and how to address an outage of the cloud infrastructure itself.

Third-party solutions

One option is to buy IRP management solutions from service providers, who offer a set of guidelines and best practices that companies can follow when they have a security breach. They have the capabilities to take feeds from SIEM tools and even from threat intelligence services. The solution includes support for privacy laws and regulations adopted in different countries.

Security considerations – checklist

This list can be used to make sure all major security issues are addressed with a CSP either during due diligence or as part of a request for proposals (RFP):

Data storage

- Understand what access rights to consumer data the CSP has.
- Check if it provides encryption/tokenisation, including a zero-knowledge policy on consumer data.
- Make sure there is support for data breach and forensic investigations.
- Establish the storage location of data for backup or disaster recovery options.
- Consider how cloud storage is segmented, and whether techniques such as VSAN or VVOLs are used.
- Make sure you know how data is destroyed when moving among virtual storage objects.
  Location of data
- An organisation with data in the cloud has to comply both with the laws of the nation hosting the cloud data and with their own local laws, including data export laws.
- Before entering the cloud, consider creating a data storage plan or cloud strategy including a governance framework.
- Also consider distributing encryption keys within the enterprise and not providing these to the CSPs. This can help you secure data and understand who is accessing it.
- Sensitive information that should not be viewed by anybody outside the enterprise, or that can identify an individual, should always be stored using cloud hosting within the local nation.
- When storing information in other countries, the most important factor is the CSP itself. Choose a reputable provider that is widely recognised, particularly for hosting sensitive information in the cloud.

Data security

- First, evaluate the sensitivity levels of all of your data.
- Understand the encryption or tokenisation techniques offered for data security and make careful choices based on the business requirements, compliance aspects and data protection requirements.
- Also consider using a cloud access security broker.
- For data privacy purposes, find out who has access to the encrypted data, and learn about the associated key management process.

Incident response planning

- Build a thorough understanding of the roles and responsibilities of both parties: this is paramount for effective and timely incident response.
- Obtain granular clarity of the CSP's IRP and ensure that a responsible individual has been identified for all time periods in case of an incident.
- In the SaaS model, make sure you clearly understand the incident response process and monitoring controls, and whether that information is accessible to the consumer. (Remember that the consumer is entirely dependent on the CSP in this model.)
- In the IaaS model, ensure the security controls deployed by the consumer are in line with the cloud deployment architecture.
- In the case of private cloud deployment, obtain a clear understanding of the response and investigative impacts of the underlying technology.
- Consider buying solutions for incident response planning management to support existing processes.

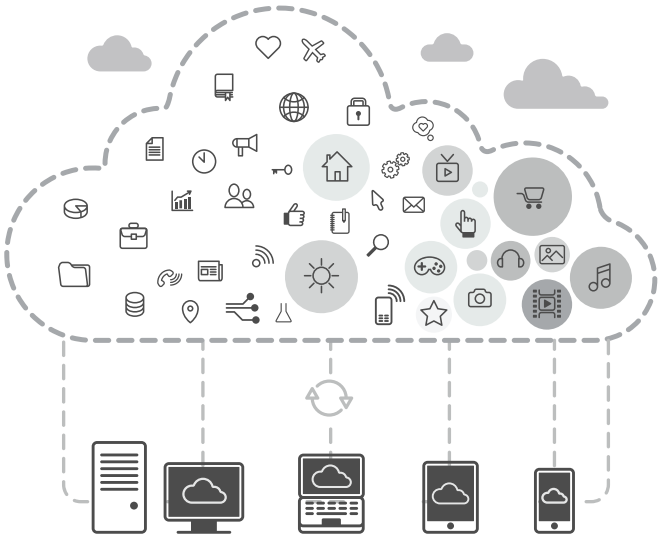

Conclusion

Prospective users of cloud need to understand and evaluate the risks associated with different cloud service and delivery models, paying particular attention to business requirements, data protection laws and compliance regimes.

Following an established framework for the security architecture design is good practice and provides traceability from either requirements or risks through to technical implementation. Organisations should use this framework as a basis for addressing four key topics: location of data, data security, incident response planning and data storage.

