

INFORMATION SECURITY BRIEFING 01/2010

CLOUD COMPUTING

MARCH 2010

This briefing note is based upon a research document compiled on behalf of CPNI by Deloitte. The findings presented here have been subjected to an extensive peer review process involving technical advisers from CPNI, our information exchange groups and wider industry.

Disclaimer:

Reference to any specific commercial product, process or service by trade name, trademark, manufacturer, or otherwise, does not constitute or imply its endorsement, recommendation, or favouring by CPNI. The views and opinions of authors expressed within this document shall not be used for advertising or product endorsement purposes.

To the fullest extent permitted by law, CPNI accepts no liability for any loss or damage (whether direct, indirect or consequential and including, but not limited to, loss of profits or anticipated profits, loss of data, business or goodwill) incurred by any person and howsoever caused arising from or connected with any error or omission in this document or from any person acting, omitting to act or refraining from acting upon, or otherwise using, the information contained in this document or its references. You should make your own judgement as regards use of this document and seek independent professional advice on your particular circumstances.

Contents

1. Executive summary	4
2. What is cloud computing?	6
2.1 Cloud computing characteristics	6
2.2 Attributes of the cloud	7
2.3 Alternative views of the cloud's key attributes	7
2.4 The delivery models of cloud computing	8
2.5 The services and sub-services of cloud computing	9
2.6 Examples of 'the cloud'	9
3. What are the drivers of cloud computing?	11
3.1 Drivers of cloud computing	11
3.2 Benefits of cloud computing	12
4. Cloud computing architecture	13
4.1 Service architectures	13
4.2 Software as a Service (SaaS)	14
4.3 Platform as a Service (PaaS)	14
4.4 Infrastructure as a Service (IaaS)	15
5. Cloud computing maturity	16
5.1 Adoption of cloud computing	16
5.2 Maturity of the cloud	16
5.3 Vendor maturity and impacts on adoption	18
6. Evolution of cloud computing	20
6.1 History	20
6.2 Evolution of cloud technologies	21
7. Risks of cloud computing	24
7.1 Purpose and aim of section	24
7.2 Overview of risks	24
8. Business risks	28
8.1 Overview of business risks	28
8.2 Business risks associated with vendor or public clouds	28
8.3 Private clouds	29
8.4 Hybrid clouds	30
8.5 Community clouds	30
9. Security in the cloud	32
9.1 Cloud threats	33
9.2 Types of attackers	35
9.3 Security risks	36
9.4 Assessing the security of a third party cloud provider	40

9.5	Emerging cloud security threats	41
9.6	Examples of cloud security incidents	42
9.7	Mitigating advice	43
10.	Reliability and resilience	45
10.1	Overview of resilience issues	45
10.2	Benefits of cloud computing to continuity planners	45
10.3	Systemic and specific risks	45
10.4	Delivering resilience in the cloud.....	46
10.5	Delivering resilience through testing	46
10.6	Mitigating advice	47
11.	Usability and performance	48
11.1	Latency.....	48
11.2	Reducing latency.....	49
11.3	Network access.....	49
11.4	Network availability	50
11.5	Network performance.....	51
11.6	Monitoring of network performance.....	51
11.7	Mitigation advice	52
12.	Regulations and legislation.....	53
12.1	Overview of regulatory and legislation issues	53
12.2	Rights to data	54
12.3	Outsourcing contracts	55
12.4	Outsourcing, subcontracting and the FSA	55
12.5	Processing personal data in the cloud	56
12.6	Mitigation advice	57
13.	Organisational change	58
13.1	Organisational change management	58
13.2	Changing roles and responsibilities	58
13.3	Software development and testing methodologies	60
13.4	Mitigating advice	61
14.	Security testing	62
14.1	The objective: Information and technology risk management.....	62
14.2	The approach	62
14.3	Testing cloud services.....	63
14.4	Testing cloud delivery models.....	64
14.5	The solution.....	64
15.	The future of cloud computing	66
15.1	Drivers for future change.....	66
15.2	Predictions	68
16.	Glossary.....	70

1. Executive summary

This guidance provides a detailed overview of cloud computing, focusing on the potential benefits and risks as well as identifying mitigation advice to reduce vulnerability. The briefing is aimed at information security practitioners from organisations within the National Infrastructure as well as government agencies.

The key findings within this briefing are summarised as follows:

- There are conflicting descriptions of cloud computing and industry is still searching for a clear definition to encapsulate this profound but subtle technological evolution.
- Cloud computing offers customers considerable benefits in terms of being able to scale up or down IT services (applications, platform or infrastructure) on demand.
- Cloud services are leased and therefore customers do not incur capital costs of IT resources and equipment as they would in traditional IT service models.
- In cloud computing, IT operations are outsourced to the cloud; the risk is not. Accountability for customer (and business) sensitive data resides with the cloud customer.
- There is a lack of accepted cloud computing standards at an EU or worldwide level.
- There are wide ranging legal and regulatory issues in cloud computing covering rights to data, security loopholes, outsourcing and subcontracting. In particular, national laws and regulations governing interception and disclosure of data in jurisdictions in which data is stored, or transmitted across, differ considerably over who has access to that data.
- Third party cloud provider assurance and risk assessment activities are critically important for customers storing data in the cloud. The large number of third parties involved in the cloud, and its geographical dispersion, means that risk assessment activities are likely to be more complex, time consuming and costly.
- There are a number of IT data recovery risks associated with hosting data in multi-tenanted data centres, including the corruption of customer data, overloading of computing resources and proving the service meets disparate IT disaster recovery requirements.

The key recommendations for customers of cloud computing are:

- Customers should consider both customer-managed security controls such as encryption and identity management, as well as contractually agreed standards covering the right to audit, use of physical security, protective monitoring, data segregation controls and vulnerability management processes to secure their data in the cloud.
- Customers should give particular consideration to the laws governing the interception and disclosure of their data for all jurisdictions in which their data is stored or transmitted across.
- Customers should pursue a programme of assurance activities on their cloud providers to ensure contractually agreed standards are being met.

2. What is cloud computing?

There is, to date, no universally agreed industry definition of cloud computing and it is usual to find conflicting descriptions in any nascent industry. Cloud computing is a term used to describe a set of IT services that are provided to a customer over a network on a leased basis and with the ability to scale up or down their service requirements. Usually cloud computing services are delivered by a third party provider who owns the infrastructure. This section explores some of the alternative definitions for the cloud and begins by looking at the cloud's key characteristics.

2.1 Cloud computing characteristics

There is a level of consensus emerging around the characteristics of cloud computing, or the capabilities that must be adhered to an offering to be considered a cloud. These include:

- **Pay as you go** – payment is variable based on the actual consumption by the customer.
- **Highly abstracted** – server hardware and related network infrastructure is highly abstracted from the users.
- **Multi-tenant** – multi-tenant architectures allow numerous customer enterprises to subscribe to the cloud computing capabilities while retaining privacy and security over their information.
- **Immediately scalable** – usage, capacity, and therefore cost, can be scaled up or down with no additional contract or penalties.

There is a widely held view that the cloud is not a new concept. Indeed, many of the technologies and services associated with cloud computing, such as Web 2.0 or virtualisation¹, have been in existence for some time². What is different in the cloud is that these technologies are being implemented in new ways to provide *dynamic*, *scalable* and *virtualised* computing infrastructure, platforms and software.³

Cloud computing combines a number of computing concepts and technologies for Service Oriented Architecture (SOA), which may include Web 2.0 and the virtualisation of services and communication infrastructure. These technologies have allowed cloud customer organisations to achieve: improved utilisation and efficiency of their service providers' infrastructure through the controlled sharing of computing resources with other customers (multi-tenancy); and, greater flexibility to scale up and down IT services. In some respects, cloud computing represents the maturing of these technologies and is a marketing term to represent that maturity and the cloud services they provide.

¹ CPNI: CPNI Technical Note 1/2009 Security Considerations for Server Virtualisation, www.cpni.gov.uk/Docs/tn-1-09-security-server-virtualisation.pdf (January 2009)

² Further discussion is given in section 6. Evolution of cloud computing

³ ISF Briefing: Cloud Computing, www.securityforum.org

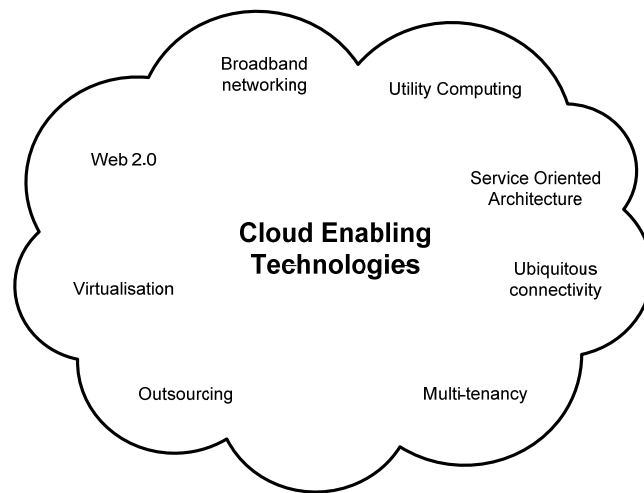


Figure 1 – The enabling and maturing technologies of cloud computing

2.2 Attributes of the cloud

There are differing views on the number and description of the cloud's key attributes. For this Information Security Briefing, the cloud is defined by a minimum of three attributes:

1. **Hardware management is highly abstracted;**
2. **Infrastructure costs are incurred as variable (operating) expense; and**
3. **Infrastructure capacity is elastic (i.e. it can be scaled up or down).**

2.3 Alternative views of the cloud's key attributes

There are alternative definitions of the cloud's key attributes. The US National Institute of Standards and Technology (NIST) define cloud computing with five attributes:⁴

1. **On demand self-service.** A consumer can unilaterally provision computing capabilities such as server time and network storage, as needed without requiring human interaction with each service's provider.
2. **Ubiquitous network access.** Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms such as mobile phones, laptops, and PDAs.
3. **Location independent resource pooling.** The provider's computing resources are pooled to serve all consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand. The customer generally has no control or knowledge over the exact location of the provided resources. Examples of resources include storage, processing, memory, network bandwidth, and virtual machines.

⁴ National Institute of Standards and Technology: NIST Definition of Cloud Computing v15, <http://csrc.nist.gov/groups/SNS/cloud-computing/index.html> (published October 2009)

4. **Rapid elasticity.** Capabilities can be rapidly and elastically provisioned to quickly scale up, and rapidly released to quickly scale down. To the consumer, the capabilities available for rent often appear to be infinite and can be purchased in any quantity at any time.
5. **Pay per use.** Capabilities are charged using a metered, fee-for-service, or advertising based billing model to promote optimisation of resource use. Examples are: measuring the storage, bandwidth, and computing resources consumed, and charging for the number of active user accounts per month. Clouds within an organisation accrue cost between business units and may or may not use actual currency.

Besides NIST, there are a number of other leading organisations that have defined the cloud:

- The University of California at Berkeley ⁵
- The Information Security Forum ⁶
- Forrester ⁷
- O'Reilly

2.4 The delivery models of cloud computing

Cloud computing services are normally delivered in one of four ways, depending on the level of ownership and the technical architecture:

Delivery Model	Description
Vendor cloud (External)	Vendor (or provider) cloud computing services can be accessed across the Internet or a private network, using one or more data centres, shared among multiple customers, with varying degrees of data privacy control. Sometimes called “public” cloud computing.
Private cloud (Internal)	Computing architectures modelled on vendor clouds, yet built, managed and used exclusively by a single enterprise; uses a shared services model with variable usage of a common pool of virtualised computing resources. Data is controlled within the enterprise.
Hybrid cloud	A mix of vendor cloud services, internal cloud computing architectures, and IT infrastructure, forming a hybrid model that uses industry good practice technologies to meet specific needs.
Community cloud	Community clouds are used across organisations that have similar objectives and concerns, allowing for shared infrastructure and services. Community clouds can be deployed using any of the three methods outlined above, simplifying cross-functional IT governance.

Table 1 - Delivery models

⁵ Armburst, Fox, Griffith, Joseph, Katz, Konwinski, Lee, Patterson, Rabkin, Stoica, Zaharia; Above the Clouds: A Berkeley View of Cloud Computing, (University of California at Berkeley, February 2009)

⁶ ISF Briefing: Cloud Computing, www.securityforum.org

⁷ TechRadar™ For Infrastructure & Operations Professionals: Cloud Computing, (Forrester, 2009).

2.5 The services and sub-services of cloud computing

Clouds are commonly described in terms of the functionality offered. The table below provides a summary of the three main types of cloud computing services.

Service Type	Description
Software-as-a-Service (SaaS)	SaaS covers the range of applications that are licensed for use as services provided to customers on demand typically across the Web and it is currently the largest component of the cloud computing market. SaaS predates the recent term cloud computing by several years.
Platform-as-a-Service (PaaS)	The PaaS model makes all of the facilities required to support the complete life cycle of building and delivering web applications and services entirely available from the Internet.
Infrastructure-as-a-Service (IaaS)	IaaS is the delivery of computer infrastructure as a service. Rather than purchasing servers, software, data centre space, or network equipment, customers instead buy those resources as a fully outsourced service.

Table 2 - Cloud services

2.6 Examples of ‘the cloud’

At a simplistic level, cloud computing represents a way to architect and remotely manage computing resources such as database services or end user applications. A simple example of a cloud would be a third party managed email service where the service is rented and is highly elastic (i.e. the service can be scaled up or scaled down). An organisation only has to establish an account with a cloud provider to be able to instantly begin using its services. Cloud services can range from simple software deployment such as web email, or photo sharing, but can also include other types of more sophisticated computing solutions:

- Cloud applications might be interactive web applications. Applications in the cloud might utilise a regional database.
- The cloud may have a web service infrastructure and message queues.
- Cloud applications might need to interoperate with CRM or e-commerce application services, necessitating construction of a custom technology stack to deploy within the cloud if these services are not already provided.
- The cloud might involve new types of long term digital storage technologies that possess improved reliability and resilience capabilities.
- The cloud might include the remote hosting and use of custom or third party software systems.
- The cloud might automatically increase or decrease computing resources as a function of business intelligence about resource demand using automation and virtualisation.

While not all of these capabilities exist in today’s clouds as fully automated solutions, a good number of these *can* be provided. The table below provides some examples of existing cloud offerings, listed according to the type of service and functionality offered.

Software as a Service	Platform as a Service	Infrastructure as a Service
<p>Organisations can access a wide range of applications, operating systems and services. These services frequently support collaborative working and the interlinking of services (mash ups).</p> <ul style="list-style-type: none"> • Zoho • Salesforce.com • Basecamp • Ulteo • Google Apps 	<p>All of the facilities required to support the complete life cycle of building and delivering Web applications and services are entirely available from the Internet.</p> <ul style="list-style-type: none"> • Windows Azure • Google App Engine • Aptana Cloud 	<p>Rather than purchasing servers, software, data centre space, or networking equipment, customers lease those resources as a fully outsourced service.</p> <ul style="list-style-type: none"> • Dropbox • Amazon Web Services • Mozy • Akamai

Table 3 - Examples of cloud services

3. What are the drivers of cloud computing?

Cloud computing has been considered as one of the most hyped IT terms in recent years⁸. Interest has been growing in cloud computing steadily since 2006 and is continuing to gather momentum across the IT industry. Initially, this interest has been vendor driven, but is now being led and influenced by potential customers of this technology. This section examines these market driving forces and the perceived benefits of adopting the cloud.

