INTERNAL AFFAIRS
Te Tari Taiwhenua

New Zealand Government

Cloud Computing

Information Security and Privacy Considerations

April 2014

Contents

Table of figures

Table of tables

1 Introduction

In October 2013, Cabinet agreed on a cloud computing risk and assurance framework for government agencies, to sit within the wider ICT Assurance Framework. The agreed approach is based on the following principles:

- case-by-case consideration, by agency chief executives with Government Chief Information Officer (GCIO) oversight, of all cloud computing decisions, whether hosted onshore or offshore, that balances the risk and benefits appropriately;
- agency Chief Executives are ultimately responsible for decisions to use cloud services
- no data above RESTRICTED should be held in a public cloud, whether it is hosted onshore or offshore;
- all agencies in the State services are expected to follow a uniform and robust information management process that includes:
    - classifying the information
    - undertake a risk assessment using the agency's own processes, if they have them, or those supplied by the GCIO in the *Risk Assessment Processes: Information Security*
    - if the system is likely to be a cloud service, Public and non-Public Service departments must use the guidelines in this paper to ensure appropriate and consistent consideration of cloud computing issues (including privacy and security);
- the GCIO has oversight of all-of-government and agency cloud solutions to provide assurance that the guidance and risk assessment process has been correctly followed by agencies; and
- when necessary, the GCIO may direct Public and non-Public Service departments to modify their use of cloud services.

This document presents information security and privacy implications that need to be carefully considered and managed by agencies seeking to take advantage of cloud computing. The process is mandatory for Public and non-Public Service departments as part of the robust information management process listed above, however, all State services agencies are expected to follow the process.

The process does not attempt to qualify or quantify the risks associated with the adoption of cloud services - rather it is designed to support agencies when they are performing a risk assessment. This document enables agencies to systematically identify, analyse and evaluate the information security and privacy risks associated with cloud services, and provides controls to effectively manage those risks.

Although it presents the most common areas of concern associated with cloud computing, the risks identified in this document should not be considered exhaustive and agencies are encouraged to identify and assess any other risks that may be unique to their business context or the cloud services that they are planning to use.

2 Overview of Cloud Computing

There are many different definitions for cloud computing. The New Zealand government has adopted the National Institute of Science and Technology (NIST) definition that defines cloud computing as:

"A model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g. networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction."[1]

This section provides a brief overview of the essential characteristics of cloud computing together with the cloud service and deployment models. It is recommended that agencies familiarise themselves with the NIST definitions to ensure that they are able to identify and understand the risks associated with different cloud service and deployment models.

2.1 Essential Characteristics

The following provides an overview of the five essential characteristics for cloud computing as defined by NIST:

- **On-Demand Self-Service** – customers are able to provision resources (e.g. a virtual server or email account) without any interaction with the service provider's[2] staff.
- **Broad Network Access** – customers are able to access resources over networks such as the Internet using a ubiquitous client (e.g. a web browser) from a range of client devices (e.g. smartphones, tablets, laptops).
- **Resource Pooling** – the service provider's computing resources are pooled to serve multiple customers. Typically, virtualisation technologies are used to facilitate multi-tenancy and enable computing resources to be dynamically assigned and reallocated based on customer demand.
- **Rapid Elasticity** – resources can be quickly provisioned and released, sometimes automatically, based on demand. Customers can easily increase or decrease their use of a cloud service to meet their current needs.
- **Measured Service** – customers pay only for the resources they actually use within the service. Typically the service provider will supply customers with a dashboard so that they can track their usage.

2.2 Service Models

The following provides an overview of the three cloud service models defined by NIST together with some real world examples for each:

- **Infrastructure as a Service (IaaS)** – the provision of computing resources (i.e. processing, memory, storage and network) to allow the customer to deploy and run their own operating

[1] The NIST Definition of Cloud Computing: http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf

[2] Commonly referred to as a Cloud Service Provider or CSP

systems and applications. Typically, virtualisation technologies are used to enable multiple customers to share the computing resources. The service provider is only responsible for managing and maintaining the underlying infrastructure hardware and virtualisation hypervisor[3]. Examples of IaaS offerings include the government IaaS platforms, Amazon Web Services (AWS), Elastic Cloud Compute (EC2), Google Compute Engine and Rackspace Compute.

- **Platform as a Service (PaaS)** – the provision of standardised operating systems and application services (e.g. web server or database platform) delivered on IaaS services to enable customers to deploy and run their own applications developed using programming languages supported by the service provider. The service provider is responsible for managing and maintaining the underlying infrastructure hardware, virtualisation hypervisor, operating systems and standard application services. Usually, customers can only make predefined configuration changes to the standard operating systems and application services but remain responsible for managing and maintaining their applications. Examples of PaaS offerings include the government Desktop as a Service (DaaS), Google App Engine, Microsoft Windows Azure, Force.com and Oracle Database Cloud.

- **Software as a Service (SaaS)** – the provision and consumption of the service provider's standardised application services (e.g. email or customer relationship management) usually on a pay-per-use basis using a web browser or thin client application[4]. The service provider is solely responsible for managing and maintaining the application, platforms and underlying infrastructure. Customers can typically only make predefined configuration changes to the application and manage user permissions to their own data. Examples of SaaS offerings include the government Office Productivity as a Service (OPaaS), Microsoft Office 365, Google Apps, Salesforce.com and Oracle Applications Cloud.

2.3 Deployment Models

The following provides an overview of the four cloud delivery models defined by NIST:

- **Public Cloud** – the provision and use of services that are hosted, operated and managed by a service provider. Public cloud services are typically delivered over the Internet from one or more of the service provider's data centres. They are offered to the general public and rely on multi-tenancy (i.e. multiple customers sharing the service providers resources) to drive economies of scale and deliver the maximum potential cost efficiencies. However, they usually offer a low degree of control and oversight of the security provided by the service.

- **Private Cloud** – the provision of services exclusively for the use of a single organisation (i.e. there is no multi-tenancy). A number of private cloud patterns have emerged and the following provides an overview of the most common patterns:

    - **Dedicated** – the service is owned, operated and managed by the organisation and is hosted within its premises or co-located within a data centre facility;

[3] A hypervisor is a specialised operating system that enables server hardware to run multiple guest operating systems concurrently

[4] A light-weight application that performs minimal processing which relies on a server component to perform information processing activities

- **Managed** – the service is owned by the organisation but is operated and managed on its behalf by a service provider. The service may be hosted within the organisation's premises or co-located within the service provider's facility;

- **Virtual** – the service is owned, operated, managed and hosted by a service provider but the organisation is logically isolated from other customers.

When compared to the other deployment models, private clouds (usually with the exception of virtual private clouds) provide a greater degree of control and oversight of the security provided by the service. However, they also provide the lowest cost efficiencies because the organisation must invest capital to purchase the hardware and software required to meet its anticipated peak usage. Further, costs to maintain hardware over time as it is superseded or falls out of warranty may also be borne directly by the Customer.

Note: A virtualised compute environment is not considered a private cloud if it does not exhibit the five essential characteristics (see Essential Characteristics) for cloud computing.

- **Community Cloud** – a community cloud is essentially a private cloud that is shared by a number of organisations that have similar business objectives and/or requirements such as different government agencies within a specific sector. They attempt to achieve a similar level of security control and oversight as those provided by private clouds whilst trying to offer some of the cost efficiencies offered by public clouds.

- **Hybrid Cloud** – a hybrid cloud is created when an organisation uses a combination of two or more of the other cloud deployment models to implement its cloud strategy. For example, an organisation might choose to publish its websites from the public cloud at the same time as it continues to deliver its business critical applications from an in-house private cloud.

2.3.1 Responsibility for Security in Cloud Computing Environments

Figure 1 highlights the party that is responsible for implementing and managing information security controls across the different cloud service models.