In all these areas, customers need to work closely with providers to put the right security in place. Some responsibilities are shared while others rest with either party but it is vital to know who is responsible for what, and ensure that the right governance processes, and mechanisms for sharing information, are in place.

There is evidence that security considerations are still inhibiting many organisations' adoption of cloud, but we believe that with the right approach to security they can pursue cloud strategies safely and reap the benefits.

List of abbreviations

Abbreviation	Description	Abbreviation	Description
AES	Advanced Encryption Standard	IDS/IPS	Intrusion Detection System/Intrusion Prevention System
API	Application Program Interface	IRAM	Information Risk Analysis Methodology
BCRs	Binding Corporate Rules	IRP	Incident Response Planning
CASB	Cloud Access Security Broker	ISF	Information Security Forum
CCM	Cloud Controls Matrix	NIST	National Institute of Standards and Technology
CPU	Central Processing Unit	NSA	National Security Agency
CRM	Customer Relationship Management	OS	Operating System
CSA	Cloud Security Alliance	PaaS	Platform as a Service
CSP	Cloud Service Provider	PCI-DSS	Payment Card Industry Data Security Standard
DB	Database	RAW	Risk Analyst Workbench
DC	Data Centre	RPO/RTO	Recovery Point Objective/Recovery Time Objective
DDOS	Distributed Denial of Service	SaaS	Software as a Service
DMTF	Distributed Management Task Force	SABSA	Sherwood Applied Business Security Architecture
DR	Disaster Recovery	SIEM	Security Incident and Event Management
EKM	Enterprise Key Management	SLA	Service Level Agreement
ERP	Enterprise Resource Planning	SRM	Security Reference Model
GCHQ	Government Communications Headquarters	TCO	Total Cost of Ownership
HE	Homomorphic Encryption	TOGAF	The Open Group Architecture Framework
HIPAA	Health Insurance Portability and Accountability	VSAN	Virtual Storage Area Network
HSM	Hardware Security Module	VVOLs	Virtual Volumes
IaaS	Infrastructure as a Service	XSS	Cross Site Scripting
IAF	Integrated Architecture Framework		

Capgemini
CONSULTING.TECHNOLOGY.OUTSOURCING

SOGETI

For more details contact:

Sunil Iyengar
Senior Cybersecurity Consultant | Cybersecurity Unit
sunil.iyengar@capgemini.com

Ian Cole
Head of Cybersecurity Consulting & Projects | Cybersecurity Unit
ian.cole@capgemini.com

About Capgemini & Sogeti

With more than 145,000 people in over 40 countries, Capgemini is one of the world's foremost providers of consulting, technology and outsourcing services. The Group reported 2014 global revenues of EUR 10.573 billion. Together with its clients, Capgemini creates and delivers business and technology solutions that fit their needs and drive the results they want. A deeply multicultural organization, Capgemini has developed its own way of working, the Collaborative Business Experience™ and draws on Rightshore®, its worldwide delivery model.

Sogeti is a leading provider of technology and software testing, specializing in Application, Infrastructure and Engineering Services. Sogeti offers cutting-edge solutions around Testing, Business Intelligence & Analytics, Mobile, Cloud and Cyber Security. Sogeti brings together more than 20,000 professionals in 15 countries and has a strong local presence in over 100 locations in Europe, USA and India. Sogeti is a wholly-owned subsidiary of Cap Gemini S.A., listed on the Paris Stock Exchange.

Capgemini and Sogeti are experts in IT infrastructure and application integration. Together, we offer a complete range of cybersecurity services to guide and secure the digital transformation of companies and administrations. Our 2,500 professional employees support you in defining and implementing your cybersecurity strategies. We protect your IT, industrial systems, and the Internet of Things (IoT) products & systems. We have the resources to strengthen your defenses, optimize your investments and control your risks. They include our security experts (Infrastructures, Applications, Endpoints, Identity and Access Management), and our R&D team that specializes in malware analysis and forensics. We have ethical hackers, five multi-tenant security operation centers (SOC) around the world, an Information Technology Security Evaluation Facility, and we are a global leader in the field of testing.

Find out more:

www.capgemini.com/cybersecurity or www.sogeti.com/cybersecurity