3.1 Drivers of cloud computing

The pressures to decrease IT costs and increase agility are driving enterprises to consider the adoption of cloud computing services. For small and medium sized organisations in particular, cloud computing can help reduce both capital and revenue expenditure by replacing traditional packaged software and hardware procurements with the purchase of complete IT services which can scale and flex to meet changing business needs.

Driver	Reason
Reduce total IT spend without compromising service quality	<ul style="list-style-type: none">• Current financial climate and budget pressure.• Lower up front capital expenditure costs compared to on-premise solutions. Note: there will be costs associated with data migration to the cloud.• Fewer assets, such as hardware and software licences, on the balance sheet.• Different profile of in-house IT organisation required, potentially at reduced cost.• Costs are treated as operating expense, not capital expenses.
Economies of Scale	<ul style="list-style-type: none">• Small to medium sized organisations using cloud services could realise economies of scale by utilising computing solutions typically found in a larger organisation at a unit price which they could not negotiate on their own.
Gain flexibility and speed in implementations	<ul style="list-style-type: none">• Shift in IT from supporting the infrastructure to providing innovative services for business functions.• Software and hardware maintenance and upgrades will typically be handled by cloud providers.• Bring new users on board without the need for business cases to obtain approval to spend capital and without the lead times associated with hardware purchases.

Table 4 - Cloud drivers

⁸ Gartner, Inc, "Gartner Highlights 27 Technologies in the 2008 Hype Cycle for Emerging Technologies" (2008), <http://www.gartner.com/it/page.jsp?id=739613>

3.2 Benefits of cloud computing

The table below lists the main advantages for each of the cloud delivery models outlined in Chapter 2. What is cloud computing?

Model	Benefit
Vendor Cloud (external)	<ul style="list-style-type: none"> • Quick startup time; no capital investment required. • Allows outsourcing of non-core functions to a service provider. • Leverages a highly scalable provider infrastructure. • Uses a reliable and standardised software stack. • Lower initial fees, variable costs, billed by usage.
Private Cloud (internal)	<ul style="list-style-type: none"> • Quick startup and flexibility of resource allocation; requires capital investment. • On-premise data and systems; allows direct support of governance and compliance, security, data privacy, etc; limited opportunities for reduction of staffing. • Maybe a good choice when possible to leverage existing staff and investments; allows control of service levels and operational reporting. • Cost savings through leveraging virtualisation and more effective use of assets to increase resource utilisation and lower internal costs.
Hybrid Cloud (mixed)	<ul style="list-style-type: none"> • Quick startup, but the integration of vendor and private cloud adds complexity. • Can permit control of data and reduction of non-core focus. • Allows selection of scalable provider infrastructure when needed; can allow internal control when required. • Allows fine-grained sourcing of most appropriate technology and cost profiles; integration may constrain savings potential.
Community Cloud	<ul style="list-style-type: none"> • Sharing service costs between organisations. • Can be architected to permit information sharing between organisations without passing data into external network environments.

Table 5 - Cloud benefits

4. Cloud computing architecture

4.1 Service architectures

Three primary types of cloud service models were introduced earlier these are:

- Software-as-a-Service (SaaS)
- Platform-as-a-Service (PaaS)
- Infrastructure-as-a-Service (IaaS)

There are a specific set of sub-services that describe specialisations of the above cloud computing service models. These sub-services are described in the table below:

Sub-Service Type	Description
IaaS: DataBase-as-a-Service (DBaaS)	DBaaS allows the access and use of a database management system as a service.
PaaS: Storage-as-a-Service (STaaS)	STaaS involves the delivery of data storage as a service, including database-like services, often billed on a utility computing basis, e.g., per gigabyte per month.
SaaS: Communications-as-a-Service (CaaS)	CaaS is the delivery of an enterprise communications solution, such as Voice Over IP, instant messaging, and video conferencing applications as a service.
SaaS: SECURITY-as-a-Service (SECaaS)	SECaaS is the security of business networks and mobile networks through the Internet for events, database, application, transaction, and system incidents.
SaaS: Monitoring-as-a-Service (MaaS)	MaaS refers to the delivery of second-tier infrastructure components, such as log management and asset tracking, as a service.
PaaS: Desktop-as-a-Service (DTaaS)	DTaaS is the decoupling of a user's physical machine from the desktop and software he or she uses to work.
IaaS: Compute Capacity-as-a-Service (CCaaS)	CCaaS is the provision of "raw" computing resource, typically used in the execution of mathematically complex models from either a single "supercomputer" resource or a large number of distributed computing resources where the task performs well.

Table 6 - Cloud sub-services

4.2 Software as a Service (SaaS)

SaaS is a delivery model allowing for on-demand licensing of software services providing a cost-effective alternative across the web. SaaS has the following attributes:

- **Accessibility and reliability:** ability to easily, consistently, and frequently access service offerings within SaaS when required due to the critical use of software services for end users of supported business operations.
- **Standardised IT-based capability:** ability to provide the same quality of service as existing on-site software vendors, such as timely deployment of critical patches, configurability due to multi-tenancy on the cloud.
- **Customer service and enterprise presence:** sustainable market presence within the SaaS service offering and ability to provide customer service comparable to licensed software products, such as SAP and Oracle.

Service offering attributes	Supplier Examples
Accessibility and reliability	Salesforce.com – Salesforce.com continues to build value in its six-tier, user-based pricing model with its premier support delivering 24/7 live phone support and priority phone queues, two-business-hour response time, an assigned customer service representative, application programme interface (API) support, outsourced admin services, and CRM health checks.
Cost	Google, Microsoft – Vendors with large economies of scale for cloud-based infrastructure, such as Google and Microsoft, will drive prices down for software applications, such as office productivity suites and customer relationship management software.
Enterprise presence	Salesforce.com – Salesforce.com is a leading example of the SaaS cloud computing model. It services nearly 40,000 customers around the world and is growing rapidly.

Table 7 - SaaS offering attributes

4.3 Platform as a Service (PaaS)

PaaS is a delivery model that manages the facilities required to support the complete life cycle of building and delivering applications and services from the web.

- **Application base and support:** ability to provide services to support the development life cycle with all the required tools to provide a quality product, including but not limited to, version control, source code control and integration with existing tools.
- **Elasticity and scalability:** ability to provide the same level of service for development processes while quickly allowing upward or downward scalability of services as development teams ramp up or down during phases of the SDLC.
- **Developer affinity, customer service, and enterprise presence:** ability for developers to utilise existing skill sets with various tools available to provide the same quality of service throughout application development; sustainable market presence within PaaS.

Service offering attributes	Supplier Examples
Elasticity and scalability	Google, Microsoft, and salesforce.com promise to deliver highly reliable PaaS services.
Standardised IT-based capability	Google, salesforce.com, Amazon Web Services – One of the key strengths of PaaS offerings is the foundation in data centre architectures pioneered by the likes of Google and Amazon .
Customer service	Google Apps, salesforce.com – Established PaaS products, such as Bungee Connect, Caspio Bridge, Google App Engine, and salesforce.com are providing value to application development and program management teams.

Table 8 - PaaS offering attributes

4.4 Infrastructure as a Service (IaaS)

IaaS is a cloud computing model that facilitates the delivery of computer resources as a service. IaaS enables a customer to buy resources as a fully outsourced service rather than purchasing servers, software, data centre space, or network equipment. It has the advantage of near instantaneous scalability in turn providing a cost-effective and flexible solution. IaaS has the following attributes:

- **Elasticity and scalability:** ability to retain the same level of service while quickly allowing infrastructure components to be scaled upward or downward in a timely fashion throughout the entire software development lifecycle.
- **Standardised IT-based capability:** ability to provide the same quality of service as existing on-site infrastructure, such as communication infrastructure, help desk availability, servers, storage, etc. so that the service levels to the end users are not impacted.
- **Customer service and enterprise presence:** sustainable market presence within IaaS service offerings and the ability to provide more robust infrastructure services without sacrificing service requirements in critical business areas.

Service offering attributes	Supplier Examples
Elasticity and scalability	Amazon Web Services (AWS) – includes the Elastic Computer (EC2), Simple Storage Service (S3), and Simple DB. Akamai – includes scalable solutions for web applications.
Standardised IT-based capability	AWS – is a service provider for end user requisition of computing power, storage, and other services.
Customer service	IBM and HP have an enterprise presence in IT infrastructure, while Amazon and Rackspace are developing their customer service expertise in this area under the new service delivery method.

Table 9 - IaaS offering attributes

5. Cloud computing maturity

This section provides an overview of cloud service provider maturity. In particular, maturity as a set of specific measurable service criteria, the perception of maturity, and the impact of adoption on an organisation are discussed.

5.1 Adoption of cloud computing

The rate at which organisations embrace cloud computing services is linked to the perceived maturity and stability of the cloud services on offer from today's providers. Established enterprise IT vendors and niche cloud providers are jostling to position their cloud computing services' technical and security components and supporting services as best placed for their existing customer base and those interested in exploring the benefits of the cloud. There is commercial pressure on businesses to adopt cloud computing models but customers need to ensure that their cloud services are driven by their own business needs rather than by providers' interests, which will be driven by short-term revenues and sales targets and long-term market share aspirations.

5.2 Maturity of the cloud

For each cloud model there are a set of functional, process and technical maturity levels that can be defined to measure a cloud service's ability to function as required within a given business environment. Most established cloud providers publish the maturity of their own cloud services and this information provides valuable insights for prospective cloud customers, who would otherwise need to carry out a lengthy analysis of the cloud service against their specific requirements.

Charting a migration path to a cloud computing service requires a clear understanding of the maturity and viability of current cloud categories. An indicative representation of maturity during 2009 is shown in Figure 2 - Cloud maturity:

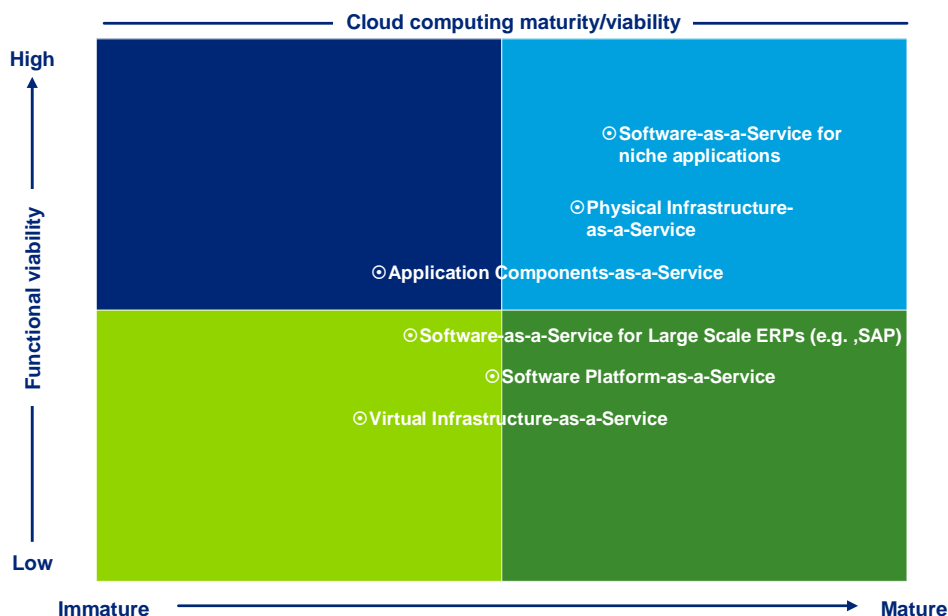


Figure 2 - Cloud maturity⁹

The rapid adoption of cloud computing services has been impacted by concerns over data security, data access, network latency, service levels, provider lock-in, and service availability. The maturity of a cloud computing environment provides adopting organisations with an understanding of the suitability of the cloud service and the level of investment required by the customer in order to address any challenges around security, network latency, performance and so on. Assessing the maturity level of each of these ‘maturity aspects’ is one qualitative approach to gauging the maturity of a given cloud.

Each of the cloud maturity factors has questions associated with it that potential customers ought to consider as part of their adoption process.

The table below provides some high level assessment criteria in determining a cloud’s maturity:¹⁰

Maturity aspect	Assessment of maturity
Functionality	<ul style="list-style-type: none"> • Can the proposed cloud service adequately support the current business model and any expected growth/reduction and change within the business plan?
Security	<ul style="list-style-type: none"> • Can the cloud provider demonstrate relevant security certification with standards such as ISO27001 or PCI DSS given their specific scope?
Availability	<ul style="list-style-type: none"> • Can the cloud service deliver demonstrate acceptable and measurable uptime consistent with the expected trading operations of your business?
Network performance	<ul style="list-style-type: none"> • Does the cloud service provider support adequate network bandwidth and latency to deliver acceptable performance to your users?

⁹ Source: Deloitte Touche Tomhatsu global webinar on cloud computing (July 2009)

¹⁰ Note that this is not an exhaustive list of assessment criteria for determining a cloud’s maturity.

Resilience	<ul style="list-style-type: none"> Does the cloud service provide multiple locations from which it stores data backups and resilient hardware in order to recover from incidents including environmental hazards such as earthquakes or flooding?
Organisational and Financial stability	<ul style="list-style-type: none"> Does the provider have a sound history of cloud service delivery? What is their financial position? Are they a likely target for acquisition/merger? Is the provider's security culture aligned to your own?
Service Level Agreements (SLAs)	<ul style="list-style-type: none"> Does the cloud service provider give a comprehensive SLA regarding the service, including specific security elements? What is the provider's historical track record of achievement against this or similar SLAs for other customers?

Table 10 - Maturity levels

Other industry bodies and research organisations have proposed alternative ways of assessing a cloud's maturity level. According to a recent Jericho publication, there are several "cloud formations" – or forms of cloud computing which can be used to determine the maturity of a cloud.¹¹ Each cloud formation resembles different characteristics including degrees of flexibility, different collaborative opportunities and risks. A cloud's maturity can be distinguished by four criteria. First, whether the cloud is outsourced or in-sourced; second, whether the cloud is perimeterised¹² or de-perimeterised¹³; third, whether the cloud is open or proprietary; and fourth; whether the cloud is external or internal.

5.3 Vendor maturity and impacts on adoption

The maturity of a cloud service is also characterised by the level of adoption associated with that service. There are three stages of maturity that can define the level of adoption of a cloud service:¹⁴

- **Technology pilots** – Pilot services will not provide a fully functional service and therefore cannot be considered for use within existing production business environments. There will be many risks that prohibit adoption of the technology;

¹¹ Jericho Forum, Cloud Cube Model: Selecting Cloud Formations for Secure Collaboration

¹² According to the Jericho Forum, "Perimeterised implies continuing to operate within the traditional IT perimeter, often signalled by "network firewalls"... "In effect, when operating in the perimeterised areas, you may simply extend your own organisation's perimeter into the external cloud computing domain using a VPN and operating the virtual server in your own IP domain, making use of your own directory services to control access. Then, when the computing task is completed you can withdraw your perimeter back to its original traditional position. We consider this type of system perimeter to be a traditional, though virtual, perimeter."

¹³ According to the Jericho Forum, "De-perimeterised, assumed that the system perimeter is architected following the principles outlined in the Jericho Forum's Commandments and Collaboration Oriented Architectures Framework. The terms Micro-Perimeterisation and Macro-Perimeterisation will likely be in active use here – for example in a de-perimeterised frame the data would be encapsulated with meta-data and mechanisms that would protect the data from inappropriate usage. COA-enabled systems allow secure collaboration. In a de-perimeterised environment an organisation can collaborate securely with selected parties (business partner, customer, supplier, outworker) globally over any COA capable network."

¹⁴ Whilst these characteristics indicate the level of maturity of a cloud service, the rate of adoption and customer confidence in that service is not necessarily attributed to these.

however, this does not withstand the potential for future development into an acceptable business implementation. Therefore, it is likely that pilots will be limited to small non-production services or concept driven developments.