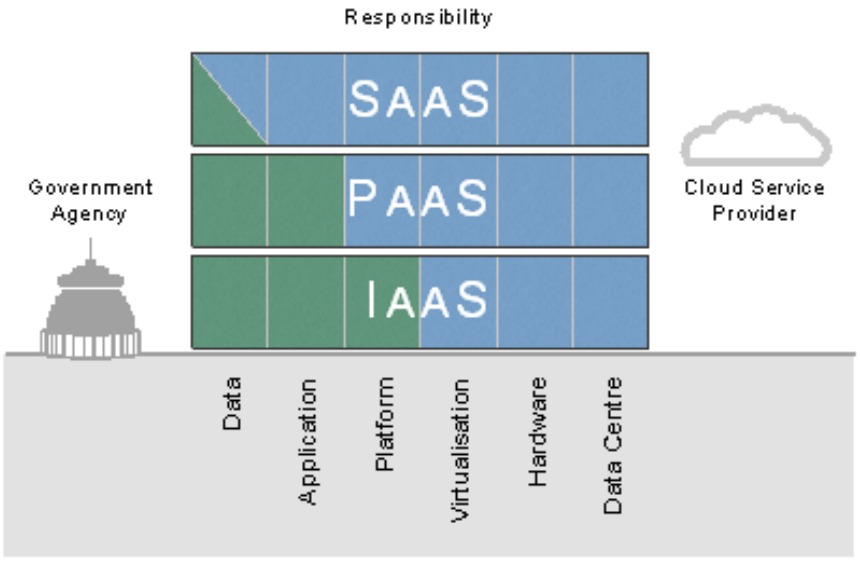

Figure 1 - Responsibility for Information Security Controls by Cloud Service Model

The following provides an overview of the responsibility boundary for each of the service models:

- **IaaS** – the service provider is responsible for the implementation, management and maintenance of the information security controls up to, and including, the virtualisation hypervisor layer (i.e. the underlying infrastructure). Customers are responsible for ensuring that there are appropriate security controls in place to protect and maintain all of the components built on top of the hypervisor including the guest operating system, application services and the applications they deploy within the IaaS environment.

- **PaaS** – The PaaS service model builds upon IaaS to include the guest operating system and application services. Therefore the service provider is also responsible for implementing, managing and maintaining the security controls to protect these components. Customers are responsible for ensuring that the applications that they deploy on the PaaS environment are secure.

- **SaaS** – the customer has very limited control over security in the SaaS service model. Generally they will maintain responsibility for managing their user accounts to ensure that they are only assigned the permissions required to perform their duties. The service provider is responsible for ensuring all other security controls are in place and provide an appropriate level of protection.

Note: it is important to understand that although agencies can outsource responsibility to a service provider for implementing, managing and maintaining security controls they cannot outsource their accountability for ensuring their data is appropriately protected.

3 Security and Privacy Considerations

This section describes the core considerations for any agency planning a deployment of a cloud computing service. Each area is described in some detail followed by a list of key considerations to assist agencies in developing an assessment of their risk position for a proposed service.

3.1 Value, Criticality and Sensitivity of Information

In order to be able to assess the risks associated with using a cloud service, agencies must recognise the value, criticality and sensitivity of the information they intend to place in the service.

Agencies are required to classify official information in accordance with the guidance published in '*Security in the Government Sector 2002 (SIGS)'* [5]. They are also required to protect official information in line with the guidance published in the '*New Zealand Information Security Manual (NZISM)'*[6].

The under-classification of data could result in official information being placed in a cloud service that does not have appropriate security controls in place and therefore cannot provide an adequate level of protection. Conversely, over-classification could lead to unnecessary controls being specified leading to excessive costs resulting in suitable cloud services being rejected. Therefore it is critical that an agency accurately assesses the value, criticality and sensitivity of its data, and correctly classifies it to ensure that it is appropriately protected.

Key Considerations

1. Who is the business owner of the information?
2. What are the business processes that are supported by the information?
3. What is the security classification of the information based on the NZ government guidelines for protection of official information?
4. Are there any specific concerns related to the confidentiality of the information that will be stored or processed by the cloud service?
5. Does the data include any personal information?
6. Who are the users of the information?
7. What permissions do the users require to the information? (i.e. read, write, modify and/or delete)
8. What legislation applies to the information? (e.g. Privacy Act 1993, Official Information Act 1982, Public Records Act 2005)
9. What contractual obligations apply to the information? (e.g. Payment Card Industry Data Security Standard (PCI DSS))

[5] Available from http://www.security.govt.nz/assets/media/Security_in_the_Government_Sector_2002.pdf

[6] Available from http://www.gcsb.govt.nz/newsroom/nzism/NZISM_2011_Version_1.01.pdf

10. What would the impact on the business be if the information was disclosed in an unauthorised manner?

11. What would the impact on the business be if the integrity of the information was compromised?

12. Does the agency have incident response and management plans in place to minimise the impact of an unauthorised disclosure?

13. What would the impact on the business be if the information were unavailable?

    a. What is the maximum amount of data loss that can be tolerated after a disruption has occurred? This is used to define the Recovery Point Objective.

    b. What is the maximum period of time before which the minimum levels of services must be restored after a disruption has occurred? This is used to define the Recovery Time Objective.

    c. What is the maximum period of time before which the full service must be restored to avoid permanently compromising the business objectives? This is used to define the Acceptable Interruption Window.

3.2 Data Sovereignty

The use of cloud services located outside of New Zealand's jurisdiction, or owned by foreign companies, introduces data sovereignty risks. This means that any data stored, processed or transmitted by the service may be subject to legislation and regulation in those countries through which data is stored, processed and transmitted. Similarly, a foreign owned service provider operating a service within New Zealand may be subject to the laws of the country where its registered head offices are located.

The laws that could be used to access information held by the service provider vary from country to country. In some instances when a service provider is compelled by a foreign law enforcement agency to provide data belonging to their customers, they may be legally prohibited from notifying the customer of the request. Therefore it is critical that an agency identify the legal jurisdictions in which its data will be stored, processed or transmitted. Further, they should also understand how the laws of those countries could impact the confidentiality, integrity, availability and privacy of the information.

If the service provider outsources or sub-contracts any aspect of the delivery of the service to a third-party, agencies must also identify whether this introduces additional data sovereignty risks.

Privacy information that is held in legal jurisdictions outside of New Zealand may be subject to the privacy and data protection laws of the countries where the cloud service is delivered. Privacy and data protection laws can vary considerably from country to country. Therefore it is important that agencies assess how the laws of those countries could affect the privacy of their employees and/or customers' information.

Key Considerations

14. Where is the registered head office of the service provider?

15. Which countries are the cloud services delivered from?

16. In which legal jurisdictions will the agency's data be stored and processed?

17. Does the service provider allow its customers to specify the locations where their data can and cannot be stored and processed?

18. Does the service have any dependency on any third parties (e.g. outsourcers, sub-contractors or another service provider) that introduce additional jurisdictional risks? If yes, ask the service provider to provide the following details for each third party involved in the delivery of the service:

    a. The registered head office of the third party;

    b. The country or countries that their services are delivered from; and

    c. The access that they have to client data stored, processed and transmitted by the cloud service.

19. Have the laws of the country or countries where the data will be stored and processed been reviewed to assess how they could affect the security and/or privacy of the information?

20. Do the laws actually apply to the service provider and/or its customer's information? (e.g. some privacy laws exempt certain types of businesses or do not apply to the personal information of foreigners.)

21. Do the applicable privacy laws provide an equivalent, or stronger, level of protection than the Privacy Act 1993? If no, are customers able to negotiate with the service provider to ensure that the equivalent privacy protections are specified in the contract?

22. How does the service provider deal with requests from government agencies to access customer information?

    a. Do they only disclose information in response to a valid court order?

    b. Do they inform their customers if they have to disclose information in response to such a request?

    c. Are they prevented from informing customers that they have received a court order requesting access to their information?

Once agencies have identified the legal jurisdictions where their data will be held, they should assess whether or not it is appropriate to store their data in the service. This may require them to

seek specialist legal and/or security advice. Agencies without access to specialist resources are encouraged to seek advice from the Government Chief Information Officer (GCIO).

3.3 Privacy

Agencies planning to place personal information[7] in a cloud service should perform a Privacy Impact Assessment (PIA)[8] to ensure that they identify any privacy risks associated with the use of the service together with the controls required to effectively manage them.

Cloud services may make it easier for agencies to take advantage of opportunities to share information. For example, sharing personal information with another agency may be achieved by simply creating user accounts with the appropriate permissions within a SaaS solution rather than having to implement a system-to-system interface to exchange information. Although cloud services have the potential to lower the technical barriers to information sharing agencies must ensure that they appropriately manage access to personal information and comply with the requirements of the Privacy Act 1993.

Service providers typically use privacy policies to define how they will collect and use personal information about the users of a service. US service provider's privacy policies usually distinguish between Personally Identifiable Information (PII) and non-personal information. However, it is important to note that both are considered personal information under the Privacy Act 1993. Agencies must carefully review and consider the implications accepting a service provider's privacy policy.

Key Considerations

23. Does the data that will be stored and processed by the cloud service include personal information as defined in the Privacy Act 1993[9]?

If no, skip to question 28.

24. Has a PIA been completed that identifies the privacy risks associated with the use of the cloud service together with the controls required to effectively manage them?

25. Is the service provider's use of personal information clearly set out in its privacy policy? Is the policy consistent with the agency's business requirements?