- **Early adopters** – A service which has established itself as providing a useful level of functionality will often be adopted by those with immediate challenges that cloud computing services appear to address. For example businesses with an imminent and expensive infrastructure requirement may wish to reduce costs by adopting cloud computing. However a lack of clarity around maturity of other factors such as security and legislative requirements mean that early adopters are at risk of deploying a service that does not yet meet all their requirements
- **Stable technology** – Stability of a service is a major sign of maturity, however, in its own right does not make a cloud computing service immediately suitable for deployment. The impact of stability can lead organisations to miscalculate the risks when planning implementation. The fact that a service is stable does not mean that it answers all the functional, security and legislative requirements around cloud computing. A stable service should, however, be able to provide more formal SLAs and better integration to your existing systems.

6. Evolution of cloud computing

This chapter examines the history of cloud computing from the 1950's, and the work done by AT&T in the area of telephone networking and the evolution of these technologies which are present in today's cloud implementations.

6.1 History

Indeed the cloud is not as new as it seems. The cloud symbol that permeates virtually all cloud computing literature is more than 50 years old, as indeed are the concepts that were recognised as early as the 1950s in the work done by AT&T in the area of telephone networking. At that time, AT&T had already begun to develop an architecture and system where data would be located centrally and accessed by business through redesigned telephones and updated telephone network. While the service did not materialise the concepts and advantages were understood and relentlessly pursued through to this day.

The pursuit of centralised, abstracted IT services progressed over the decades with the advent and adoption of technologies such as Internet Service Providers (ISP – where servers were located at the Internet access point) and Application Service/Infrastructure Providers (ASP – where infrastructure was rented to a customer at an offsite location, but used most of the time by the one, paying customer). Other IT services historically offered include Time Sharing Systems, Co-Location, Hosting, and Outsourcing.

As with any evolution, the step from ASP to cloud computing is subtle yet disruptively important. While ASPs managed to offsite infrastructure for a customer, they were bound to the concept that the infrastructure capacity was predetermined and inflexible; ASP customers were required to declare the quantity of compute and storage capacity needed upfront. If the customer's computing needs grew or contracted the hardware had to be scaled up or down with an associated delay and up-front investment.

One of the main principles of cloud computing, from Software-as-a-Service to Storage on demand, is that the computing capacity varies immediately and transparently with the customer's need, and clients no longer need plan, configure, and use fixed quantities of computing equipment, reducing associated costs, lead-times, and financial risks.

6.2 Evolution of cloud technologies

From a computing standpoint, many of the technologies and technical concepts of the cloud can be traced back. All of the following types of computing technologies and architectures share similarities with the benefits and architecture present in today's cloud implementations:

1. **Mainframe and thin client computing** – Mainframe computing is a highly reliable, powerful, centrally located form of computing service. A user of a mainframe system may access applications using a thin client; a specialist terminal for users to interact with and operate a mainframe system. These classic 'green screen' thin client interfaces were the first instances of client-server style computing. Mainframe computing is still widely deployed today and is an effective standard for providing businesses with reliable and large scale processing power. The advantages of mainframes and modern Unix systems are also applicable to modern cloud computing architectures:
 - a. **Resilient highly available architectures** – Each mainframe system is designed to run at a high level of utilisation without failure, and to support hardware upgrading whilst still in operation.
 - b. **Mainframes can host multiple operating system instances** – Each mainframe can effectively provide virtual instances of operating systems and application environments. This is a crucial requirement for supporting scalability within cloud computing.
 - c. **Grid and supercomputing** – The development of high powered computers and large scale parallel, or grid, computing has been driven by the need for number crunching processing power. The use of specialist supercomputers, or large numbers of computers configured to run in parallel in a 'grid' permits the operators to model and solve complex problems such as predicting the weather or decrypting data encrypted with strong encryption algorithms. Systems designed for these purposes are generally expensive and designed for a particular purposes (i.e. highly targeted), though this is not always the case.
2. **Scalability and on demand processing power** – The use of a supercomputer or grid computing service provides a level of scalability to those needing resources that may be too cost prohibitive to purchase in house. However, there may be higher risks in establishing an internal grid system based on using spare CPU cycles. Capacities of existing systems would need to be closely monitored to avoid potential impact on existing services. The processing power within these facilities can be shared and provided to multiple users concurrently to execute complex software programs, which cannot use traditional computing infrastructure. One salient example of this is the SETI¹⁵ (Search for Extraterrestrial Intelligence) project, a scientific research programme involving the use of thousands of individual users' computer systems to form a single distributed computing environment with increased processing power.

¹⁵ <http://setiathome.ssl.berkeley.edu/>

3. **Utility computing** – Computing services that can be metered and billed to customers in the same way that electricity or telephony systems operate, are known as utility computing services. Utility computing services offer a commercial and multipurpose computing platform for high volume and scalable computing services and are a yet another precursor to modern cloud computing services. The concept of utility computing is also associated with the commercialisation of problem solving in supercomputing systems.
4. **Eliminating capital expenditure costs** – A utility service absolves its customers of investment in high cost hardware. This model was popular before computing hardware costs were lowered and achieved mass market availability. As the cost of acquiring, managing and supporting computing facilities is considered to be high, outsourcing technology operations has re-emerged as a popular way of managing an efficient business.
5. **The Internet and worldwide web (WWW)** – The invention of web pages accessible by remote computers was initially part of a scientific research facility used to share information. The WWW drove consumer interest in the Internet, caused by growth in the home PC industry, and, improvements in network technology implementation resulting in greater bandwidth availability. The concept of freely and globally accessible information is an attractive proposition to the public and business users, which has provided wide ranging social benefits such as high speed, global communications and knowledge sharing and education.
6. **Harnessing improved network technologies** – Developments in Wide Area Network (WAN) technologies and the improved access to websites has enabled the advances in information and eCommerce services experienced today. The network is the connector to the cloud and improvements in this technology are a primary requirement for increasing the uptake of cloud computing as a concept.
7. **Enabling global and enterprise accessibility** – The principal of availability and accessibility enabled by the network and hardware infrastructure is a crucial element of cloud computing, whether it be vendor or private clouds. The infrastructure and architecture behind large scale Internet sites that appeared in the late 1990's which manage a high volume of traffic with minimal downtime was the main precursor to interactive next generation websites and present day vendor cloud computing innovations.

8. **Dedicated cloud operating systems** – The purpose of the operating system with respect to developing an application for a PC is to provide a standardised, supported and testable environment that enables developers to use tools to quickly create an application. The cloud operating system provides the same environment and services from which to architect and run cloud based applications.

The history of cloud computing has shown that organisations tend to build computing services from the ground up and focus their efforts on their own data centres using a selection of hardware and software components that they must manage. The concept of the cloud as a pre-existing standard environment is not something most industries or businesses have fully embraced. However, the benefits of a common standard are likely to emerge over time, enabling both small and large organisations to adopt a top-down approach in the adoption of cloud based services.

9. **Web 2.0 and cloud computing operating systems** – The global presence of the Internet and the introduction of wireless networking and mobile devices featuring always on Internet connectivity has raised expectations of users and demand for services over the internet. Social networking sites, video and voice communications, and location based services are part of everyday life and Web 2.0 is the label applied to the interactive Internet. It is arguably a layer of new technologies built upon the existing foundations of the web.

However, the architectures required by service providers to enable Web 2.0 has created an IT service that is differentiated by resilience, scalability, reusability, interoperability, security and open platform development. This has effectively become the backbone of cloud computing and is considered by a number of vendors and services to be an operating system layer of its own.

7. Risks of cloud computing

7.1 Purpose and aim of section

This section of the briefing introduces the risks associated with cloud computing. The following topics are discussed:

- Business risks
- Security in the cloud
- Reliability and resilience
- Usability and performance
- Regulations and legislation
- Organisational change

Following each section above there is a summary table containing mitigating advice for the risk discussed. No overall conclusion is drawn following all of the topics discussed as the landscape of each topic is in flux. An executive summary of the risks discussed throughout this section is provided below in Table 11 – Cloud Risks.

There are two additional sections that follow the discussion on risks. These are on security testing and the future of cloud computing.

7.2 Overview of risks

Although cloud computing is portrayed as a generally valuable consideration for enterprise IT integration, adoption of cloud computing models carry a number of risks. The table below provides a summary of these many of which are discussed in upcoming sections of this report:

Risks	Description of risk
Availability: Service Availability and Recoverability	<ul style="list-style-type: none">• Cloud providers may not be able to match in-house IT service availability, Recovery Time Objectives (RTO), and Recovery Point Objectives (RPO).• Cloud providers may drastically change their business model or discontinue cloud services, impacting customers.
Availability: Complexity	<ul style="list-style-type: none">• Complexity introduced by a cloud computing environment can result in more components that must be managed, and more complex recovery procedures.
Availability: Single-Points-of-Failure	<ul style="list-style-type: none">• Even if the cloud environment is architecturally designed for high-availability, single-points-of-failure may exist in the access path to the cloud.

Risks	Description of risk
Availability: Data Replication	<ul style="list-style-type: none"> Due to technical architecture complexity, and, potentially, restrictions by the cloud provider, replicating data back to the customers' enterprise or to another provider may be difficult.
Availability: Testing constraints	<ul style="list-style-type: none"> Due to concerns about confidentiality and impact on other customers, cloud providers may place heavy constraints on disaster recovery testing activities.
Availability: Over-subscription Risk	<ul style="list-style-type: none"> In the event of an incident, other customers may receive higher priority in recovery activities. As cloud providers shift from investment mode to capture market share, to cost cutting mode, to reach profitability, capability may become constrained.
Access: Multi-tenancy	<ul style="list-style-type: none"> Data is possibly exposed to third parties due to a lack of granular access controls in the cloud, potentially allowing unauthenticated parties access to confidential data.
Access: Data access	<ul style="list-style-type: none"> Data may be stored in the cloud without proper customer segregation allowing possible accidental or malicious disclosure to third parties.
Access: Secure Data Deletion	<ul style="list-style-type: none"> Customer data that was required to be deleted may still be retained on backup servers or storage located in the cloud without customers' knowledge.
Authentication: External Authentication	<ul style="list-style-type: none"> Where ownership and maintenance of credential repositories is the responsibility of an external party, security good practices cannot be guaranteed without SLAs.
Authentication: Federated Authentication	<ul style="list-style-type: none"> Organisations may implement Single Sign On (SSO) applications used by multiple business partners but the SSO may also grant access to sensitive internal information if configured incorrectly and without any monitoring.
Authentication: Key Management	<ul style="list-style-type: none"> Any activity related to key generation, exchange, storage, safeguarding, use, vetting, and replacement that results in disclosure will provide access to infrastructure and data.
Authentication: Cloud to Cloud Authentication	<ul style="list-style-type: none"> One cloud provider may rely on a second cloud provider to authenticate a user's identity based on the first cloud passing a Security Assertion Markup Language (SAML) assertion to the second cloud at the request of a user. Based strictly on the assertion, the second cloud provider may grant the user access to cloud resources. Incorrectly implemented SAML assertions can be susceptible to the following attacks: DoS, Man-in-the-Middle, Replay, and Session Hijacking.
Regulatory: Audit Rights	<ul style="list-style-type: none"> Customers may have no or limited rights to perform audits, and review performance against contracts or SLAs.

Risks	Description of risk
Regulatory: Compliance	<ul style="list-style-type: none"> • Migration to the cloud can infer a more complex regulatory environment for some customer businesses.
Regulatory: Certification	<ul style="list-style-type: none"> • The scope of certifications such as PCI DSS and ISO27001 may be increased to consider parts of the cloud provider infrastructure that cannot be removed from the certification scope.
Integrity: Shared Environments	<ul style="list-style-type: none"> • Where customer data in the cloud is in a shared environment alongside data from other customers, additional security testing may be required to prevent data corruption.
Integrity: Data Monitoring	<ul style="list-style-type: none"> • Changes to customer data without the knowledge of the data owners may be caused by interoperability issues with the cloud provider's data storage component technologies.
Integrity: Data Encryption	<ul style="list-style-type: none"> • Data at rest (if not encrypted) accessed by third parties due to faulty access controls is subject to loss of integrity.
Privacy: Legal uncertainties	<ul style="list-style-type: none"> • Multiple jurisdictions increase regulatory complexity. • Conflicting legal provisions can create significant uncertainty in assessing compliance and risk. • The Privacy and Data Protection legal landscape continues to evolve at a rapid pace. • Data sharing agreements may be required before moving data to the cloud. • Business associate agreements may need to be considered (HIPAA). • Data controllers and third parties may need to be considered (EU DPD).
Privacy: Individual Rights/Confidentiality	<ul style="list-style-type: none"> • Strict terms of service are particularly important in the cloud to preserve individual privacy/confidentiality and to meet regulatory requirements to which the customer is subject. • The cloud facilitates the ability to use/share data across organisations and therefore increase secondary uses of data that may require additional consent/authorisation. • Data is easily accessible by a larger group of users and must be strictly controlled (Protect data at rest).
Privacy: Breach/Disclosure	<ul style="list-style-type: none"> • Centralised data stores may be especially prone to security breaches. • Timely discovery and reporting of a breach by the cloud provider may be challenging.
Operational Security: Vulnerability Management	<ul style="list-style-type: none"> • One security vulnerability on the right component has the potential to exposure large numbers of corporations' critical assets.
Operational Security: Asset Management	<ul style="list-style-type: none"> • Assets in the cloud may not be managed to an adequate standard and could leak critical company information or cause data exposures.

Risks	Description of risk
Operational Security: Incident Response	<ul style="list-style-type: none"> Ownership, responsibilities, and actions during incident response are not well defined.
Operational Security: Security Management	<ul style="list-style-type: none"> A complete information security management system may not be defined between the cloud provider and customer. Security testing is critical in testing the integrity of the cloud service.

Table 11 – Cloud Risks

An additional source of discussion on risks from cloud computing can be found in the Cloud Computing Risk Assessment¹⁶ from European Network and Information Security (ENISA).

¹⁶ <http://www.enisa.europa.eu/act/rm/files/deliverables/cloud-computing-risk-assessment>

8. Business risks

8.1 Overview of business risks

There is a wide range of perceived business risks associated with the adoption of cloud computing technologies. These typically fall into the categories of:

- Technical risk
- Commercial risk
- Legal/compliance risk
- Operational risk

This chapter provides an overview of the perceived business risks associated with each type of cloud computing.

8.2 Business risks associated with vendor or public clouds

8.2.1 Lack of standards and interoperability

Experience of the dotcom boom has shown how new computing concepts are a magnet for investment. The first to market with a new technology solution can bring large financial rewards, but it also creates a focus on short term wins rather than sustainable, long-term market solutions. Brand new market technologies also expose customers to supply chain risks as there may only be a minority of organisations developing or supporting the technology. There are risks associated with service failure, data loss and vendor lock-in and whilst there are no standards governing cloud implementations, these are likely to be inhibitors to the widespread adoption of cloud computing solutions.

As the marketplace matures, it is likely that provider standards will emerge that will improve operational integrity; enhance service agreements; and mitigate the risks around provider lock-in.

8.2.2 Shared computing resources and segregating data

Complex and geographically dispersed supply chains also present organisations with challenges in understanding which other organisations and individuals have access to their data and infrastructure. As with traditional outsourcing arrangements, cloud customers need to understand which organisations have access to their data for the duration of its lifecycle, from creation through to secure destruction. One of the key advantages of the vendor cloud computing model is the ability of customers to maximise the use of available resources through the sharing of computing resources.

However, there is an increased risk that customer data could be accessed by other customers sharing the cloud's services. Segregation controls are needed to ensure that access to data within the cloud is properly managed.