26. Does the service provider notify its customers if their data is accessed by, or disclosed to, an unauthorised party? Does this include providing sufficient information to support cooperation with an investigation by the Privacy Commissioner?

27. Who can the agency, its staff and/or customers complain to if there is a privacy breach?

[7] Personal information is information about an identifiable, living individual. This may include information even when additional steps or knowledge are required to identify an individual from the information.

[8] The Privacy Impact Assessment Handbook is available from http://privacy.org.nz/assets/Uploads/Privacy-Impact-Assessment-Handbook-June2007.pdf

[9] See http://www.legislation.govt.nz/act/public/1993/0028/latest/DLM296639.html

In addition to this, the Office of the Privacy Commissioner (OPC) has published guidance for small to medium organisations that are considering placing personal information in a cloud service[10]. Agencies are encouraged to review and ensure that they understand the guidance.

3.4 Governance

3.4.1 Terms of Service

Cloud computing is essentially a form of outsourcing and like all outsourcing arrangements, it introduces governance challenges. However, unlike traditional outsourcing models it may not always be possible for customers to fully negotiate all contract terms with the service provider, especially in the case of public cloud services (e.g. Google Apps, Microsoft Office 365, Amazon Web Services).

The primary governance control available to agencies is the service provider's Terms of Service (or contract), the associated Service Level Agreement (SLA) and Key Performance Indicators and metrics demonstrating the service performance. These must be carefully reviewed to ensure that the service can meet the agency's obligations to protect the confidentiality, integrity and availability of its official information and the privacy of all personally identifiable information it intends to place within it.

To be able to exercise any level of control over the data that is held in the cloud service agencies must maintain ownership of their data and know how the service provider will use the data when delivering the service. Service providers may use customers' data for their own business purposes (e.g. for generating revenue by presenting targeted advertising to users or collecting and selling statistical data to other organisations). Although the use of customer data is usually limited to consumer rather than enterprise contracts it is important to determine whether the service provider will use the data for any purpose other than the delivery of the service. Therefore, the service provider's Terms of Service must be reviewed to ensure that they clearly define the ownership of data, how it will be used in the delivery of the service and whether the service provider will use it for any purpose other than the delivery of the service.

It is not uncommon for a service provider to rely on components from other service providers. For example, a SaaS service may be hosted on an IaaS offering from a different provider. It is essential to identify any dependencies that the service provider has on third-party services to gain a complete understanding of the risks introduced through the adoption of a service.

Key Considerations

28. Does the service provider negotiate contracts with their customers or must they accept a standard Terms of Service?

29. Does the service provider's Terms of Service and SLA clearly define how the service protects the confidentiality, integrity and availability of official information and the privacy of all personally identifiable information?

30. Does the service provider's Terms of Service specify that the agency will retain

[10] See http://privacy.org.nz/making-the-right-choices-in-cloud-computing-new-privacy-commissioner-guidance/

ownership of its data?

31. Will the service provider use the data for any purpose other than the delivery of the service?

32. Is the service provider's service dependent on any third-party services?

3.4.2 Compliance

The NZISM advises agencies to formally assess and certify that their information systems have been deployed with sufficient controls to protect the confidentiality, integrity and availability of the information they store, process and transmit before accrediting them for use.

As discussed, it may not be possible for customers to negotiate the terms of the contract with a service provider. As a result, an agency may not be able to stipulate any specific security controls to protect its data, or to directly verify that the service provider has sufficient controls in place to protect its data. Even if is possible to directly verify that a service provider has controls, it may not actually be practical to do so if the service is hosted in a data centre outside New Zealand. Therefore customers must typically rely on the service provider commissioning a third-party audit.

Agencies need to consider which certifications are useful and relevant, and whether or not they increase their confidence in the service provider's ability to protect their information. It is essential that an agency understand if certification to an internationally recognised standard or framework provides any assurance that the service provider meets its security requirements. For example, the *Statement for Standards for Attestation Engagements (SSAE) No. 16 Service Organization Control (SOC) 2 Type II* allows the service provider to limit the scope of the audit. Similarly, service providers that are certified as being compliant with the requirements defined in ISO/IEC 27001 are able to define the scope of the audit using a Statement of Applicability. Therefore agencies need to check exactly what controls are covered by the audit by asking the service provider for a copy of the latest external auditor's report (including the scope or Statement of Applicability), and the results of all recent internal audits.

Access to information related to audits varies amongst service providers. Some are willing to provide customers (including potential customers) with full audit reports under a non-disclosure and confidentiality agreement. Whereas others will only provide the certificate to demonstrate that they have passed the audit. The more transparent the service provider is, the easier it is for agencies to assess if the provider has suitable security practices and controls in place to meet their requirements.

Another potential source of information relating to the security controls that a service provider has in place is the *Cloud Security Alliance's Security, Trust & Assurance Register* (*CSA STAR*). The level of assurance provided depends on the level that the service provider has achieved on the *CSA's Open Certification Framework* (OCF).

The first level is self-assessment. To achieve this, service providers submit a completed *Consensus Assessments Initiative Questionnaire* (*CAIQ*) or *Cloud Controls Matrix* (*CMM*) report that asserts their compliance with the CSA cloud security controls. While these reports can provide agencies with an insight into the service provider's security controls and practices, the CSA only verifies authenticity of the submission and performs a basic check of the accuracy of its content

before adding it to the registry. The CSA does not guarantee the accuracy of any entries. As a result, the fact that a provider is listed on the *CSA STAR Self-Assessment* is an indication that the provider has sought to assert some level of diligence with a registration body but does not actually provide any assurance that they have adequate security practices or controls in place.

The second levels are *CSA STAR Certification* and *Attestation*. To achieve these levels service providers undergo third party auditing by an approved Certification Body. The *CSA STAR Certification* is based on ISO/IEC 27001 and the controls specified in the CMM. The maturity of the service provider's Information Security Management System (ISMS) is assessed and given a rating (i.e. Bronze, Silver or Gold) if they are found to have adequate processes in place. Similarly, the *CSA STAR Attestation* is based on SSAE 16 SOC 2 Type II and is supplemented by the criteria defined in the CMM. The service providers are regularly assessed based on the controls that they assert are in place and their description of the service.

The third level is continuous monitoring and assessment of the cloud service's security properties using the CMM and the CSA's *Cloud Trust Protocol* (CTP). This is currently in development and is not anticipated to be available until 2015. The goal of *CSA STAR Continuous* is provide on-going assurance about the effectiveness of the service provider's security management practices and controls.

The Institute of Information Technology Practitioners (IITP) has published the *New Zealand Cloud Computing Code of Practice*[11] that provides a standardised method for New Zealand based service providers to voluntarily disclose information about the security of their service(s). The *Cloud Code* is designed to ensure that service providers are transparent in the positioning of their services to their clients. However, it does not provide any specific assurance that they have adequate security practices or controls in place. Therefore, an agency should only use the *Cloud Code* for informational purposes and should not rely on it to replace its own due diligence.

When relying on certification performed by another party (e.g. a third-party auditor or another government agency) it is important for agencies to understand the scope and limitations of the certification and to assess whether they need to perform further assurance activities. For example, agencies deploying services on one of the approved government IaaS platforms must perform a certification and accreditation review of the components they implement as part of their project (e.g. guest operating systems and applications).

Key Considerations

33. Does the service provider's Terms of Service allow the agency to directly audit the implementation and management of the security measures that are in place to protect the service and the data held within it?

    a. If yes, does this include performing vulnerability scans and penetration testing of the service and the supporting infrastructure?

    b. If no, does the service provider undergo formal regular assessment against an internationally recognised information security standard or framework by an

[11] See http://www.nzcloudcode.org.nz/ for details

independent third-party? (E.g. are they certified as being compliant with ISO/IEC 27001? Have they undergone an ISAE 3402 SOC 2 Type II?)

34. Will the service provider allow the agency to thoroughly review recent audit reports before signing up for service? (E.g. will the service provider provide the Statement of Applicability together with a copy of the full audit reports from their external auditor, and the results of any recent internal audits?)

35. Will the service provider enable potential customers to perform reference checks by providing the contact details of two or more of its current customers?

36. Is there a completed CAIQ or CMM report for the service provider in the CSA STAR?

37. Has the service provider undergone a CSA STAR Certification and/or Attestation? Have they published the outcome of the audit?

38. Has the service provider published a completed Cloud Computing Code of Practice?

39. What additional assurance activities must be performed to complete the certification and accreditation of the cloud service?

3.5 Confidentiality

There are many factors that may lead to unauthorised access to, or the disclosure of, information stored in a cloud service. However, it is important to note that the vast majority of these are not unique to cloud computing.

As highlighted in Figure 1 the cloud service model (i.e. IaaS, PaaS or SaaS) will determine which party is responsible for implementing and managing the controls to protect the confidentiality of the information stored, processed or transmitted by the service. Similarly, the cloud deployment model (i.e. public, private, community or hybrid) will affect a customer's ability to dictate its control requirements.

3.5.1 Authentication and Access Control

An agency may find that as its use of cloud services increases so will the identity management overhead. The adoption of multiple cloud services may place an unacceptable burden on users if the agency does not have an appropriate identity management strategy. For example, each cloud service that is adopted could result in users requiring another username and password). A discussion of the approaches to identity management is beyond the scope of this document. However, agencies are encouraged to develop an approach to identity and access management that supports their adoption of cloud services, by both their employees and customers. This should include consideration of the security implications and risks.

The broad network access characteristic of cloud computing amplifies the need for agencies to have strong identity lifecycle management practices. This is because users can typically access the information held in a cloud service from any location, which could present a significant risk as employees or contractors may still be able access the service after they have ceased to be employed. Therefore agencies should maintain a robust process for managing the lifecycle of identities that ensures:

- Permissions are approved at the appropriate level within the organisation.
- Role Based Access Control (RBAC) is sufficiently granular to control permissions.
- Users are only granted the permissions they require to perform their duties.
- Users do not accumulate permissions when they change roles within the organisation.
- User accounts are removed in a timely manner when employment is terminated.

In addition, agencies should regularly audit user accounts and the permissions granted to the accounts within the cloud services they have adopted to ensure that redundant accounts are removed and that users continue to only be granted the permissions they require to perform their duties.

Ubiquitous access also means that users can access the information held in the cloud service from any location using many different devices. Agencies must carefully consider the associated information security implications and assess what controls are required to adequately protect their information. For example, an agency adopting a SaaS based Customer Relationship Management (CRM) solution may determine that it needs to restrict access to specific features and functionality (e.g. downloading customer records or saving reports) when users access the service from outside the agency's premises or using a non-agency owned and managed device.

Another area of concern when adopting cloud services is whether passwords provide a sufficient level of assurance that the person authenticating to the service is the owner of the user account. Agencies must determine whether they require a stronger authentication mechanism (e.g. multi-factor authentication) that provides sufficient confidence that the party asserting the identity is the authorised user.

Key Considerations

40. Does the agency have an identity management strategy that supports the adoption of cloud services? If yes, does the cloud service support the agency's identity management strategy?

41. Is there an effective internal process that ensures that identities are managed throughout their lifecycle?

42. Is there an effective audit process that is actioned at regular intervals to ensure that user accounts are appropriately managed?

43. Have the controls required to manage the risks associated with the ubiquitous access provided by the cloud been identified?

44. Does the cloud service meet those control requirements?

45. Is there a higher level of assurance required that the party asserting an identity is the authorised user of the account when authenticating to the service? (I.e. is multi-factor authentication necessary?)

3.5.2 Multi-Tenancy

The resource pooling characteristic of cloud computing means that cloud services typically use some form of multi-tenancy. This enables service providers to deliver services at a lower cost than traditional delivery models by allowing multiple customers (tenants) to share the same compute resources and/or instance of an application. While resource pooling and sharing has obvious benefits in terms of costs it does introduce security risks that must be understood by agencies wishing to leverage the benefits of cloud computing. The risks associated with multi-tenancy are typically related to either infrastructure virtualisation or data commingling.

Virtualisation is a key technology in the delivery of cloud services as it enables information systems to be abstracted from the underlying hardware using a hypervisor (i.e. software that enables a host server to run multiple guest operating systems concurrently). The most often cited area of concern within a virtualised environment is that a malicious party could exploit a vulnerability within the hypervisor to gain access to another customers' information (e.g. by performing a 'guest-to-host' or 'guest-to-guest' attack).

Virtualisation has made it easy to take a snapshot (i.e. a copy of a running server's memory and disk at a point in time for backup and redundancy purposes). If the snapshots are not appropriately protected, a malicious party may be able to gain unauthorised access to the information stored on the virtual machine's local drives and all encryption keys and data stored in memory. As a result, the service provider's architecture, implementation and ongoing management and monitoring of the virtualisation environment together with their patch and vulnerability management practices are key to ensuring the security of information stored and processed within the cloud service.

Another common concern in IaaS and PaaS environments is that the customer with the weakest security practices and controls may determine the security of the entire environment (the lowest common denominator problem). For example, a co-tenant that does not harden its operating systems and applications could define the security of the environment to the lowest common denominator if there are not appropriate controls in place to isolate customer's virtual machines and networks from each other.

SaaS and PaaS services use logical controls within the application or platform and supporting infrastructure to isolate access to each customer's data. However, the data is usually commingled within the application, database and back-up media. This places complete reliance on the quality of the design, implementation and enforcement of access controls within the platforms and applications.

The on-demand self-service characteristic of cloud computing introduce security concerns because the registration processes to become a customer are not always robust in confirming a customer's identity (i.e. web-based self-registration). This weakness can allow a malicious party to register for a service to then use it for malicious or fraudulent activities that may include attempting to subvert the access controls to gain unauthorised access to another customer's data.

An agency must be sufficiently assured that other customers using a cloud service cannot subvert the service provider's controls to gain access to its data. As discussed, this can be difficult as the "*as a service*" nature of cloud computing often means a lack of transparency regarding the security controls and practices that the service provider has in place to protect their customers' data. Consequently there is again a strong dependency on third-party audit reports and penetration testing.

Key Considerations

46. Will the service provider allow the agency to review a recent third-party audit report (e.g. ISO 27001 or ISAE 3402 SOC 2 Type II) that includes an assessment of the security controls and practices related to virtualisation and separation of customer's data?

47. Will the service provider permit customers to undertake security testing (including penetration tests) to assess the efficacy of the access controls used to enforce separation of customer's data?

48. Does the service provider's customer registration processes provide an appropriate level of assurance in line with the value, criticality and sensitivity of the information to be placed in the cloud service?

3.5.3 Standard Operating Environments

Although not unique to cloud computing it is important to acknowledge that one of the biggest causes of information security incidents is poorly configured and managed information systems. While the service provider is entirely responsible for ensuring that their SaaS solution is appropriately configured and managed, the responsibility is shared between the agency and the service provider in the other cloud service models (i.e. IaaS and PaaS). Agencies that do not have defined and documented build and hardening standards for operating systems and applications they are planning to deploy on IaaS or PaaS cloud services may find it difficult to effectively protect their systems against unauthorised access.

Where an agency decides to delegate the build and hardening of the operating systems and applications to the service provider, it must determine whether it is appropriate to accept the provider standards or define its own. Irrespective of the approach that is selected by the agency it is recommended that a penetration test be undertaken to ensure that services are initially deployed in a secure manner.

Key Considerations

49. Are there appropriate build and hardening standards defined and documented for the service components the agency is responsible for managing?

50. Can the agency deploy operating systems and applications in accordance with internal build or hardening standards? If no, does the service provider have appropriate build and hardening standards that meet the agency's security requirements?

    a. Does the virtual image include a host-based firewall configured to only allow the ingress and egress (inbound and outbound) traffic necessary to support the service?

    b. Does the service provider allow host-based intrusion detection and prevention service (IDS/IDP) agents to be installed within the virtual machines?

51. Does the service provider perform regular tests of its security processes and controls?

Will they provide customers with a copy of the associated reports?

52. Can a penetration test of the service be performed to ensure that it has been securely deployed?

3.5.4 Patch and Vulnerability Management

Improved patch and vulnerability management is often cited as one of the main benefits of moving to the cloud. Vulnerabilities present a significant risk to any information system, particularly those that are exposed to the Internet. The ubiquitous access provided by cloud services mean that it is very important that agencies ensure that these services are patched in a timely manner.

It is important to identify which party is responsible for patching each component of a cloud service (e.g. the application, operating system, hypervisor software, Application Programming Interface (API) etc.). As discussed, the cloud service model (i.e. SaaS, PaaS or IaaS) usually dictates which party is responsible for the management and maintenance of individual components.

Where the service provider is responsible the agency must ensure that Terms of Service and SLA specify the maximum time period permitted between a patch being released by a vendor and being applied to all affected systems (i.e. the maximum exposure window).

Where the agency is responsible for applying patches it must ensure that it has an effective patch management process and monitors appropriate sources for vulnerability alerts (e.g. the vendor's website and/or mailing list, Common Vulnerabilities and Exposures (CVE) databases and the National Cyber Security Centre (NCSC) website) to ensure patches are identified and deployed in a timely manner.

Key Considerations

53. Is the service provider responsible for patching all components that make up the cloud service? If no, has the agency identified which components the service provider is responsible for and which it is responsible for?