8.2.3 Legal and regulatory risks

There are a wide range of legal and regulatory issues associated with the adoption of vendor clouds including subcontracting, rights to data and vendor lock-in. Since cloud providers may be utilising computing resources in foreign jurisdictions, data in the cloud may be subject to the laws and regulations of another country. These laws may conflict with the cloud customer's legal or regulatory obligations in their home country. These risks are discussed in more detail in Chapter 12.0 Regulations and legislation

Case Study:

Earlier this year a small Internet start-up, providing PaaS to clients wishing to create custom online database applications, filed for bankruptcy blaming the current tough economic climate.

A major shareholder promptly stepped in to acquire the remaining assets with the intention of limiting the services to internal use only.

In this instance*, customers were able to access their source data before the service was terminated. However, since the service was built on a proprietary platform, the source data did not enable the customers to continue to use their applications, and they were unable to move to another provider without rewriting all their applications.

* <http://blogs.zdnet.com/SAAS/?p=668>

8.2.4 Challenges of undertaking due-diligence in vendor clouds

There are a number of additional challenges in carrying out due-diligence on vendor cloud providers. Customers will be faced with increased costs and complexity:

- **Increased cost of due-diligence** – Cloud computing architectures are inherently more geographically dispersed than traditional outsourcing models. Due-diligence enquiries on cloud providers could feasibly involve several third party providers located in several different jurisdictions, dramatically increasing costs incurred should on-site visits be required.
- **Complexity of due-diligence activities** – Given the complexities inherent in cloud architectures and in particular, the outsourcing of virtual technologies and use of virtualisation technologies; there is a risk that due-diligence enquiries will be a greater cost for customers, potentially involving a number of third parties from several jurisdictions. Risks may be difficult to identify and/or quantify.

8.3 Private clouds

Private clouds allow organisations to retain greater control of data and supplier choices and, therefore, provide a greater degree of control of risk. Subject to the size of customer, private clouds may not, however, offer the same cost benefits as vendor clouds.

Organisations considering implementing their own private clouds need to consider carefully the risks and benefits. The benefits of implementing a private cloud should be clear from experience of maintaining an existing, in-house service model.

8.4 Hybrid clouds

The primary reason for utilising the services of a hybrid cloud is to realise the collective benefits of vendor clouds, private clouds and traditional IT services. Hybrid clouds offer flexibility in choosing a combination of cloud and in-house services suited to business needs. This is particularly beneficial where only some services are suitable to be provided by a cloud, allowing the business to retain direct control over the remaining services and assets. Hybrid cloud risks include the management of internal and external change control processes.

8.5 Community clouds

From an organisational perspective, a community cloud could be viewed as a shared service private cloud which brings together non-competing organisations with a common interest and risk appetite. Community cloud users share a common interest of seeking to exploit economies of scale whilst minimising the costs of adopting a private cloud or the risks associated with vendor clouds. A management organisation would normally be contracted to oversee the operation of a community cloud.

Mitigating advice

Issue	Description of mitigating advice
<p>Cost of due-diligence</p>	<p>The level to which due-diligence enquiries are undertaken should be determined by the value of the contract and the level of risk the customer is exposing itself to by entering into a contract with the provider</p>
<p>Due-diligence and risk assessment</p>	<p>Due-diligence enquiries should consider, as a minimum:</p> <ul style="list-style-type: none"> • Whether contracts give customers the right to audit • Whether the security environment meets customers' security standards and scope of requirements, covering: <ul style="list-style-type: none"> ○ Segregation controls for shared hosting/IT resource ○ Security of data in transit ○ Security testing ○ External certifications and accreditations such as ISO27001, AAF or SAS70 • Whether resilience IT DR meets the customer's IT DR requirements. • Whether there are potential conflicts between the regulatory and legal obligations of the customer and the provider and the risks involved by engaging with that provider. • Security and legal obligations of storing or processing data in an offshore jurisdiction. • Consider the financial risks of engaging with a provider to ensure that security is not compromised.

Contracts and SLA's	<ul style="list-style-type: none">• Robust contracts should be in place with cloud providers. These should have additional emphasis on:<ul style="list-style-type: none">• Data protection and security;• Data controls and ownership;• Geographic and jurisdictional constraints of the services;• Support to enable exit management; and SLAs.
----------------------------	---

Table 12 - Mitigation advice

9. Security in the cloud

Security in the cloud is achieved, in part, through third party controls and assurance much like in traditional outsourcing arrangements. But since there is no common cloud computing security standard there are additional challenges associated with this. Many cloud vendors implement their own proprietary standards and security technologies, and implement differing security models, which need to be evaluated on their own merits. In a vendor cloud model, it is ultimately down to adopting customer organisations to ensure that security in the cloud meets their own security policies through requirements gathering, provider risk assessments, due-diligence, and assurance activities.

Thus, the security challenges faced by organisations wishing to use cloud services are not radically different from those dependent on their own in house managed enterprises. The same internal and external threats are present and require risk mitigation or risk acceptance. This section examines the information security challenges that adopting organisations will need to consider, either through assurance activities on the vendor or public cloud providers or directly, through designing and implementing security controls in a privately owned cloud. In particular, this chapter examines:

- The threats against information assets residing in cloud computing environments.
- The types of attackers and their capability of attacking the cloud.
- The security risks associated with the cloud, and where relevant consideration of attacks and countermeasures.
- Emerging cloud security risks.
- Example cloud security incidents.

9.1 Cloud threats

The threats to information assets residing in the cloud can vary according to the cloud delivery models used by cloud user organisations. The table below provides an overview of threats for cloud customers categorised according to the Confidentiality, Integrity and Availability (CIA) security model and their relevance to each of the cloud delivery models.

Threat	Description
Confidentiality	
Insider user threats: <ul style="list-style-type: none"> Malicious cloud provider user Malicious cloud customer user Malicious third party user (supporting either the cloud provider or customer organisations) 	The threat of insiders accessing customer data held within the cloud is greater as each of the delivery models can introduce the need for multiple internal users: SaaS – Cloud customer and provider administrators PaaS – Application developers and test environment managers IaaS – Third party platform consultants
External attacker threats: <ul style="list-style-type: none"> Remote software attack of cloud infrastructure Remote software attack of cloud applications Remote hardware attack against the cloud Remote software and hardware attack against cloud user organisations' endpoint software and hardware Social engineering of cloud provider users, and cloud customer users 	The threat from external attackers may be perceived to apply more to public Internet facing clouds, however all types of cloud delivery model are affected by external attackers, particularly in private clouds where user endpoints can be targeted. Cloud providers with large data stores holding credit card details, personal information and sensitive government or intellectual property, will be subjected to attacks from groups, with significant resources, attempting to retrieve data. This includes the threat of hardware attack, social engineering and supply chain attacks by dedicated attackers.
Data Leakage: <ul style="list-style-type: none"> Failure of security access rights across multiple domains Failure of electronic and physical transport systems for cloud data and backups 	A threat from widespread data leakage amongst many, potentially competitor organisations, using the same cloud provider could be caused by human error or faulty hardware that will lead to information compromise.
Integrity	
Data segregation: <ul style="list-style-type: none"> Incorrectly defined security perimeters Incorrect configuration of virtual machines and hypervisors 	The integrity of data within complex cloud hosting environments such as SaaS configured to share computing resource amongst customers could provide a threat against data integrity if system resources

Threat	Description
	are not effectively segregated.
User access: <ul style="list-style-type: none"> • Poor identity and access management procedures 	Implementation of poor access control procedures creates many threat opportunities, for example that disgruntled ex-employees of cloud provider organisations maintain remote access to administer customer cloud services, and can cause intentional damage to their data sources.
Data quality: <ul style="list-style-type: none"> • Introduction of faulty application or infrastructure components 	The threat of impact to data quality is increased as cloud providers host many customers' data. The introduction of a faulty or mis-configured component required by another cloud user could potentially impact the integrity of data for other cloud users sharing infrastructure.
Availability	
Change management: <ul style="list-style-type: none"> • Customer penetration testing impacting other cloud customers • Infrastructure changes upon cloud provider, customer and third party systems impacting cloud customers 	As the cloud provider has increasing responsibility for change management within all cloud delivery models, there is a threat that changes could introduce negative effects. These could be caused by software or hardware changes to existing cloud services.
Denial of Service threat: <ul style="list-style-type: none"> • Network bandwidth distributed denial of service • Network DNS denial of service • Application and data denial of service 	The threat of denial of service against available cloud computing resource is generally an external threat against public cloud services. However the threat can impact all cloud service models as external and internal threat agents could introduce application or hardware components that cause a denial of service.
Physical disruption: <ul style="list-style-type: none"> • Disruption of cloud provider IT services through physical access • Disruption of cloud customer IT services through physical access • Disruption to third party WAN providers services 	The threat of disruption to cloud services caused by physical access is different between large cloud service providers and their customers. These providers should be experienced in securing large data centre facilities and have considered resilience among other availability strategies. There is a threat that cloud user infrastructure can be physically disrupted more easily whether by insiders or externally where less secure office environments or remote working is standard practise.

Threat	Description
Exploiting weak recovery procedures: <ul style="list-style-type: none"> • Invocation of inadequate disaster recovery or business continuity processes 	The threat of inadequate recovery and incident management procedures being initiated is heightened when cloud users consider recovery of their own in house systems in parallel with those managed by third party cloud service providers. If these procedures are not tested then the impact upon recovery time may be significant.

Table 13 - Cloud security threats

The table above provides the basis for the discussion of the security threats and countermeasures throughout the following sections.

9.2 Types of attackers

Many of the security threats and challenges in cloud computing will be familiar to organisations managing in house infrastructure and those involved in traditional outsourcing models. Each of the cloud computing service delivery models' threats result from the attackers that can be divided into two groups:

Internal Attacker	<ul style="list-style-type: none"> • An internal attacker has the following characteristics: <ul style="list-style-type: none"> ○ Is employed by the cloud service provider, customer or other third party provider organisation supporting the operation of a cloud service ○ May have existing authorised access to cloud services, customer data or supporting infrastructure and applications, depending on their organisational role ○ Uses existing privileges to gain further access or support third parties in executing attacks against the confidentiality, integrity and availability of information within the cloud service.
External attacker	<ul style="list-style-type: none"> • An external attacker has the following characteristics: <ul style="list-style-type: none"> ○ Is not employed by the cloud service provider, customer or other third party provider organisation supporting the operation of a cloud service ○ Has no authorised access to cloud services, customer data or supporting infrastructure and applications ○ Exploits technical, operational, process and social engineering vulnerabilities to attack a cloud service provider, customer or third party supporting organisation to gain further access to propagate attacks against the confidentiality, integrity and availability of information within the cloud service.

Table 14 - Cloud attackers

Although internal and external attackers can be clearly differentiated, their capability to execute successful attacks is what differentiates them as a threat to customers and vendors alike.

For the purposes of this briefing attackers have been categorised into four types. Each of these categories is based on ability to instigate a successful attack, rather than on the type of threat they present (i.e. criminal, espionage or terrorism):

- **Random** – the most common type of attacker uses simple tools and techniques. The attacker may randomly scan the Internet trying to find vulnerable computers. They will deploy well known tools or techniques that should be easily detected.
- **Weak** – semi-skilled attackers targeting specific servers / cloud providers by customising existing publicly available tools for specific targets. Their methods are more advanced as they attempt to customise their attacks using available exploit tools.
- **Strong** – organised, well financed and skilled groups of attackers with an internal hierarchy specialising in targeting particular applications and users of the cloud. Generally this group will be an organised crime group specialising in large scale attacks.
- **Substantial** – motivated, strong attackers not easily detected by the organisations they attack, or even by the relevant law enforcement and investigative organisations specialising in eCrime or cyber security. Mitigating this threat requires greater intelligence on attacks and specialist resources in response to detection of an incident or threat.

9.3 Security risks

The security risks associated with each cloud delivery model vary and are dependent on a wide range of factors including the sensitivity of information assets, cloud architectures and security controls involved in a particular cloud environment. The following sections discuss these risks in a general context, except where a specific reference to the cloud delivery model is made.

The table below summarises the security risks relevant in the cloud:

Risk	Description
Privileged user access	Cloud providers generally have unlimited access to user data, controls are needed to address the risk of privileged user access leading to compromised customer data.
Data location and segregation	Customers may not know where their data is being stored and there may be a risk of data being stored alongside other customers' information.
Data disposal	Cloud data deletion and disposal is a risk, particularly where hardware is dynamically issued to customers based on their needs. The risk of data not being deleted from data stores, backups and physical media during decommissioning is

Risk	Description
	enhanced within the cloud.
e-Investigations and Protective monitoring	The ability for cloud customers to invoke their own electronic investigations procedures within the cloud can be limited by the delivery model in use, and the access and complexity of the cloud architecture. Customers cannot effectively deploy monitoring systems on infrastructure they do not own; they must rely on the systems in use by the cloud service provider to support investigations.
Assuring cloud security	Customers cannot easily assure the security of systems that they do not directly control without using SLAs and having the right to audit security controls within their agreement.

Table 15 - Cloud security risks

9.3.1 Privileged user access

Once data is stored in the cloud, the provider has access to that data and also controls access to that data by other entities (including other users of the cloud and other third party suppliers). Maintaining confidentiality of data in the cloud and limiting privileged user access can be achieved by at least one of two approaches by the data owner: first, encryption of the data prior to entry into the cloud to separate the ability to store the data from the ability to make use of it; and second, legally enforcing the requirements of the cloud provider through contractual obligations and assurance mechanisms to ensure that confidentiality of the data is maintained to required standards. The cloud provider must have demonstrable security access control policies and technical solutions in place that prevent privilege escalation by standard users, enable auditing of user actions, and support the segregation of duties principle for privileged users in order to prevent and detect malicious insider activity.

Encryption of data prior to entry into the cloud poses two challenges. For encryption of data to be an effective means of maintaining data confidentiality, decryption keys must be segregated securely from the cloud environment to ensure that only an authorised party can decrypt data. This could be achieved by storing keys on segregated systems in house or by storing keys with a second provider.

An additional challenge around encryption in the cloud is to prevent manipulations of encrypted data such that plain text, or any other meaningful data, can be recovered and be used to break the cipher. This constraint in encryption technology¹⁷ means that cloud providers must not be granted unlimited ability to store and archive encrypted data. If the cloud user organisation permits the cloud service provider to handle unencrypted data, then the cloud service provider must provide assurance that the data will be protected from unauthorised access, both internally and externally. Within the cloud, the generation and use of cryptographic keys for each cloud customer could be used to provide another level of

¹⁷ Homomorphic encryption schemes are a means of alleviating this constraint. They permit defined manipulation of the plaintext *without* needing to decrypt the ciphertext, and therefore they can be used to maintain a segregation between computations applied to encrypted data and access to the plaintext. At the time of writing, no practical homomorphic encryption scheme exists, although Gentry proposed a scheme which satisfies the requirements for a fully homomorphic scheme under certain conditions in 2009. [Gentry, Fully homomorphic encryption using ideal lattices; *41st ACM Symposium on Theory of Computing 2009*:169-178, http://domino.research.ibm.com/comm/research_projects.nsf/pages/security.homoenc.html]

protection above and beyond data segregation controls. However, providers need robust key management processes in place and the challenge for customers then becomes gaining assurances over that process.

A strong or substantial attacker could exploit weak encryption policies, and privileged cloud provider management access, to recover customer data using a complex software or hardware attack on user endpoint devices, or cloud infrastructure devices. This attack may involve long term compromise of the cloud provider supply chain, or social engineering of a particular cloud customer user.

The use of encryption technology may also be subject to limitations or specific requirements depending on the jurisdiction in which the cloud provider will be storing cloud customers' data. For example in some countries the use of encryption technologies may be restricted based upon the type of encryption or its purpose of operation. Cloud customers should review whether the application of encryption as mandated by the local jurisdiction of the cloud provider is acceptable and does not enhance risk to their data.

For example in the UK the Regulatory Investigatory Powers Act¹⁸ (RIPA) can impose a legal obligation to disclose encryption keys to enable access to data by security and law enforcement agencies. Cloud customers should ensure that they understand their obligations within all of the jurisdictions used by the cloud provider, and have policies and procedures in place to deal with specific external enquiries with respect to encrypted data.

9.3.2 Data location and segregation

Data location and data segregation are of particular importance in the cloud given the disparate physical location of data and shared computing resource. Cloud users may be under statutory, regulatory or contractual obligations to ensure that data is held, processed and managed in a certain way. There are a number of associated security risks:

- The cloud provider being required to disclose data (and potentially decryption keys) or hand over physical media to a third party or statutory authority;
- Development of liabilities to pay tax to local authorities as a result of processing sales or other transactions within their jurisdictions;
- Environmental hazards such as earthquakes, flooding, and extreme weather affecting the security of customer data; and
- Macro-economic hazards such as hyper-inflation or deep recession affecting the providers' services and personnel.

Central storage arrangements in cloud computing also provide attackers with a far richer target of information. In a single attack, attackers could potentially gain access to confidential information belonging to several customer organisations. If adequate segregation of data isn't applied many customers may find themselves suffering a security breach due an incident that should have been limited to a single customer.

¹⁸ <http://security.homeoffice.gov.uk/ripa/encryption/>

Virtualisation is one of a number of enabling technologies of cloud computing that itself is a run-time method of segregation for processing data. Many of the security concerns and issues associated with virtualisation are relevant in cloud computing, regardless of whether or not the cloud service provider employs virtualisation technologies. Security of data depends on having adequate security controls in each of the layers of the virtualised environment. In addition, secure deletion of memory and storage must be used to prevent data loss in a multi-tenant environment where systems are reused.

The hypervisor layer between the hardware and virtual machine / guest OS has privileged access to layers above. It also has a great deal of control over hardware, and increasingly so, as hardware manufacturers implement hypervisor functions directly into chipsets and CPUs. Cloud users therefore need to assess cloud service providers' use and operation of virtualisation technologies and whether the risk profile can be tolerated. This topic is discussed in more detail in the CPNI technical paper on virtualisation.¹⁹

9.3.3 Data disposal

Cloud services that offer data storage typically provide either guarantees or service-level objectives around high availability of that data. Cloud providers achieve this by keeping multiple copies of the data. Where the cloud customer has a requirement to delete data, cloud-based storage may be inappropriate for that data at all points in its lifecycle.

Depending on the type of data hosted in the cloud, customers may require providers to delete data in accordance with industry standards such as the UK, HMG Information Assurance Standard 5, or in the United States, NIST Special Publication 800-88²⁰. Unless the cloud architecture specifically limits the media on which the data may be stored and the data owner can mandate use of media sanitisation techniques on that media in line with the required standards, customers may need to preclude their data from being transmitted in the cloud.

9.3.4 e-Investigations and protective monitoring

Implementing protective monitoring in the cloud presents challenges for both cloud customers and providers given the disparate location of physical data and the high number of providers involved. Whilst cloud enabling technologies are designed to place a security perimeter between the cloud service systems and the cloud users, vulnerabilities in this layer of security cannot be ruled out altogether. There is a risk of insider threats and attacks on the cloud and this is likely to require expertise in e-Investigations and protective monitoring.

Effective protective monitoring of cloud-based information assets is likely to require integration between monitoring tools employed by the cloud provider as well as tools employed by the cloud user. Tracing actions back to accountable users and administrators in the cloud may require an integrated or federated (mutual trust) identity management and associated logging system which permits unambiguous identification of all authorised individuals with access to the cloud resources.

¹⁹ <http://www.cpni.gov.uk/Docs/tn-01-09-security-server-virtualisation.pdf>

²⁰ http://csrc.nist.gov/publications/nistpubs/800-88/NISTSP800-88_rev1.pdf

Managing identity and access in the cloud for an enterprise is likely to require integration with a pre-existing identity management system. Federating the cloud customer's identity management system with the cloud provider's identity and access management system is one solution.

Protective monitoring of the cloud will, in certain cases, depend on cloud customers' ability to trace actions back to all authorised identities in both the cloud and customer IT environment. This is likely to require a federated identity management approach which encompasses the cloud users as well as the cloud service provider.

The technology supporting federated identity management is currently in its very early stages, with several competing standards vying for dominance in a landscape of numerous proprietary identity management technologies.

Access to accurate information is, of course, vital in investigating incidents. Having access to data within protective monitoring logging systems, and the ability to carry out forensic investigations on computing devices and other infrastructure within a cloud environment may be difficult for cloud customers pursuing an investigation. Therefore customers should address this issue within their contractual agreements with providers, and understand how their provider implements protective monitoring within their cloud environment. Customers placing specific requirements within a contract relating to investigations will need to consider how their investigation team integrates with their equivalent investigations team within the cloud provider organisation. This is of particular importance where investigations are taking place on multi-tenant systems and providers have a responsibility to protect other customers' data. Generally collecting digital evidence within the cloud should be the responsibility of the cloud provider, and it should be handed over as part of the chain of custody of evidence to the customer for their own investigation purposes. Should customers request more direct access to specific data devices that are part of a shared customer infrastructure, then the provider may choose to change the architecture of that customer's service which may substantially increase the costs to the customer and may impact the original business case for choosing cloud services.

9.4 Assessing the security of a third party cloud provider

One of the most significant challenges for vendor cloud customers in particular is assurance over the security controls of their cloud provider. This is exacerbated by the fact that there is currently no common industry cloud computing security standard from which customers can benchmark their providers. Customers are primarily concerned with:

- **Defining security requirements** – The customers' information security requirements are derived from the organisation's own policy, legal and regulatory obligations, and may carry through from other contracts or SLAs that that company has with its customers.
- **Due diligence on cloud service providers** – Prospective cloud customers should undertake proper due-diligence on providers before entering into a formal relationship. Detailed due-diligence investigations can provide an unbiased and valuable insight into a providers' past track record, including its financial status, legal action taken against the organisation and its commercial reputation. Certification schemes such as ISO27001 also provide customers with some assurances that a cloud provider has taken certain steps in its management of information security risks.

- **Managing cloud supplier risks** – The outsourcing of key services to the cloud may require customer organisations to seek new and more mature approaches to risk management and accountability. Whilst cloud computing means that services are outsourced, the risk remains with the customer and it is therefore in the customer's interest to ensure that risks are appropriately managed according to their risk appetite. Effective risk management also requires maturity both in vendor relationship management processes and operational security processes.

9.5 Emerging cloud security threats

The following examples give some insight into the emerging security threats that are being detected and researched by academics, security companies and both cloud providers and customers alike.

- **Side channel attacks** – An emerging concern for cloud delivery models using virtualisation platforms is the risk of side channel attacks causing data leakage across co-resident virtual machine instances. This risk is evolving, though currently is considered to be in its infancy, as the virtual machine technologies mature. However, it is possible that attackers who fail to compromise endpoints or penetrate cloud infrastructure from outside the cloud perimeter, may consider this technique. Acting as a rogue customer within a shared cloud infrastructure to access other customers' data²¹.
- **Denial of service attacks** – Availability is a primary concern to cloud customers and as such it is equally of concern to providers who must design solutions to mitigate this threat. Traditionally denial of service has been associated with network layer distributed attacks flooding infrastructure with excessive traffic in order to cause critical components to fail or to consume all available hardware resources. Within a multi-tenant cloud infrastructure there are more specific threats associated with denial of service, these include:
 - Shared resource consumption – attacks that deprive other customers of system resources such as thread execution time, memory, storage requests and network interfaces can cause a targeted denial of service.
 - Virtual machine and hypervisor exploitation – attacks that exploit vulnerabilities in the underlying hypervisor, or operating system hosting a virtual machine instance will allow attackers to cause targeted outages or instability.

Attacks using these methods are designed to circumvent traditionally well defended cloud architectures that have concentrated on securing against external network based denial of service attacks.

²¹ <http://people.csail.mit.edu/tromer/papers/cloudsec.pdf>

- **Social networking attacks** – With the increased popularity of business and personal social networking sites the risk of advanced social engineering attacks is increased. Cloud computing systems are targeted due to their large customer data stores. The complex set of relationships between cloud providers, customers, suppliers and vendors means that many employees of these organisations will be listed on social networking sites and be connected to each other. Attackers can setup identities to gain trust, and use online information to determine relationships and roles of staff to prepare their attacks. A combination of technical attacks and social engineering attacks can be deployed against a target user by taking advantage of the people they know and the online social networks they use.
- **Mobile device attacks** – The use of smartphones has increased and cloud connectivity is now no longer limited to laptop or desktop computing devices. Attacks are now emerging that are targeted for mobile devices and rely on features traditionally associated with laptops and desktops, including:
 - rich application programming interfaces (APIs) that support network communications and background services;
 - always on wireless Internet access; and
 - large local data storage capabilities.

As mobile devices now have these equivalent features Internet based spyware, worms or even physical attacks may be more likely to occur against mobile devices, as they are potentially a less risky target to an attacker that wishes to remain undetected. This is generally supported by the fact that most mobile devices do not have the equivalent security features enabled, or in some cases available. For example mature anti-malware, anti-virus or full disk encryption technologies are not widespread on smartphones.

- **Insider and organised crime threat** – Cloud providers will store a range of different data types, including credit card and other financial and personal data. All of this data may be aggregated from multiple customers and therefore be extremely valuable to criminals. There is a risk that insiders are deliberately used to gain access to customer data and probe systems in order to assist any external attackers that require additional information in order to execute complex Internet based attacks. Cloud customers should ensure that providers are aware of this threat and have rigorous identity validation and security vetting procedures built into their recruitment process.

9.6 Examples of cloud security incidents

There have been a number of publicly disclosed security incidents in the cloud. Some high profile public examples of security incidents include:

- **Availability** – a major search and online advertising organisation was forced to make an embarrassing apology in February 2009 when its email service collapsed in Europe. Similar incidents recurred on a smaller scale several times during 2009. There were several unexpected repercussions of this service outage, believed to have been a

- **Confidentiality** – a mature cloud services provider was targeted by a phishing attack in November 2007. An employee fell victim to this phishing attack which captured his company login credentials. The attackers then used the credentials to harvest confidential customer contact data and sent phishing emails to the organisation's customers in the form of fake sales invoices.
- **Availability** – In February 2008, a cloud storage service went down for almost four hours, wreaking havoc on several companies that use and depend on the cloud. The provider described the cause as an unexpected spike in customer transactions, the security of the service was compromised due to a lack of availability of the data for cloud customers.

9.7 Mitigating advice

Mitigating advice	Description
Assurances over security in the cloud	<ul style="list-style-type: none"> • Define the organisation's security requirements and the security requirements applicable to information assets in the cloud. • Carry out appropriate due diligence investigations of cloud providers. • Manage risks associated with using cloud services.
Undertaking threat assessments in the cloud	<ul style="list-style-type: none"> • Assess the threats to information assets in the cloud. • Assess the threats to the cloud providers, including threats aimed at other cloud users.
Encryption of data	<ul style="list-style-type: none"> • If data must be kept confidential from the cloud provider: <ul style="list-style-type: none"> ○ It must be encrypted prior to entering the cloud ○ Keys must be managed separately from the cloud provider who has access to the data. • Review whether there are specific legal limitations or requirements upon the use of encryption in the local jurisdiction of the cloud provider.
Assess cloud provider's controls for segregating data between cloud users.	<ul style="list-style-type: none"> • Encryption can protect data stored in the cloud if the cloud provider can demonstrate robust key management and security. • Assess the cloud provider's mechanisms for segregating data when unencrypted (for example, when being processed).
Assess whether data disposal can be assured in the cloud.	<ul style="list-style-type: none"> • Secure deletion of data is usually not feasible. • Determine whether strong encryption under a key which is subsequently securely deleted is an acceptable substitute for secure deletion.

Mitigating advice	Description
<p>Identify all geographic locations for data storage, processing and transfer</p>	<ul style="list-style-type: none"> • Statutory and regulatory obligations on the cloud user and cloud provider will be dependent on geographic location of data and computing facilities. • Other risks (environmental, political, economic) will also vary by geographic location.
<p>Assess effectiveness of protective monitoring of the cloud</p>	<ul style="list-style-type: none"> • Tracing actions back to individual users may be impossible without tight integration of the cloud user's and cloud provider's protective monitoring. • Assess the cloud provider's capability and willingness to protectively monitor information assets in the cloud.
<p>Assess identity management and interoperability concerns</p>	<ul style="list-style-type: none"> • Standards are immature and do not have widespread acceptance. • Manage the risks around becoming locked-in to proprietary technologies and suppliers.

Table 16 - Security mitigating advice

10. Reliability and resilience

10.1 Overview of resilience issues

Cloud computing provides an additional option for business continuity planners. It is, in many respects, simply an outsourcing solution for either production and/or recovery of technology by a third party provider. It therefore must be evaluated as a recovery option and tested in the same way as all other solutions as stipulated in both BS 25999 the business continuity management standard and BS 25777 the ICT continuity standard. As is common with all outsourcing arrangements, operational aspects of the service may be provided by the third party provider but the risk remains with the business.

10.2 Benefits of cloud computing to continuity planners

One of the benefits of cloud computing is that it makes sophisticated technology solutions and purpose built data centres with their resilient infrastructure and trained staff available and affordable to small and medium sized companies. This may be a significant improvement to the resilience of many companies, provided that the risks surrounding the service are identified and mitigated.

However, access to a purpose built data centre is not a push button resilience solution and there is a danger that it may be seen as this. Business continuity planners will need to define the recovery requirements and test them as they do with all solutions. Small and medium sized businesses with limited IT budgets may look at cloud computing as a far more resilient solution than they currently own and may move to cloud computing on the basis that it increases their resilience from nothing to something. While this may be worthwhile for those with nothing currently in place, the risks, from a continuity perspective, should be understood first.

10.3 Systemic and specific risks

One of the systemic risks in cloud computing is the increased load on data centres. Cloud computing emerged as a result of high levels of redundancy in existing data centres. Data centre owners recognised that the spare capacity in their facilities could be used and sold to others without significantly impacting their resilience. However the business continuity planner should consider the implications for resilience. As soon as the work load of the data centre is increased, built in redundancy is inevitably reduced.

There are four specific areas of resilience risk that should be considered:

- **Single site.** If a single data centre is being used to provide services, which is unlikely in a cloud, the risks of that site being unavailable need to be firmly understood. Data centre failure is not necessarily improbable and business continuity planners should take this into account when evaluating the cloud as an option.
- **Shared risk.** Increasing the utilisation of the data centre increases the workload on the infrastructure and people. Specifically, the higher the utilisation the higher the demand will be on power, air-conditioning, bandwidth and staff. Recovering services may take longer due to the competing demands on resources by multiple tenants.
- **Contagion.** Multi-tenanted data centres run the risk of corruption. Issues that develop for one cloud customer may/can spill over and impact others, tying up infrastructure or staff, and damage caused to shared hardware by one cloud customer can impact another.
- **Contention.** A spike in usage by one or several cloud customers, for example, as a result of successful marketing, can tie up bandwidth or stress infrastructure impacting all users.