54. Does the service provider's Terms of Service or SLA include service levels for patch and vulnerability management that includes a defined the maximum exposure window?

55. Does the agency currently have an effective patch and vulnerability management process?

56. Has the agency ensured that all of the components that it is responsible for have been incorporated into its patch and vulnerability management process?

57. Is the agency subscribed to, or monitoring, appropriate sources for vulnerability and patch alerts for the components that it is are responsible for?

58. Does the service provider allow its customers to perform regular vulnerability assessments?

59. Do the Terms of Service or SLA include a compensation clause for breaches caused by vulnerabilities in the service? If yes, does it provide an adequate level of compensation

should a breach occur?

3.5.5 Encryption

Encryption is often presented as the solution for addressing confidentiality risks within the cloud. There are however, a number of important limitations that should be understood and considered by agencies planning adoption of cloud services. Agencies must determine their specific requirements for protecting information using encryption. Careful consideration must be given to:

- What information needs to be encrypted? All information held by the cloud service or only certain data types, or database rows, columns or entities?
- Why does the information need to be encrypted? For example, is encryption required to achieve compliance with a policy or standard?
- How should the information be encrypted? For example, what protocols and algorithms should be used?
- Who will encrypt the information and manage the encryption keys? The agency or the service provider?
- Where should the information be encrypted and decrypted? Within the agency, on the client devices or within the cloud service?
- When does the information need to be encrypted and decrypted? In transit, by the application (message encryption) and/or at rest?

While encryption is an effective control for protecting the confidentiality of data at rest, for data to be processed by a business rule within an information system, generally it must be unencrypted. As a result, it may be impractical or impossible to encrypt data stored within a cloud service that actually processes information (as opposed to simple storage).

Where a cloud service is capable of storing data in an encrypted format it is important to know which party (the agency or the service provider) is responsible for managing the encryption keys. It is important to note that if the service provider has access to, or manages, the encryption keys then they will be able to decrypt and access the information held in the cloud service. This has data sovereignty implications if encryption is used to treat risks related to information being stored outside New Zealand.

The party that manages the encryption keys must have an effective key management plan. Key management is essential to ensure that encryption keys are protected from being compromised, which could result in unauthorised disclosure or the agency no longer being able to access its information. It may also affect an agency's ability to meet its obligations under the Official Information Act 1982 and the Public Records Act 2005. The NZISM specifies the key management practices required to effectively manage cryptographic keys.

The interception of data in transit is an inherent risk whenever sensitive information traverses a network, especially a network not owned or managed by the agency such as the Internet or a service provider's network. Agencies must ensure that the cloud service encrypts all sensitive data in transit (including authentication credentials) using only approved encryption protocols and algorithms. Agencies relying on encryption should consider whether the encryption protocol, algorithm and key length used are appropriate. The NZISM specifies the encryption protocols and algorithms (together with recommended key lengths) that are approved for use by agencies for specific information classifications.

Key Considerations

60. Have requirements for the encryption of the information that will be placed in the cloud service been determined?

61. Does the cloud service use only approved encryption protocols and algorithms (as defined in the NZISM)?

62. Which party is responsible for managing the cryptographic keys?

63. Does the party responsible for managing the cryptographic keys have a key management plan that meets the requirements defined in the NZISM?

3.5.6 Cloud Service Provider Insider Threat

Unauthorised access to sensitive information by the service provider's employees is a common concern for organisations planning to use cloud services. The controls required to manage this risk are no different from those used to protect against malicious insiders within the agency or a traditional outsource provider.

Agencies should ascertain whether the service provider has appropriate procedures in place to ensure its personnel are reliable, trustworthy and do not pose a security risk to its clients. The level of assurance available to agencies may vary significantly depending on the physical location of the service provider's service and its employees. For example, a New Zealand based service provider will be able to perform a standard Ministry of Justice criminal history check for all employees and require staff that manage system components that store, process or transmit the agency's data to gain New Zealand Security Intelligence Service security clearance (e.g. CONFIDENTIAL, SECRET or TOP SECRET). However, where a service is delivered or supported from another country these New Zealand specific checks will not be possible. In such circumstances agencies must consider whether the alternatives available to the service provider can provide an equivalent level of assurance.

Whilst vetting may prevent a service provider from employing someone that has a history of being untrustworthy, it does have its limitations. For example, vetting that reveals a criminal record may result in a potential employee being rejected. However, candidates that are untrustworthy but have never been caught or haven't been convicted may not be identified. Similarly, a previously trustworthy employee may become untrustworthy if they become disgruntled or their personal circumstances change. These risks can be effectively managed if the service provider logs and monitors employees' activities and enforces separation of duties so that any malicious activity requires collusion from multiple sources making it less likely.

Logging and monitoring employees' activities is an important control to manage the risks associated with malicious insiders. Logging should cover all relevant activities performed by the service provider's employees that have logical or physical access to equipment or media that contains customer data. The service provider should monitor and review logs to identify any suspicious activity that requires investigation. In addition to this, duties should be separated to ensure that logs are protected from unauthorised modification and deletion (e.g. the administrator of a service component should not be granted modify or delete rights to the Security Information Event Monitoring (SIEM) service).

Key Considerations

64. Does the service provider undertake appropriate pre-employment vetting for all staff that have access to customer data? Does the service provider perform on-going checks during the period of employment?

65. If the service provider is dependent on a third-party to deliver any part of their service, does the third-party undertake appropriate pre-employment vetting for all staff that have access to customer data?

66. Does the service provider have a SIEM service that logs and monitors all logical access to customer data?

67. Does the service provider enforce separation of duties to ensure that audit logs are protected against unauthorised modification and deletion?

68. Do the Terms of Service or SLA require the service provider to report unauthorised access to customer data by its employees? If yes, is the service provider required to provide details about the incident to affected customers to enable them to assess and manage the associated impact?

3.5.7 Data Persistence

It can be difficult to permanently delete data from a multi-tenanted cloud service when the organisation scales down or terminates its use of the service. If data is not securely deleted a future compromise of the service may still expose agency information. Similar issues arise if the service provider does not have processes to ensure that ICT equipment and storage media (e.g. hard disk drives, backup tapes etc.) are securely wiped before redeployment or disposal. Therefore it is essential that organisations establish that the service provider has robust and demonstrable data destruction and disposal processes in place.

Key Considerations

69. Does the service provider have an auditable process for the secure sanitisation of storage media before it is made available to another customer?

70. Does the service provider have an auditable process for secure disposal or destruction of ICT equipment and storage media (e.g. hard disk drives, backup tapes etc.) that contain customer data?

3.5.8 Physical Security

Physical security controls are vital to ensure that information is physically protected from unauthorised access by both malicious service provider personnel and third parties. Effective information security is dependent on the efficacy of the physical controls implemented to protect the service provider's offices, datacentres and physical assets.

SIGS, the NZISM and the Protective Security Manual (PSM) define the minimum physical security controls that must be in place to adequately protect official information based on its classification.

However, as discussed it may not be possible or practical to directly assess the physical controls that the service provider has implemented to protect its customers data within a cloud service. An agency may be limited to reviewing a third party audit report.

Key Considerations

71. If it is practical to do so (i.e. the datacentre is within New Zealand), can the service provider's physical security controls be directly reviewed or assessed by the agency? If no, will the service provider allow the agency to review of a recent third party audit report (e.g. ISO 27001 or ISAE 3402 SOC 2 Type II) that includes an assessment of their physical security controls?

72. Do the service provider's physical security controls meet the minimum requirements as defined in the New Zealand government's security guidelines[12] to protect the information stored in the cloud service?

3.6 Data Integrity

Service providers can provide significantly different levels of protection against data loss or corruption. Some providers include data backup services as part of the base service offering, others offer them as an additional cost service and some do not offer them at all (e.g. Google Apps for Business does not provide any back-up services without a subscription to Google Apps Vault at additional cost). As a result, it is important to identify what level of protection the service provider offers and to assess whether or not they meet the agency's business requirements for recovering from data loss and corruption incidents.

It is essential to identify how the service provider protects its customers from data loss or corruption as it can indicate the level of protection provided. If the service provider replicates customer data to another datacentre in near real-time (e.g. every 5 minutes) a corruption could be replicated before it is detected. Similarly, if data is backed-up to tape on a daily basis then a Recovery Point Objective (RPO) of less than 24 hours may not be possible.

Agencies should ascertain the level of granularity offered for data restoration (e.g. can a single file or email be restored or are customers limited to restoring an entire mailbox or database?). In addition to this, they should identify and understand the process for initiating a restore. For example, can a user restore an email or file they have accidentally deleted or will an authorised staff member need to log a call with the service provider?

Data loss or corruption could lead to information being permanently lost. This may mean that agencies are unable to meet their obligations under the Public Records Act 2005 and the Official Information Act 1982. Agencies are advised to assess whether the planned data backup and archiving strategy supports their compliance efforts. Agencies without specialised knowledge in these Acts are encouraged to seek advice from Archives New Zealand and/or the Ministry of Justice to ensure compliance.

[12] Security in the Government Sector (SIGS), the New Zealand Information Security Manual (NZISM) and the Protective Security Manual (PSM)

It is important to realise that the use of cloud services may not preclude the need for an agency to develop, implement and test its own data backup strategy to ensure that it can sufficiently recover from an incident that results in data loss or corruption.

Key Considerations

73. Does the service provider provide data backup or archiving services as part of their standard service offering to protect against data loss or corruption? If not, do they offer data backup or archiving services as an additional service offering to protect against data loss and corruption?

74. How are data backup and archiving services provided?

75. Does the SLA specify the data backup schedule?

76. Does the data back-up or archiving service ensure that business requirements related to protection against data loss are met? (I.e. does the service support the business Recovery Point Objective?)

77. What level of granularity does the service provider offer for data restoration?

78. What is the service provider's process for initiating a restore?

79. Does the service provider regularly perform test restores to ensure that data can be recovered from backup media?

80. Does the agency need to implement a data backup strategy to ensure that it can recover from an incident that leads to data loss or corruption?

81. Does the proposed data backup and archiving strategy support the agency in meeting its obligations under the Public Records Act and Official Information Act?

3.7 Availability

3.7.1 Service Level Agreement

The service provider's SLA typically specifies the level of expected availability performance as a percentage. It is important for agencies to understand exactly what the defined percentage means and to assess whether or not these levels meet the requirements for availability (e.g. 99.9% up time over a year allows for up to 9 hours of unscheduled outages without breaching the SLA).

The SLA should include the details of any scheduled outage windows. This will ensure that the service provider cannot schedule long outages (including emergency outages) with little or no notification without breaching the SLA.

Where scheduled outage windows are defined in the SLA they should be reviewed to ensure that they will not have an adverse impact on business operations. For example, if an SLA includes a 3 hour scheduled outage on the first Tuesday of each month between 17:00 and 20:00 Eastern Daylight Time, the outage would occur between 10:00 and 13:00 on Wednesday in New Zealand. Some service providers use technologies to enable them to perform maintenance activities without

the need for an outage, however, agencies should not assume that this is the case simply because scheduled outages are not defined in the SLA.

Another important consideration is the adequacy of the compensation provided if the SLA is breached and the method for calculating penalties over a service period. Typically an SLA for cloud services will specify minimal compensation such as service credits or discounted invoices. Agencies should review any compensation clauses taking into account the impact on the business if the service was unavailable to determine if the level of reparation is sufficient.

Key Considerations

82. Does the SLA include an expected and minimum availability performance percentage over a clearly defined period? If yes, are the business requirements for availability met? (I.e. does the service support the business's Recovery Time Objective and Acceptable Interruption Window?)

83. Does the SLA include defined, scheduled outage windows?

    a. If yes, do the specified outage windows affect New Zealand business operations?

    b. If no, has the service provider implemented technologies that enable them to perform maintenance activities without the need for an outage?

84. Does the SLA include a compensation clause for a breach of the guaranteed availability percentages? If yes, does this provide an adequate level of compensation should the service provider breach the SLA?

3.7.2 Denial of Service Attacks

Denial of Service (DoS) attacks are an inherent risk for all Internet facing services. The use of cloud services may increase the risk of such an attack eventuating as the aggregation of multiple agencies onto a single service may present a more attractive target for attackers. Similarly, an agency may suffer associated or collateral damage in an attack against a service provider or a co-tenant. A DoS attack may be launched against the service provider or the agency itself.

Typically it is difficult to protect against traffic based DoS attacks as they are intended to consume all the available resources and effective defences rely on blocking the source of the attack as close to the attackers location as possible. However, the use of cloud services may lessen the impact of some forms of DoS attacks as service providers have spare network bandwidth and computing capacity. In addition to this some service providers use protocols and technologies (e.g. Anycast, Application Delivery Networks and Content Delivery Networks) together with geographically dispersed datacentres to distribute network traffic and computer processing around the world.

The elastic nature of cloud services may also cause financial impacts. A successful DoS attack may force a service to scale exponentially resulting in abnormally high charges for resource use. This is usually referred to as Economic Denial of Service (EDoS) or bill shock. Agencies using cloud services that scale to meet demand can effectively reduce the risk of unexpected charges by ensuring that they set boundaries to limit the resources that can be consumed to those required to meet their anticipated peak usage.

Key Considerations

85. Does the service provider utilise protocols and technologies that can protect against DDoS attacks? If yes, does enabling the service provider's DDoS protection services affect the answer to questions 15, 16 and 17?

86. Can the agency specify or configure resource usage limits to protect against EDoS/bill shock?

3.7.3 Network Availability and Performance

The availability and performance of cloud services are heavily dependent on the supporting network infrastructure. The available bandwidth, latency, reliability and resiliency of local and international network connections could have a significant impact on user experience.

Agencies should evaluate the network connectivity between their users and the cloud service to ensure availability and performance requirements are met. This may be difficult if public networks (such as the Internet) are utilised in the delivery of the service, however, agencies should confirm that the network services they directly manage, or subscribe to, provide an adequate level of availability and bandwidth, together with sufficiently low latency and packet loss to meet the needs of the business.

Key Considerations

87. Do the network services directly managed, or subscribed to by the agency provide an adequate level of availability?

88. Do the network services directly managed, or subscribed to by the agency provide an adequate level of redundancy/fault tolerance?

89. Do the network services directly managed, or subscribed to by the agency provide an adequate level of bandwidth (network throughput)?

90. Is the latency between the agency network(s) and the service provider's service at levels acceptable to achieve the desired user experience? If no, is the latency occurring on the network services directly managed, or subscribed to by the agency? Can the issue be resolved either by the network service provider or the agency?

91. Is the packet loss between the agency network(s) and the service provider's service at levels acceptable to achieve the desired user experience? If no, is the packet loss occurring on a network services directly managed, or subscribed to by the agency? Can the issue be resolved either by the network service provider or the agency?

3.7.4 Business Continuity and Disaster Recovery

The use of cloud services introduces a reliance on the service provider's business continuity and disaster recovery plans. Therefore it is important to confirm that the service provider has adequate plans in place and to understand the level of continuity and recovery provided by them. It is also

important to realise that the use of cloud services does not preclude the need for an agency to develop, implement and test its own business continuity and disaster recovery plans to ensure that it can continue to operate during a service outage.

As the cloud computing market is relatively immature, agencies should consider how they would recover business operations should a service provider go out of business or withdraw a service. They should ensure that the service provider uses common or de facto data format standards and provides a method to extract data in a format usable by the agency.

Key Considerations

92. Does the service provider have business continuity and disaster recovery plans?

93. Will the service provider permit the agency to review of its business continuity and disaster recovery plans?

94. Do the service provider's plans cover the recovery of the agency data or only the restoration of the service?

95. If the service provider's plans cover the restoration of agency data, is the recovery of customer data prioritised? If so, how? Are customers prioritised based on size and contract value?

96. Does the service provider formally test its business continuity and disaster recovery plans on a regular basis? If yes, how regularly are such tests performed? Will they provide customers with a copy of the associated reports?

97. Does the agency have its own business continuity and disaster recovery plan in place to ensure that it can recover from a service outage, the service provider going out of business or withdrawing the service?

98. Does the agency require its own data backup strategy to ensure that it can recover from a service outage, the service provider going out of business or withdrawing the service?

99. Are the backups (whether performed by the service provider or the agency) encrypted using an approved encryption algorithm and appropriate key length?

3.8 Incident Response and Management

The level of visibility and control of security incidents is likely to vary considerably across the cloud service models. The service provider is typically responsible for all incident management activities involving a SaaS solution, however, when an incident relates to a system located on an IaaS solution the customer is usually responsible for the incident management activities related to the platform, application and data and the service provider is only responsible for the activities directly related to the infrastructure components they manage. Similarly, the cloud deployment model (i.e. public, private, community or hybrid) adopted by the agency could significantly affect its visibility and control over the incident management activities. For example, customers of public cloud services normally have less visibility and control over incident management activities than those that have implemented a private cloud.

It is not reasonable to expect service providers to implement a separate incident response and management plan for each of their customers, therefore agencies need to gain an appropriate level of assurance that a service provider is capable of effectively and efficiently responding to an information security incident, as even the most meticulously planned, implemented and managed preventative controls can fail to stop a risk from eventuating. As a result, agencies need to review the service provider's Terms of Service and SLA to identify what, if any, support they provide to their customers during an information security incident.