10.4 Delivering resilience in the cloud

Cloud computing offers an opportunity to simplify production and improve recovery times for business continuity planners. Cloud based computing may be a cost effective route but this will depend on the specific requirements of the business and the ability of the provider to meet those needs.

However, businesses still need to define their service level and recovery requirements, and this remains the essential first step. Once defined, the options, including cloud based computing models should be evaluated and solutions identified. The recovery solutions when incorporating cloud computing need to consider the additional complications of managing recovery through a third party provider. Incident management including invocation/activation of recovery plans and escalation procedures need to be clearly understood and practised regularly. Recovery processes are likely to be more involved and will require a high degree of coordination, communication and speedy, informed decision making if they are to be effective.

10.5 Delivering resilience through testing

A robust IT disaster recovery plan validated through testing remains essential. To achieve resilience providers must understand and be able to demonstrate through thorough testing how the components of their cloud interact and actively control its configuration. Without this, a cloud customer can have no assurance that the resilience levels they require can be met.

It is therefore critical for the continuity planner to evaluate the risks surrounding the cloud service and either accepts or mitigates them. Specific areas to focus on when evaluating cloud computing services are:

- **Single or multiple site hosting.** Understanding the risk of single points of failure for example a single site failure against the capability and complexity of a multi-site solution.
- **Infrastructure.** Confirm the power, cooling, fire suppression, communications links, environmental factors, and physical security suitable for the business requirement.
- **Segregation.** To what extent is hardware shared (single points of failure) and what is the risk of failure caused by other third party providers? Understanding the risk of a single large platform shared by multiple parties against that of multiple smaller platforms and the speed of recovery for both options.
- **Service level agreements.** Usage, agreed service levels, data back-up, disaster recovery and testing in the cloud should not be assumed. The requirements should be defined by the customer and written in to service level agreements. Testing of the solution is vital. A data centre may fail over along with supporting storage area network within a matter of minutes but if each customer requires a manual re-start there may be significant time delay involved in a full recovery. This will only become apparent through testing.

10.6 Mitigating advice

Mitigating advice	Description
Defining ITDR requirements	<ul style="list-style-type: none"> • Define the business requirements. • Ensure the technical architecture supporting the cloud provides you with the required resilience level. • Do not be reliant on the SLA - have a tested BCM and IT DR plan in place. • Cloud computing is based on a shared operating environment so ensure your data is safe. Identify who else might have access to your data? Your cloud provider will have system administrators who will most likely have access and this point should be factored into the decision for moving to a cloud based service operating model. • Do your homework and know what the full list of options is being provided by the cloud provider. • Data corruption – how many copies of your data does the third party provider have and how will they take steps to protect your data? What type of backups will be used, incremental or full and how long could it take to reconstruct your data?
Achieving effective ITDR solutions	<ul style="list-style-type: none"> • Test the solution regularly.

Table 17 - Resilience mitigating advice

11. Usability and performance

Organisations choosing to use cloud services must specify requirements for their cloud applications and infrastructure that satisfy the needs of the users executing business processes. Users must be able to operate systems within reasonable tolerance levels, which may include responsiveness of the user interface, time taken to execute operations and overall availability of cloud applications. Cloud user organisations will have an expectation of system performance gained from experience in managing their own in-house systems and dealing with third party suppliers for hardware, software and telecommunications services.

These expectations for modern network-aware software applications exist whether those systems are being hosted in house, or remotely from a private or vendor cloud. This section describes the specific issues and challenges associated with accessing cloud-based services over Wide Area Networks (WANs), including the Internet, which could impact service usability and performance in user operational environments.

11.1 Latency

By its nature the cloud is a physically separate entity often geographically distant from the IT systems and end users it provides services to. This distance, and the extent of IP switching between two remote points, could be a major factor in the performance of cloud computing when integrated in existing business environments that require fast and responsive systems.

Definition of latency

“The elapsed time between sending a data packet across a network to a remote client and the time it is received.”

If the systems and networks are not designed and configured correctly cloud computing can increase latency to an unacceptable duration which may have a detrimental impact on a given service. Effectively latency issues are caused by the distributed nature of cloud computing involving WANs with data traversing across several physical transmission technologies and geographical boundaries including terrestrial and celestial links. Latency is also affected by network capacity, delays caused by network congestion and packet prioritisation resulting from traffic shaping techniques (Quality of Service (QoS) for example). Over a large network where a data packet transits across multiple links, the potential latency from source to destination is defined as the sum of the minimum latency for each link in the route across the network.

Latency can also be incurred by poor network design, hardware faults and software bugs. Common issues with latency include:

- Applications with a high bandwidth requirement will suffer the most as large numbers of delayed data packets may cause bottlenecks in parts of the network.
- Voice and video communications can be disrupted by data packet delays.
- Services can fail if latency causes timeouts, or is outside the tolerance level of a hardware component or software service.

- Cloud computing services will be heavily dependent on the latency of the network, and this should be considered before moving any critical applications to a hosted cloud model.

11.2 Reducing latency

The impact of latency can be mitigated depending on the architecture of the cloud service in question. Considering the use of a cloud for the following example services requires thought to be given to reducing latency:

- **Data Access** – For cloud services that host high volumes of application data such as that used by design and publishing tools, or video editing applications, latency will be caused by the time taken to transfer large files across limited network bandwidth. Avoiding this latency can be achieved through deploying caching servers to buffer the content stored within the cloud locally, removing delays. This principle also additionally supports maintaining a local backup copy of cloud data, so is an interesting consideration as it improves usability and availability.
- **Video-streaming** – use of QoS techniques to prioritise traffic over available bandwidth. The cloud service featuring video streaming must be able to enable scalability per customer when high bandwidth applications are required within a SLA.
- **Security** – Introducing cryptographic hardware support, such as securing remote access using SSL hardware accelerator equipment. The cloud service provider can deploy hardware to offload the computing capacity required for secure remote access, removing the burden from the existing server infrastructure.

11.3 Network access

Accessing systems across a network is a huge benefit for users sharing resources and modern Local Area Networks (LANs) cope well with the general business tasks that they are designed for. Once connected, the network (both LAN and WAN) is an enabler to many solutions including storage and transfer of data, and communications and remote working. This is in contrast to traditional standalone systems that rely on removable media devices for data transfer, which cannot communicate with other systems and require a user to be physically present at the machine to operate it.

As network access is of the utmost importance to users of cloud computing, losing the network is an extremely undesirable event. Immediately, systems may be unable to access data, data that was being accessed may now be incomplete, and in some cases work will be lost with no means of recovery. For example where there is communication occurring between a customer relationship management system and users, the impact of loss of the network may impact business critical services, or directly impact customer service satisfaction.

Addressing this issue and maintaining fully resilient network access is a difficult task for cloud providers to achieve within their own infrastructure. However, as cloud users are external, they connect through networks not under the control of the cloud service provider. Therefore despite the provider having implemented and tested their incident recovery procedures, this does not preclude a network component beyond their perimeter failing and introducing a

temporary loss of connectivity to their service. Users of cloud services should understand the resilience available in their own network infrastructure and also in any network infrastructure between the user and service provider networks, as a pre-requisite to implementing cloud services.




Figure 3 - Network boundaries

Figure 3 - Network boundaries shows the networks that both the cloud user and provider organisations must monitor to ensure the performance of cloud services. For each of these networks in house and third party providers may be involved in meeting service agreements upon which the availability and performance on networks are met. Resilience needs to be included within the overall design to minimise SPOF's (single points of failure).

For example if using SaaS for email hosting, and connectivity to the cloud hosting service is lost, consider having a business continuity and disaster recovery strategy. This may be as simple as manually changing mail server settings to revert to a fallback mail server, until the cloud service resumes. New services such as Gears²² from Google are beginning to consider the need for offline support to web application services. Offline accessibility is critical for the continuity of operations in the event of a cloud service outage.

The loss of a network connection would also impact specialist intelligent network services that may be used to support organisational helpdesks, and sales and marketing operations that are critical to the business. For example if a hosted service is providing the telephony network for a sales centre, losing the automated greeting and service selection function may impact the sales function until the issue is resolved.

Losing a network connection therefore may not always mean a total loss of service; however it always almost does imply a negative impact in terms of financial loss and possible reputation damage until the issue is resolved.

11.4 Network availability

Experiencing network access loss is a problem for cloud user organisations; however, a benefit of the cloud is being able to remotely access resources within the cloud from any location that provides adequate network coverage. When planning the use of a cloud service for remote workers (including home workers) to be accessible whilst out of the work premises and without a loss of service, consider the end point systems to be used and the network connectivity that will be relied upon for accessing the cloud, and in which geographical locations they will be required to operate.

²² http://code.google.com/apis/gears/gears_faq.html#whatIsGears

The cloud is only as accessible as the network technologies that are used to provide a connection to it.

11.5 Network performance

For cloud providers hosting infrastructure, platform and software services they must specify their network performance capability in detail along with the specific cloud services they offer. Users should consider their network bandwidth requirements and verify that their cloud provider implements sufficient service levels in order to reserve resource capacity on a communications link, in this case specifically for cloud service traffic. Bandwidth bursting, when total available bandwidth exceeds a threshold, may also be necessary to maintain service. This can be applied at a granular level by cloud providers to protect their customers based on level of service that they have purchased and the criticality of their cloud applications.

QoS should also be implemented by cloud users to guarantee their own access to external services. For example general company Internet usage by employees could at peak times occasionally cause saturation of available bandwidth impacting the users of cloud services. Priority should be given to certain types of traffic, or bandwidth should be increased during peak times only, to reduce costs.

11.6 Monitoring of network performance

The importance of network access, performance and reliability has been highlighted throughout this section. It is therefore important to monitor core network services to maintain usability by the business. This duty to monitor the network extends to both the cloud services provider and the cloud services customer.

By monitoring access, usage, administrative and fault logs, an assessment of the network performance can be made based on a set of predefined metrics. The tolerances of these metrics and resulting actions required in the event of a breach of tolerance should be defined using information derived from pilot implementations and from data provided by the cloud provider on historical performance. A full network monitoring design, including real time alarm capabilities, can then be factored into a production implementation of cloud computing. The act of monitoring the network must also be accompanied by an incident management and recovery plan. This will be used to define investigation and restoration of services within a controlled process.

A communications plan will also be required both internally, to advise on executing recovery and fallback strategies, and externally where customers experience downtime attempting to access their cloud services.

11.7 Mitigation advice

The following key areas relating to usability and performance should be taken into consideration when reviewing the use of business applications over WANs:

Mitigation Advice	Description
Latency	<ul style="list-style-type: none"> • Include regular latency testing as an element of due diligence when considering cloud services; • Understand the network latency within the LAN environment and test the bandwidth, latency and jitter requirements of applications that may be proposed for relocation to the cloud; • Define the security requirements for your cloud service and then determine any additional infrastructure and resultant latency.
Quality of Service (QoS)	<ul style="list-style-type: none"> • Understand the minimum network bandwidth requirements for cloud services and stress test this repeatedly to establish a QoS baseline; • Define a granular bandwidth burst level on the cloud network link to ensure other non cloud traffic does not impact services; • Verify with cloud providers their approach to QoS for different customers and service levels for shared infrastructure.
Resilience	<ul style="list-style-type: none"> • Consider the approach to resilience of the LAN and Internet gateway connectivity and develop and test a BCDR plan; • Internet connectivity has an expected failure rate, consider use of dedicated WAN links or an approach to hosting local fallback services in the event of an extended period of network access loss.
Availability	<ul style="list-style-type: none"> • Ensure that service level agreements are in place for each part of the network that has responsibility for providing access to cloud services. This includes making sure availability is part of the agreement with the cloud services provider. • Understand where and how cloud users will typically interact with the cloud: <ul style="list-style-type: none"> • Do the end point devices have sufficient capability to connect users to the cloud across their geographic areas of operation?
Monitoring	<ul style="list-style-type: none"> • Monitoring the network health can enable preventative maintenance to occur, reducing the likelihood of downtime; • Ensure monitoring is accompanied by appropriate incident and recovery procedures and communications plans.
Backup your data outside of the cloud	<ul style="list-style-type: none"> • Where reducing latency involves storing local data copies, investigate how this could additionally fit in with both backup and resilience policies.

Table 18 – Usability and performance mitigation advice

12. Regulations and legislation

There are a wide range of laws and regulations that cloud customers need to be concerned with in cloud computing. There are particular legal and regulatory issues around the rights to data, security loopholes, outsourcing and subcontracting. The laws and regulations governing the processing of personal or regulated data by third parties will continue to present challenges in vendor cloud arrangements by virtue of the fact that services in the vendor cloud are abstract and outsourced to a third party provider.

Since cloud environments will introduce more complex supply chains and deliver services from across multiple jurisdictions, the complexity of legal arrangements is set to increase. This will be particularly the case outside of the European Economic Area (EEA) where there have been few attempt to harmonise various national and EU laws.

This chapter highlights some of the main legal and regulatory challenges around service abstraction and the geographical boundaries of the cloud, with some specific citation to laws and real world legal scenarios.

12.1 Overview of regulatory and legislation issues

Cloud computing consumers face many of the same challenges as traditional computing environments, such as in the processing and protection of personal, commercial or government data. However, cloud computing also introduces a number of new challenges that need to be considered. The table below outlines the variety of legal and regulatory challenges that organisations have been tackling, with varying degrees of success:

Issue	Description of Issue
Legal uncertainties	<ul style="list-style-type: none">• The cloud's loosely defined, uncertain or moving geography means that consumers are faced with increased legal complexity, legal contradictions and uncertainty. The laws and regulations affecting cloud computing consumers are evolving. Certainly the privacy and data protection laws are evolving at a rapid pace.
Individual Rights and Confidentiality	<ul style="list-style-type: none">• Data may physically reside in a legal jurisdiction where the rights of data subject conflict with European principles, or are not protected at all.• The cloud facilitates the ability to use/share data across organisations and therefore increases the potential for secondary uses of data that require additional consent or authorisation.• Data in the cloud is accessible by a larger group of users and must be strictly controlled.

Breach and disclosure	<ul style="list-style-type: none"> • Centralised data stores in the cloud increase the impact of any security breach. • The timely discovery, assessment, and reporting of the breaches from within the cloud are more challenging if there are a number of providers involved.
Third party compliance	<ul style="list-style-type: none"> • Dispersion of cloud resources across a wide geographical area means that due-diligence and risk assessments on third parties are likely to be more costly and resource intensive. • Long chains of dependencies within the cloud (through several outsourcing arrangements) will also have cost implications on due-diligence and risk assessment activities. • Service providers may be willing to be flexible, but the legal responsibilities of their own third parties and customers may be strictly defined.

Table 19 - Legal and regulatory issues

12.2 Rights to data

12.2.1 Legal and regulatory issues outside of the EU

Under the UK Data Protection Act, organisations subject to UK laws need control over not only where their information is *stored*, but which jurisdictions it is *communicated across*, and accessed from, including infrastructure which may form part of an extended or transient network. In distributed cloud environments, there may be routine processes to copy data to services across the globe. Whilst this may offer advantages in terms of availability and speed, it is by no means certain that the national and federal laws of the territory in which the data physically resides protects the rights and freedoms of data subjects as it is required in European law.

There are, however, mechanisms in place to allow the transfer of personal data outside of the European Economic Area (EEA), such as the EU/US Safe Harbor Agreement. However there are also concerns that local legal or political regimes may enable governmental agencies access to data (for example the Patriot Act in the United States can be used to enforce the disclosure of UK personal data, possibly contravening UK data protection laws). There are other risks too in transferring personal data outside of the EEA; in countries which lack the rule of law, or where the judiciary is not independent, the disclosure of personal data may be unenforceable.

12.2.2 Legal and regulatory issues within the EU

In the EU, certain member countries have passed national laws permitting the lawful interception of any traffic being transferred across its territory. The Cybercrime treaty²³ is a basis for many laws within the EU on computer based interception. Cloud customers may find that by entering into a contract with a cloud provider communicating data across these countries in which they are in breach of their own national laws, particularly if personal data is concerned.

²³ <http://conventions.coe.int/Treaty/en/Treaties/html/185.htm>

As stated above, all EU member states have similar legislation. The Directive assumes that each person has the right to control his or her personal data. This concept is, however, unrecognised in many other parts of the world. The Swedish government passed a law on June 8th, 2008, enabling its government to intercept all Internet traffic transited through its territory. Even in the UK, the UK Regulation of Investigatory Powers Act 2000 gives public servants the right to obtain a warrant to access data stored on computers in the UK if this is necessary, for example, for “the interests of the economic well-being of the United Kingdom”, or to prevent or investigate a crime.²⁴ It was reported that ‘French government officials have also been forbidden to use Blackberry email devices because Blackberries send and receive email using a small number of servers in the US and United Kingdom, and the French security service feared that this might cause a threat to national security because of a risk of data interception’.²⁵

12.3 Outsourcing contracts

The idea that cloud computing abstracts customers from the security of their data in the cloud is simply wrong. In cloud computing, there is a greater need for customers to pursue SLA’s and other contractual agreements to ensure that their legal and obligations are being met and risk is reduced.

This report discusses a wide range of potential risks that customers need to mitigate through the use of contractual agreements. Yet the current user agreements for most of the prominent cloud computing services give no such assurances. According to the terms of service for Google Apps, for example, the services might be interrupted; untimely, insecure, erroneous, give inaccurate or untimely results, but Google and partners would have no liability to you.²⁶ Amazon Web Services state in their general customer agreement that “We will have no liability to you for any unauthorized access or use, corruption, deletion, destruction or loss of any of Your Content or Applications.”²⁷ These disclaimers provide little assurance for organisations seeking the benefits of vendor clouds. There is also the added problem of cloud providers outsource their services, is likely to increase legal complexity and muddy the risk, particularly if those service providers are registered in another offshore jurisdiction.

12.4 Outsourcing, subcontracting and the FSA

The FSA’s 2008 report entitled Data Security in Financial Services is an example of what a leading regulator sees as good practice when it comes to sharing data with third parties. In its report, the FSA wrote that financial firms are relying “*too much on assumptions that contractual terms are being met, with very few firms proactively checking how third parties vet their employees or the security arrangements in place to protect customer data.*” These

²⁴ Regulation of Investigatory Power Act 2000, Part II, s 28.

²⁵ Miranda Mowbray, The Fog over the Grimpen Mire: Cloud Computing and the Law (Scripted, April 2009).

²⁶ http://www.google.com/apps/intl/en/terms/standard_terms.html

²⁷ <http://aws.amazon.com/agreement/>

scenarios may be relevant to the cloud. In the report, the FSA also underlined some of the challenges around the methods by which some firms are transferring data to and from third parties. The FSA reported that in some cases, firms were using CDs or mainframe cartridges which were unencrypted.²⁸

In its 2008 report, the FSA also recommended that organisations should undertake proper due-diligence and regular audits and risks assessments of their third party suppliers. Whilst the FSA have given no specific criteria for assessing the risk of engaging with cloud providers, much of the FSA advice on third party relationships can be neatly carried over. In particular, the FSA best practices cover data segregation, specific risks associated with storing personal data offshore, risk management and accountability. However, cloud customers should expect a tightening up of regulations over third party assurances required of cloud customers, as cloud computing is adopted more widely in the market.

12.5 Processing personal data in the cloud

Organisations in the UK that process personal data, whether collecting, storing, using or disclosing it, are subject to UK and EU data protection laws and regulations. The same holds true for cloud providers; they are no different to any other type of third party provider in this respect, although the considerations that they face may be more complex.

UK data protection regulation is primarily based on the Data Protection Act 1998 (“the Data Protection Act”), which itself is based on an EU Directive (its full title being Directive 95/46/EC, on the protection of individuals with regard to the processing of personal data). There is similar legislation in each of the 27 EU member states, with some variation in each country. A number of non-EU states have similar legislative frameworks (including Switzerland, Canada, Argentina, Australia and Dubai), again with variation between each one.

The Data Protection Act places obligations on the organisation who decides how the personal data is processed (known in the Data Protection Act as the ‘data controller’), and the liability for compliance remains with the data controller, even if other suppliers are used for the processing of the data.

The data controller therefore has to be very clear in their instructions to any processor of their data, including cloud providers, and has to ensure that they manage their risk through having the right contracts and operational controls in place over their suppliers. This covers all of the principles of the Data Protection Act, including requirements around transparency of the use of personal data, the purposes for which it is used, the accuracy of the data and the relevance of the data processed. Of particular note is principle 7 of the Act²⁹, which requires businesses holding personal data to ensure “*appropriate technical and organisational measures... [are] ...taken against unauthorised or unlawful processing of personal data*” and “*measures against the accidental loss or destruction of, or damage to, personal data*”. In a cloud environment, this holds true, but the complexity of where data is stored may make it more difficult to gain the appropriate assurances over the security of the data.

²⁸ Data Security in Financial Services: Firms’ controls to prevent data loss by their employees and third-party suppliers

²⁹ UK Data Protection Act: http://www.opsi.gov.uk/acts/acts1998/ukpga_19980029_en_9#sch1-pt1

Cloud customers also need to be sure that data retention requirements are being followed by the data processor, and that disclosure procedures to raise breaches with the data controller in a timely fashion have been put in place. Where clouds diverge from more traditional hosting arrangements is that there could be a chain of dependencies in the provision of a cloud service. This increases the risk that it may not be possible to quickly escalate a breach at the end of the chain to the data owner, which may breach reporting requirements in some jurisdictions, as well as increase the damage caused to affected individuals and the overall reputational damage if the breach is publicised.

12.6 Mitigation advice

There are a number of steps that organisations can take to mitigate the increased risks that a cloud environment may present:

Mitigation advice	Description
<p>Legal Uncertainties</p>	<ul style="list-style-type: none"> Increased levels of due-diligence on cloud providers to ensure that the laws governing the rights to data in other jurisdictions in which customer data resides or is communicated across (including transient infrastructure) does not conflict with European principles (upheld in UK law) on the rights of the data subject. Specify the right to audit provider systems and infrastructure to ensure that appropriate technical and organisational measures have been carried out against accidental loss, destruction of, or damage to, personal data. Specify contracts with providers (and <i>data processors</i>) appointed to collect, store or destroy the personal data on its behalf and to ensure compliance with data protection laws and regulations.
<p>Individual Rights and Confidentiality</p>	<ul style="list-style-type: none"> Strict terms of service to preserve individual privacy/confidentiality and to meet regulatory requirements. Data sharing agreements within other users of the cloud to ensure that the rights of the individual are protected. Regular risks assessments on third party providers that focus on data segregation controls and multi-tenancy.
<p>Breach and Disclosure</p>	<ul style="list-style-type: none"> Contractual provisions to ensure that the provider is required to give appropriate notification to the customer if IT resources are being relocated to a jurisdiction, outside of the agreed listed of territories. Increased levels of due-diligence on cloud providers' ability to assess and report breaches to the consumer in a timely manner. Breach notification should be contractually defined.
<p>Third party compliance</p>	<ul style="list-style-type: none"> Third party risk assessments. Contractual agreements provisions to allow the cloud customer to protect the consumer from changes and how its data is hosted, managed or supported in order for the customer to meet compliance obligations.

Table 20 - Legal and regulatory mitigating advice

13. Organisational change

An organisation's transition from in-house managed computing to a cloud computing model will bring about significant changes at all levels. Organisations may need to form shared service and security approaches that may be unable to fit easily within current structures. Successful transition to the cloud requires organisations to embrace change, in their processes, policies, roles and responsibilities; this is a challenge in itself. This section outlines some of the challenges and risks associated with an organisation's transition to using cloud services.

13.1 Organisational change management

The adoption of cloud computing will have a significant impact on the way in which businesses operate. As cloud computing becomes prevalent within an organisation, there will be a diminishing need for internal IT infrastructure as well as the resources that manage that infrastructure. Organisations need to manage this change effectively to minimise the impact and disruption to normal business operations. In general, organisations that fail in executing a change strategy effectively have overlooked one or a number of organisational issues. In terms of cloud computing, failure can lead to models of computing that are inefficient, over budget and do not meet business needs.

In order to manage change effectively, a steering committee should support the executive level management. Steering committees are important if change is to happen smoothly and efficiently from the initial stages of adopting a cloud computing model. The main reasons for establishing steering committees are to:

- Ensure that senior management are involved in information security and information technology planning.
- Ensure a good fit between business strategy and technology and security requirements. Steering business leaders and process owners and supporting the organisational 'mind shift'.
- Support business leaders and process owners in the formulation of SLAs and benchmarks e.g. availability, performance, support and response time metrics.
- Co-ordinate the training for direct cloud customers and their end user customers.

13.2 Changing roles and responsibilities

13.2.1 Roles and responsibilities

End users, and particularly internal IT staff, have a critical role in moving to a cloud computing model. System, network and storage administrators as well as application developers will be needed to assess and investigate compatibility and interoperability issues, for example. The outsourcing of computing resources will diminish the need for large internal IT support departments. IT roles in-house will be more concerned with IT risk; vendor due-diligence; and

aligning the business objectives with the provision of the cloud through requirement and standard setting. These types of changes will inevitably be led and managed by Human Resource professionals under the direction of senior management with advice from technical subject matter experts.

The shift away from technical roles will also apply to security roles within the organisation. Security decisions in-house will be more centred on end-point security; provider risk assessments; compliance; accreditation; and, defining SLA's.

13.2.2 Training and re-education

Successful change management programmes have a strong communication and training component since all organisational staff need to be kept informed of changes that will affect their day to day responsibilities. Training material and training workshops are helpful in identifying potential issues and embedding the new mindset. Training material should include, at a minimum:

- An outline of what the proposed cloud services are e.g. Software as a Service (SaaS), Platform as a Service (PaaS) or Infrastructure as a Service (IaaS).
- An outline of how the cloud model will enable the organisation to meet its business and strategic objectives e.g. reduce costs, enhance workforce skills or downsize operational staff for a specific service.
- The operational business benefits of moving to this model of computation e.g. increased agility to supply or volume of service, access to high speed bandwidths and other business enabling technologies.
- The changes affecting staff including training requirements, changes in working patterns, processes and procedures (i.e. IT support) and organisational restructuring.
- An outline of the risks and issues associated with the particular cloud model and how these are being addressed.

One major challenge for organisations' end users in moving to a cloud computing model is embedding new ways of working. Many cloud providers do not offer classes or consulting services in conjunction with the service and there may be a general lack of available training material for the cloud's end users

Customer organisations may need to consider company-wide training to ensure that their workforce can use and realise the benefits of the cloud services. Training also mitigates the risk that IT resources are accidentally misappropriated and helps to avoid inadvertent lapses in security. Training is therefore crucial in delivering security in the initial stages of adoption, particularly for organisations processing commercially sensitive, regulated or government classified data.

13.3 Software development and testing methodologies

Software development and testing methodologies have evolved to require more dynamic, flexible tools and processes. This means development of the information security and technology environment requires more adaptability. Systems specifically need to accommodate shorter project times, be less static, more configurable and support collaborative practices, and deliver this securely (i.e. through deploying the Systems Development Lifecycle³⁰).

Using a cloud service for software development requires the customer to drive the allocation and provisioning of systems on demand. Customers need to be able to utilise the stored physical and virtual machine capacity of a cloud for more flexible automation and reuse across projects, which means less time spent configuring and finding errors when software progresses through the development lifecycle. The benefit of thinking about development resources in this way is to improve the visibility of the organisation's management team over how IT resources are being used. This, in turn, helps to improve efficiency.

Organisations implementing a PaaS cloud software development environment need to consider the following factors:

- Developers should be able to see and make progress on the code and executables in real time without having to replicate data or have to change ownership.
- Configuration management system (version control system) should be integrated with build automation tools in the same way that non-cloud developments take place.
- The branching strategy (parallel development) for development, integration and release work.
- The process for making incremental changes and carrying out system testing regularly during development and in production, and then thereafter during any significant system change.
- Collaboration and change management tools commonly used and understood by software teams. New toolsets should be adaptive to current workflows and tools as developers and engineers remain in their existing working environment.
- Support for heterogeneous environments to enable developers to be flexible across the organisation based on their specific project requirements.

Cloud testing can use web applications that control cloud model environments to simulate real web user traffic as a means of load testing and stress testing web sites. The ability and cost for some customer organisations to simulate web traffic for software testing purposes has been a hindrance to achieving high reliability for their products and services. However, the low cost and accessibility of cloud service models coupled with their inherently large computing

³⁰ According to the OpenSDLC, the Software Development Lifecycle: "provides a consistent peer-reviewed framework for the planning, definition, design, implementation, testing and operational deployment of hardware, software and management systems supporting enterprise class technology products, services, programs and projects."

resources can provide the ability to replicate real world usage of these systems at scales previously unfeasible in traditional testing environments.

There are two factors which organisations should consider before undertaking testing in a cloud software development environment:

- The cloud provider's platform elasticity model, to ensure the environment can be configured to support the testing requirements.
- The cloud provider's security monitoring services and Service Level Agreements (SLAs) to ensure that code is adequately secured and projected.

13.4 Mitigating advice

Mitigation advice	Description
Delivering organisational change	<ul style="list-style-type: none"> • There is a risk that the organisational requirements of the cloud are poorly defined and that adopted approaches do not align with the business strategy. • Change management structures (i.e. steering groups) should be established to identify and address transformational issues and transformation methodologies should be followed.
Training and re-education	<ul style="list-style-type: none"> • Without adequate staff training, there is a risk that roles and responsibilities are poorly defined leading to security breaches and data leakage. • Staff training helps to ensure that roles and responsibilities are understood; embed a new security mindset; and, maintain security standards throughout the transition process. • Staff training also enables cloud customers to realise the benefits of cloud services early in the adoption process.
Software development and testing	<ul style="list-style-type: none"> • Software development and testing in a cloud development environment gives organisations more adaptability and flexibility in the development and testing of software. • Clouds offer the benefit of shared toolsets, greater collaboration and sharing of management information between development partners.

Table 21 - Organisational change mitigating advice

14. Security testing

14.1 The objective: Information and technology risk management

Information and technology risk management is a critical factor for all organisations considering, or currently using, cloud computing services.

Effective information security controls in an increasingly public environment are a key consideration to ensure effective risk management. This enables organisations to adequately protect sensitive data and assets, defend against attackers and continue to operate as normal, without losing brand and customer confidence.

Cloud computing opens the debate on whether organisations can trust a third party provider to store and process private information assets.

- How do I know my cloud computing service provider adheres to the same standards for information security that I do?
- How do I know that my company's data is held securely by my cloud computing service provider, and that it's adequately segregated from my competitor's data who also use the same service provider?

14.2 The approach

An effective, reproducible and quantitative method for enabling information and technology risk management is to perform regular security testing of your security controls, by conducting penetration testing against IT infrastructures and applications.

Methodologies for penetration testing and vulnerability assessments are well known and mature³¹. However, cloud computing increases the risk of some attack vectors which need to be considered when scoping and planning testing.

Whether it's individuals within an organisation using a SaaS solution such as a CRM application, or enterprises using an IaaS solution to run complex or large processing tasks, security testing is an effective method to assess the suitability of controls in order to identify areas of weakness that a malicious user or attacker could exploit to compromise data or initiate a denial of service attack.