Regardless of the service or deployment model, the use of cloud services does not preclude the need for an agency to have its own incident response and management process and plans. In fact, these plans are essential as they define how the agency will handle the tasks it is responsible for including roles and responsibilities, key contacts, incident definitions and notification criteria, escalation channels, evidence collection and preservation and post incident activities.

Key Considerations

100. Does the service provider have a formal incident response and management process and plans that clearly define how they detect and respond to information security incidents? If yes, will they provide the agency with a copy of their process and plans to enable it to determine if they are sufficient?

101. Does the service provider test and refine its incident response and management process and plans on a regular basis?

102. Does the service provider engage its customers when testing its incident response and management processes and plans?

103. Does the service provider provide its staff with appropriate training on incident response and management processes and plans to ensure that they respond to incidents in an effective and efficient manner?

104. Does the service provider's Terms of Service or SLA clearly define the support they will provide to the agency should an information security incident arise? For example, does the service provider:

    a. Notify customers when an incident that may affect the security of their information or interconnected systems is detected or reported?

    b. Specify a point of contact and channel for customers to report suspected information security incidents?

    c. Define the roles and responsibilities of each party during an information security incident?

    d. Provide customers with access to evidence (e.g. time stamped audit logs and/or forensic snapshots of virtual machines etc.) to enable them to perform their own investigation of the incident?

    e. Provide sufficient information to enable the agency to cooperate effectively with

an investigation by a regulatory body, such as the Privacy Commissioner or the Payment Card Industry Security Standards Council (PCI SSC)?

f. Define which party is responsible for the recovery of data and services after an information security incident has occurred?

g. Share post incident reports with affected customers to enable them to understand the cause of the incident and make an informed decision about whether to continue using the cloud service?

h. Specify in the contract limits and provisions for insurance, liability and indemnity for information security incidents? (Note: it is recommended that agencies carefully review liability and indemnity clauses for exclusions.)

4 Appendix A – Cloud Considerations Questions

This section provides a list of the questions from Security and Privacy Considerations section of this document. The questions retain the original numbering to enable agencies to easily find the related paragraphs in the document that contain detailed information about the associated information security and privacy consideration.

Table 1 - Cloud Considerations Questions

1. Who is the business owner of the information?
2. What are the business processes that are supported by the information?
3. What is the security classification of the information based on the NZ government guidelines for protection of official information?
4. Are there any specific concerns related to the confidentiality of the information that will be stored or processed by the cloud service?
5. Does the data include any personal information?
6. Who are the users of the information?
7. What permissions do the users require to the information? (e.e. read, write, modify and/or delete)
8. What legislation applies to the information? (e.g. Privacy Act 1993, Official Information Act 1982, Public Records Act 2005)
9. What contractual obligations apply to the information? (e.g. Payment Card Industry Data Security Standard (PCI DSS))
10. What would the impact on the business be if the information were disclosed to an unauthorised party?
11. What would the impact on the business be if the integrity of the information were compromised?
12. Does the agency have incident response and management plans in place to minimise the impact of an unauthorised disclosure?
13. What would the impact on the business be if the information were unavailable? a. What is the maximum amount of data loss that can be tolerated after a disruption has occurred? This is used to define the Recovery Point Objective. b. What is the maximum period of time before which the minimum levels of services must be restored after a disruption has occurred? This is used to define the Recovery Time

Objective. c. What is the maximum period of time before which the full service must be restored to avoid permanently compromising the business objectives? This is used to define the Acceptable Interruption Window.
14. Where is the registered head office of the service provider?
15. Which countries are the cloud services delivered from?
16. In which legal jurisdictions will the agency's data be stored and processed?
17. Does the service provider allow its customers to specify the locations where their data can and cannot be stored and processed?
18. Does the service have any dependency on any third parties (e.g. outsourcers, sub-contractors or another service provider) that introduce additional jurisdictional risks? If yes, ask the service provider to provide the following details for each third party involved in the delivery of the service: a. The registered head office of the third party; b. The country or countries that their services are delivered from; and c. The access that they have to client data stored, processed and transmitted by the cloud service.
19. Have the laws of the country or countries where the data will be stored and processed been reviewed to assess how they could affect the security and/or privacy of the information?
20. Do the laws actually apply to the service provider and/or its customer's information? (e.g. some privacy laws exempt certain types of businesses or do not apply to the personal information of foreigners.)
21. Do the applicable privacy laws provide an equivalent, or stronger, level of protection than the Privacy Act 1993? If no, are customers able to negotiate with the service provider to ensure that the equivalent privacy protections are specified in the contract?
22. How does the service provider deal with requests from government agencies to access customer information? a. Do they only disclose information in response to a valid court order? b. Do they inform their customers if they have to disclose information in response to such a request? c. Are they prevented from informing customers that they have received a court order requesting access to their information?

23. Does the data that will be stored and processed by the cloud service include personal information as defined in the Privacy Act 1993?

If no, skip to question 28.

24. Has a Privacy Impact Assessment (PIA) been completed that identifies the privacy risks associated with the use of the cloud service together with the controls required to effectively manage them?

25. Is the service provider's use of personal information clearly set out in its privacy policy? Is the policy consistent with the agency's business requirements?

26. Does the service provider notify its customers if their data is accessed by, or disclosed to, an unauthorised party? Does this include providing sufficient information to support cooperation with an investigation by the Privacy Commissioner?

27. Who can the agency, its staff and/or customers complain to if there is a privacy breach?

28. Does the service provider negotiate contracts with their customers or must they accept a standard Terms of Service?

29. Does the service provider's Terms of Service and SLA clearly define how the service protects the confidentiality, integrity and availability of official information and the privacy of all personally identifiable information?

30. Does the service provider's Terms of Service specify that the agency will retain ownership of its data?

31. Will the service provider use the data for any purpose other than the delivery of the service?

32. Is the service provider's service dependent on any third-party services?

33. Does the service provider's Terms of Service allow the agency to directly audit the implementation and management of the security measures that are in place to protect the service and the data held within it?

    a. If yes, does this include performing vulnerability scans and penetration testing of the service and the supporting infrastructure?

    b. If no, does the service provider undergo formal regular assessment against an internationally recognised information security standard or framework by an independent third-party? (e.g. are they certified as being compliant with ISO/IEC 27001? Have they undergone an ISAE 3402 SOC 2 Type II?)

34. Will the service provider allow the agency to thoroughly review their recent audit reports before signing up for their service? (e.g. will they provide us with the Statement of Applicability together with a copy of the full audit reports from their external auditor, and the results of any recent internal audits?)
35. Will the service provider enable potential customers to perform reference checks by providing the contact details of two or more of its current customers?
36. Is there a completed CAIQ or CMM report for the service provider in the CSA STAR?
37. Has the service provider undergone a CSA STAR Certification and/or Attestation? Have they published the outcome of the audit?
38. Has the service provider published a completed Cloud Computing Code of Practice?
39. What additional assurance activities must be performed to complete the certification and accreditation of the cloud service?
40. Does the agency have an identity management strategy that supports the adoption of cloud services? If yes, does the cloud service support the agency's identity management strategy?
41. Is there an effective internal process that ensures that identities are managed throughout their lifecycle?
42. Is there an effective audit process that is actioned at regular intervals to ensure that user accounts are appropriately managed?
43. Have the controls required to manage the risks associated with the ubiquitous access provided by the cloud been identified?
44. Does the cloud service meet those control requirements?
45. Is there a higher level of assurance required that the party asserting an identity is the authorised user of the account when authenticating to the service? (I.e. is multi-factor authentication necessary?)
46. Will the service provider allow the agency to review a recent third-party audit report (e.g. ISO 27001 or ISAE 3402 SOC 2 Type II) that includes an assessment of the security controls and practices related to virtualisation and separation of customer's data?

47. Will the service provider permit customers to undertake security testing (including penetration tests) to assess the efficacy of the access controls used to enforce separation of customer's data?

48. Does the service provider's customer registration processes provide an appropriate level of assurance in line with the value, criticality and sensitivity of the information to be placed in the cloud service?

49. Are there appropriate build and hardening standards defined and documented for the service components the agency is responsible for managing?

50. Can the agency deploy operating systems and applications in accordance with internal build or hardening standards? If no, does the service provider have appropriate build and hardening standards that meet our security requirements?

    a. Does the virtual image include a host-based firewall configured to only allow the ingress and egress (inbound and outbound) traffic necessary to support the service?

    b. Does the service provider allow host-based intrusion detection and prevention service (IDS/IDP) agents to be installed within the virtual machines?