Cloud computing services are reliant on several enabling technologies. These technologies present a change to the risk of vulnerabilities being exploited which need to be included in the test approach.

³¹ For example, the CESG CHECK Scheme (http://www.cesg.gov.uk/products_services/iacs/check/index.shtml) is the HMG standard for conducting IT Health Checks within government, comprising methodologies and best practice for the penetration testing of IT infrastructure and applications.

14.3 Testing cloud services

The three types of cloud computing services, SaaS, IaaS and PaaS; all demand different types of security testing:

- Application penetration testing should be conducted on SaaS solutions, but may not be as useful on IaaS and PaaS solutions. Application penetration testing should be used to test the security of a custom application developed for the cloud environment. The testing required will be heavily dependent on the type of application deployed to the cloud, for example whether they use specific platform APIs or web application toolkits. Application testing would therefore need to be customised accordingly and involve asking the following questions to plan and execute a valid application test exercise:
 - Which applications are cloud user facing, and are therefore a potential target for remote attack from insiders or compromised user endpoint devices?
 - Was the application development executed using a recognised Secure Development Lifecycle (SDLC) such as the Open Web Application Security Project (OWASP)³² or Microsoft Security Development Lifecycle (SDL)³³?
 - What application layer security tools are in use within the deployment environment? Are there application layer firewalls, host intrusion detection systems (HIDS) or other execution level systems in place to prevent attacks?
 - What specific application vulnerabilities must be tested? For example this may include: user authentication, database injection, session management, code injection, or privilege escalation.
 - What third party software components do the applications depend on? and what vulnerability testing has been carried out in the development of applications using these components?
- Infrastructure vulnerability assessments, including automated scans which routinely identify missing operating system and application software security patches, should be conducted on IaaS and PaaS solutions, but not necessarily on SaaS solutions where custom code has been developed. This is because the tools are designed to check against mass market vendor operating system and application vulnerabilities affecting multiple organisations, and not focus on custom application designs within a specific customer cloud deployment. Having infrastructure and platform vulnerability testing conducted in a cloud delivery environment is crucial as a missing vendor patch can undermine any prior application testing, rendering a cloud service vulnerable to attack.

When a new application is deployed within the test and production environments of a cloud services provider, a level of trust in the security of these environments is required. Being able to trust the security of cloud environments does not normally mean third parties being

³² <http://www.owasp.org>

³³ <http://msdn.microsoft.com/en-us/security/cc448177.aspx>

permitted to enter and test the security of a cloud provider's infrastructure. The mechanism with which to manage compliance with security and trust requirements is within Service Level Agreements (SLAs) specified by the cloud user organisations. Within these agreements user organisations must specify technical requirements and any relevant accreditations that a cloud provider must adhere to such as PCI DSS or ISO27001. It is important to consider the scope of these certifications, as well as the impact they will have operationally on service management once achieved. These recognised standards require that infrastructure and operating system platforms be regularly patched for security vulnerabilities, and that good practise security implementation and maintenance techniques are used.

14.4 Testing cloud delivery models

The four cloud computing delivery models: vendor, private, hybrid and community clouds will impact the types of and extent of testing permissible.

- Private clouds, managed in-house, will allow for full penetration testing, subject to the owner's risk appetite.
- Vendor (provider) clouds will certainly have more restrictions on the types of testing allowed. In this situation it is important to carry out thorough due diligence.
- Community and hybrid cloud models are even more complicated because of the number of stakeholders involved. These models will require testing to assess the effectiveness of security controls, and the terms of the agreements in place to respond and recover from security incidents.

14.5 The solution

A hybrid approach to testing is needed to adequately test cloud computing security controls and effectively manage risk, involving the use of both technical security testing and due diligence review. Security testing should also take account of the fact that what is tested today may not be representative of tomorrow.

My cloud computing service provider won't let my security testers have physical access to their data centre, so how do we perform adequate security testing?

The following steps should be considered when defining the scope of testing:

Step	Security Testing Approach
A	The use of vulnerability assessment and manual or automated penetration testing methodologies to identify security weaknesses in cloud service provider IT infrastructure, operating systems and application software, where feasible, such as in IaaS implementations.
B	The use of application penetration testing of client software used by cloud computing services, including the penetration testing of Web 2.0 applications, where feasible, such as in SaaS scenarios.
C	Cloud computing end-point vulnerability testing of both vendors' and end users' critical end-point components, such as edge routers, perimeter firewalls and VPN configurations.
D	Integrity testing of data stored by the cloud computing service – including the review of access management, authorisation and authentication mechanisms and cryptographic key management processes and procedures, with a view to ensuring the accuracy and completeness of data, where feasible, such as in SaaS and PaaS scenarios.
E	Analysis of performance and availability of services, including bandwidth, load, capacity and connectivity of key network and server components through the use of monitoring tools.
F	Recovery and resilience testing – testing the ability of a cloud computing service to recover from a crash, connection or software, hardware or network failure and the platform's ability to failover, as one example of a recovery mechanism, to redundant components in the event of a failure. This may involve the review of service level agreements (SLAs) and contracts.
G	Review and audit of the cloud computing provider's adherence to good practice standards and guidelines, such as ISO27001/ BS 25999 and ITIL.

Table 22 - Security testing approach

The checklist shown in Table 23 provides a guideline of the scope to consider when security testing specific cloud computing services.

Test item	SaaS	IaaS	PaaS
Contracts (including NDAs, SLAs)	✓	✓	✓
Accreditation (e.g. ISO27001 Statement of Applicability, SAS70)	✓	✓	✓
Infrastructure penetration testing		✓	✓
Application penetration testing	✓		
Integrity testing (Data auditing)	✓		
Performance testing	✓	✓	✓
BCP testing	✓	✓	✓

Table 23 – Cloud computing testing checklist

15. The future of cloud computing

The purpose of this speculative and predictive section is to enable readers to consider areas for further analysis as opposed to being a definitive assessment of the direction on all aspects of cloud computing technology and provider services. Other sources of direction on the future cloud computing include the Cloud Computing Interoperability Forum (CCIF)³⁴, Cloud Security Alliance (CSA)³⁵, European Network and Information Security Agency (ENISA)³⁶ and the National Institute of Standards and Technology (NIST)³⁷.

The term cloud computing has become synonymous with many different types of Internet services over recent years, but is only now receiving the interest of CIOs and CISOs who are delivering corporate business strategies. Cloud computing has recently moved beyond a conceptual notion and become a realistic option for organisations wishing to reduce their operational infrastructure costs and the complexities of managing their own enterprise IT infrastructures. This status is supported by increased clarity on the term cloud computing driven by industry analysts, and providers striving for interoperability of their cloud services with existing technologies and business processes.

Large organisations such as IBM and Sun Microsystems are beginning to challenge the cloud market and in particular build on the success Google, Amazon and Microsoft have achieved through their innovative strategies. More importantly they are looking at delivering new and enhanced software to help further cloud computing in general practice within enterprises.

“Interest in cloud is reaching near universal proportions for IBM customers, 40% of them use some form of cloud today and almost all are interested in using it in the future” *

* IBM Value Proposition Quantitative Research Phase II Global Report, December 2008

Understanding the present status of cloud computing against historical instances where cloud computing concepts have been implemented, and how they have developed over time, is important when considering the future direction of this still-emerging delivery model.

15.1 Drivers for future change

Cloud computing depends on many technologies that each individually contribute to various cloud architectures. By analysing the business drivers for enhanced cloud computing services, it is possible to see where cloud computing may need to exploit technological developments.

³⁴ <http://www.cloudforum.org/>

³⁵ <http://www.cloudsecurityalliance.org/>

³⁶ <http://www.enisa.europa.eu/>

³⁷ <http://www.nist.gov/>

- **Resilience and reliability** – Cloud providers need to develop the reliability of their services to the point where a 100% availability level for commodity cloud delivery becomes the accepted norm.
- **Security** – Defining security of data and the perimeters in which controls are managed by organisations and cloud providers is an ongoing challenge that will shape cloud computing significantly. Security issues are particularly relevant to the cloud as it can offer a single publicly available entry point to a multitude of different organisations' sensitive data storage and processing systems.
- **Network adaptability** – As the number and type of endpoint devices accessing the cloud grows, there will be a variety of networking technologies used to transfer data to and from the cloud. Constantly changing bandwidth and latency requirements will mean that cloud service providers will need to quickly adapt and provision network bandwidth across customer cloud services utilising shared network infrastructure.
- **Power usage and Green IT** – The objective of reducing power consumption across IT infrastructure usage is fully embedded within manufacturer, provider and user organisations. Cloud Computing Futures is a new initiative in Microsoft Research to improve the efficiency of the scalable computing hardware and software infrastructure needed to deliver cloud services. Data centres and their services have grown in size and importance as Microsoft has shifted to a software-plus-services model in which an increasing number of new applications run in part, or entirely, in the “cloud” and are delivered to clients via the Internet. Achieving green IT targets are on many corporate agendas over the next few years and using cloud services can support these aims.
- **Financial** – A new way of procuring service, platform and infrastructure. In successful projects the reduction of overall expenditure, as well as being able to migrate some capital expenditure to operational expenditure.
- **Remote working** – Along with Green IT, the use of remote working facilities is supported by many organisations wishing to support and maintain an effective workforce. With employees already used to remote access for using secure networks and accessing email, there will be little further push required for them to embrace the cloud model if reliability, usability and functionality are effectively achieved.

15.1.1 General Industry trends

As well as trends specific to the cloud, general IT industry trends will also drive the change in cloud computing services and approach to future services, architectures and innovations.

- **Increasing use of mobile devices** – Laptop sales have overtaken desktops over the last few years and the trend will continue as an increasing range of mobile devices such as netbooks, PDAs and mobile phones incorporate many of the features found on a desktop based PC as little as ten years ago, including Internet access and custom application functionality.

- **Hardware capability improvements** – The inevitable improvements in processor speed and increased memory capacities across IT infrastructure will mean that the cloud will be able to support more complex environments with improved performance capabilities as standard. As an example Intel recently announced it has developed the Single-chip Cloud Computer³⁸ a research chip containing 48 cores.
- **Tackling complexity** – Despite the efforts of multiple technology vendors this challenge remains unresolved. IT architectures continue to be difficult to implement, under-utilised and expensive to operate. The massive scale of cloud computing only strengthens the need for self-monitoring, self-healing and self-configuring IT systems comprising heterogeneous storage, servers, applications, networks and other system elements
- **Legislation and security** – As larger companies consider the cloud computing model, vendors and providers will respond, but within the terms set out by their potential customers. As there still are many issues with respect to data privacy and transfer of data across international borders, the cloud computing providers need to continue to invest time and effort in order to meet the necessary laws required to operate within some of the business areas of their major customers.

15.2 Predictions

The following predictions for cloud computing in general are based on, but not exclusive to, the current drivers and general industry trends discussed above.

5 year predictions - 2015

Organisations will be able to purchase, from the same service provider, a combined cloud environment and managed endpoint service solution that effectively gives a managed IT service for the endpoint devices, along with the central cloud infrastructure. This integrated service offering though potentially available from today's providers in separate parts, will develop into an automated service where changes to the cloud requirements will be linked to those of endpoint devices and vice versa. For example requests to change the number of endpoints would automatically result in changes in scale of the existing cloud service as required to maintain functionality and performance.

The future of cloud computing will go well beyond the realms of simply benefitting businesses in flexibility and cost reduction. IBM recently unveiled the IBM Cloud Academy to develop the future of cloud computing in education. This suggested a shift to the cloud model will increase quality, increase access to educational resources and lower costs. *"Cloud computing makes it easier for those in the education industry, including students, faculty and administrators, to gain immediate access to a wide range of new educational resources and research applications and tools."*³⁹ Michael King, IBM's vice president of global education industry.

³⁸ <http://techresearch.intel.com/articles/Tera-Scale/1826.htm>

³⁹ <http://www-03.ibm.com/press/us/en/pressrelease/28749.wss>

Fully functional desktop equivalent application services, such as office applications, media editing and desktop publishing, will be hosted by providers, which integrate with desktop and laptop PCs and other endpoint devices to provide a feature rich environment available anywhere with connectivity to the cloud. This service would not be limited to remote connectivity to end user devices, but would extend to remote access to applications and operating systems within the cloud. IBM is rolling out a subscription service for hosted virtualised desktops, the IBM Smart Business Desktop Cloud. This targets large and mid-size companies who want virtualised clients but don't have the skills to implement.

10 year predictions – 2020

A global cloud computing standard for each service model will be agreed internationally to drive maturity and interoperability between major providers that also includes specific information for cloud customers on performing due diligence and meeting legislative requirements. The current lack of standards and uncertainty around cloud computing services will be significantly reduced by a global standard. Early moves towards standards are taking place, with organisations like the Cloud Computing Interoperability Forum⁴⁰ (CCIF) looking at interoperability across cloud providers.

Organisations using cloud computing will easily be able to transfer existing cloud computing services contracts between providers to gain more competitive services. This will be facilitated by fully automated and regulated agreements between providers and backed by interoperable technology standards.

⁴⁰ cloudforum.org

16. Glossary

AAF	Australian Access Federation
API	Application Programming Interface
Attacker	A person or group that seeks to impact the confidentiality, integrity and availability of data stored within a cloud.
Bandwidth	The amount of electronic data that can be transmitted in a fixed period of time across a network link.
BCM	Business Continuity Management
BCP	Business Continuity Planning
CCIF	Cloud Computing Interoperability Forum
Cloud	A set of computing services that are offered across a network link by a provider to a customer.
CPU	Central Processing Unit
CRM	Customer Relationship Management
CSA	Cloud Security Alliance
Customer	An organisation that procures cloud computing services from a provider.
ENISA	European Network and Information Security Agency
EU DPD	European Union Data Protection Directive
FSA	Financial Services Authority
HIPAA	Health Insurance Portability and Accountability Act (US)
HMG	Her Majesty's Government
Hypervisor	Software and/or firmware running on the host that creates and controls all the virtual machines. It provides the virtual machine environment.
ICT	Information Communications Technology
ISO27001	ISO/IEC 27001:2005 – Information technology – Security techniques – Information security management systems – Requirements
IT DR	Information Technology Disaster Recovery
ITIL	Information Technology Infrastructure Library
LAN	Local Area Network
Multi-tenant	A cloud computing infrastructure component that a provider has shared across multiple customers.
NDA	Non Disclosure Agreement
NIST	National Institute of Standards and Technology
OS	Operating System
PCI DSS	Payment Card Industry Data Security Standard
PDA	Personal Digital Assistant
Provider	An organisation that provides cloud computing services to customers.
QoS	Quality of Service
RIPA	Regulatory Investigatory Powers Act
SAML	Security Assertion Markup Language

SAP	Enterprise Resource Planning products vendor
SAS70	Statement on Auditing Standard 70
SOA	Service Oriented Architecture
SLA	Service Level Agreement
SPOF	Single Point of Failure
Supplier	An organisation that provides products or services to the cloud customer or provider organisations.
Third party	An organisation that is a separate entity to another organisation. For example a cloud provider may be a third party to cloud customer, and a supplier to a cloud customer or provider may be a third party of those organisations.
Vendor	An organisation that may provide cloud computing services to customers, or may also be a software or hardware manufacturer of cloud infrastructure components.
Virtual Machine	A software implementation of a machine (computer) that executes programs like a real machine.
Virtualisation	The practice of hosting multiple logically separate operating systems on a single hardware platform, with each operating system taking a share of the available underlying resources.
VOIP	Voice over IP - a system used to provide voice communications, equivalent to traditional telephony, using the Internet or another computer network to transmit speech.
VPN	Virtual Private Network
WAN	Wide Area Network