51. Does the service provider perform regular tests of its security processes and controls? Will they provide customers with a copy of the associated reports?

52. Can a penetration test of the service be performed to ensure that it has been securely deployed?

53. Is the service provider responsible for patching all components that make up the cloud service? If no, has the agency identified which components the service provider is responsible for and which it is responsible for?

54. Does the service provider's Terms of Service or SLA include service levels for patch and vulnerability management that includes a defined the maximum exposure window?

55. Does the agency currently have an effective patch and vulnerability management process?

56. Has the agency ensured that all of the components that it is responsible for have been incorporated into its patch and vulnerability management process?

57. Is the agency subscribed to, or monitoring, appropriate sources for vulnerability and patch alerts for the components that it is are responsible for?

58. Does the service provider allow its customers to perform regular vulnerability assessments?
59. Do the Terms of Service or SLA include a compensation clause for breaches caused by vulnerabilities in the service? If yes, does it provide an adequate level of compensation should a breach occur?
60. Have requirements for the encryption of the information that will be placed in the cloud service been determined?
61. Does the cloud service use only approved encryption protocols and algorithms (as defined in the NZISM)?
62. Which party is responsible for managing the cryptographic keys?
63. Does the party responsible for managing the cryptographic keys have a key management plan that meets the requirements defined in the NZISM?
64. Does the service provider undertake appropriate pre-employment vetting for all staff that have access to customer data? Does the service provider perform on-going checks during the period of employment?
65. If the service provider is dependent on a third-party to deliver any part of their service, does the third-party undertake appropriate pre-employment vetting for all staff that have access to customer data?
66. Does the service provider have a Security Information Event Monitoring (SIEM) service that logs and monitors all logical access to customer data?
67. Does the service provider enforce separation of duties to ensure that logs are protected against unauthorised modification and deletion?
68. Do the Terms of Service or SLA require the service provider to report unauthorised access to customer data by its employees? If yes, is the service provider required to provide details about the incident to affected customers to enable them to assess and manage the associated impact?
69. Does the service provider have an auditable process for the secure sanitisation of storage media before it is made available to another customer?
70. Does the service provider have an auditable process for secure disposal or destruction of ICT equipment and storage media (e.g. hard disk drives, backup tapes etc.) that contain customer data?

71. If it is practical to do so (i.e. the datacentre is within New Zealand), can the service provider's physical security controls be directly reviewed or assessed by the agency? If no, will the service provider the agency to review a recent third party audit report (e.g. ISO 27001 or ISAE 3402 SOC 2 Type II) that includes an assessment of their physical security controls?
72. Do the service provider's physical security controls meet the minimum requirements as defined in the New Zealand government's security guidelines[13] to protect the information stored in the cloud service?
73. Does the service provider provide data backup or archiving services as part of their standard service offering to protect against data loss or corruption? If no, do they offer data backup or archiving services as an additional service offering to protect against data loss and corruption?
74. How are data backup and archiving services provided?
75. Does the SLA specify the data backup schedule?
76. Does the data back-up or archiving service ensure that our business requirements related to protection against data loss are met? (I.e. does the service support the business's Recovery Point Objective?)
77. What level of granularity does the service provider offer for data restoration?
78. What is the service provider's process for initiating a restore?
79. Does the service provider regularly perform test restores to ensure that data can be recovered from backup media?
80. Does the agency need to implement a data backup strategy to ensure that it can recover from an incident that leads to data loss or corruption?
81. Does the proposed data backup and archiving strategy support the agency in meeting its obligations under the Public Records Act and Official Information Act?
82. Does the SLA include an expected and minimum availability performance percentage over a clearly defined period? If yes, are the business requirements for availability met? (I.e. does the service support the business's Recovery Time Objective and Acceptable Interruption Window?)

[13] Security in the Government Sector (SIGS), the New Zealand Information Security Manual (NZISM) and the Protective Security Manual (PSM)

83. Does the SLA include defined scheduled outage windows? a. If yes, do the days and time specified for scheduled outages windows affect business operations? b. If no, has the service provider implemented technologies that enable them to perform maintenance activities without the need for an outage?
84. Does the SLA include a compensation clause for a breach of the guaranteed availability percentages? If yes, does this provide an adequate level of compensation should the service provider breach the SLA?
85. Does the service provider utilise protocols and technologies that can protect against DDoS attacks? If yes, does enabling the service provider's DDoS protection services affect the answer to questions 15, 16 and 17?
86. Can the agency specify or configure resource usage limits to protect against EDoS/bill shock?
87. Do the network services directly managed, or subscribed to by the agency provide an adequate level of availability?
88. Do the network services directly managed, or subscribed to by the agency provide an adequate level of redundancy/fault tolerance?
89. Do the network services directly managed, or subscribed to by the agency provide an adequate level of bandwidth (network throughput)?
90. Is the latency between the agency network(s) and the service provider's service at levels acceptable to achieve the desired user experience? If no, is the latency occurring on the network services directly managed, or subscribed to by the agency? Can the issue be resolved either by the network service provider or the agency?
91. Is the packet loss between the agency network(s) and the service provider's service at levels acceptable to achieve the desired user experience? If no, is the packet loss occurring on a network services directly managed, or subscribed to by the agency? Can the issue be resolved either by the network service provider or the agency?
92. Does the service provider have business continuity and disaster recovery plans?
93. Will the service provider permit the agency to review of its business continuity and disaster recovery plans?
94. Do the service provider's plans cover the recovery of the agency data or only the restoration of the service?

95. If the service provider's plans cover the restoration of agency data, is the recovery of customer data prioritised? If so, how? Are customers prioritised based on size and contract value?
96. Does the service provider formally test its business continuity and disaster recovery plans on a regular basis? If yes, how regularly are such tests performed? Will they provide customers with a copy of the associated reports?
97. Does the agency have its own business continuity and disaster recovery plan in place to ensure that it can recover from a service outage, the service provider going out of business or withdrawing the service?
98. Does the agency require its own data backup strategy to ensure that it can recover from a service outage, the service provider going out of business or withdrawing the service?
99. Are the backups (whether performed by the service provider or the agency) encrypted using an approved encryption algorithm and appropriate key length?
100. Does the service provider have a formal incident response and management process and plans that clearly define how they detect and respond to information security incidents? If yes, will they provide the agency with a copy of their process and plans to enable it to determine if they are sufficient?
101. Does the service provider test and refine its incident response and management process and plans on a regular basis?
102. Does the service provider engage its customers when testing its incident response and management process and plans?
103. Does the service provider provide its staff with appropriate training on incident response and management process and plans to ensure that they respond to incidents in an effective and efficient manner?
104. Does the service provider's Terms of Service or SLA clearly define the support they will provide to the agency should an information security incident arise? For example, does the service provider: a. Notify customers when an incident that may affect the security of their information or interconnected systems is detected or reported? b. Specify a point of contact and channel for customers to report suspected information security incidents? c. Define the roles and responsibilities of each party during an information security incident? d. Provide customers with access to evidence (e.g. time stamped audit logs and/or forensic snapshots of virtual machines etc.) to enable them to perform their own investigation of

the incident?

e. Provide sufficient information to enable the agency to cooperate effectively with an investigation by a regulatory body, such as the Privacy Commissioner or the Payment Card Industry Security Standards Council (PCI SSC)?

f. Define which party is responsible for the recovery of data and services after an information security incident has occurred?

g. Share post incident reports with affected customers to enable them to understand the cause of the incident and make an informed decision about whether to continue using the cloud service?

h. Specify in the contract limits and provisions for insurance, liability and indemnity for information security incidents? (**Note**: it is recommended that agencies carefully review liability and indemnity clauses for exclusions.)

5 Appendix B – Additional Resources

Agencies adopting cloud computing may find following resources useful (**Note**: Internal Affairs accepts no liability for the content of these resources):

Department of Prime Minister and Cabinet (DPMC)
Security in the Government Sector (SIGS) 2002
http://www.security.govt.nz/assets/media/Security_in_the_Government_Sector_2002.pdf

New Zealand Security Intelligence Service (NZSIS)
Protective Security Manual (PSM)
(Government departments and agencies can request a copy of the PSM from NZSIS)

Government Communications Security Bureau (GCSB)
New Zealand Information Security Manual (NZISM) 2011 V1.01
http://www.gcsb.govt.nz/newsroom/nzism/NZISM_2011_Version_1.01.pdf

Office of the Privacy Commissioner (OPC)
Cloud Computing a Guide to Making the Right Choices - February 2013
http://privacy.org.nz/assets/Files/Brochures-and-pamphlets-and-pubs/OPC-Cloud-Computing-guidance-February-2013.pdf

Privacy Impact Assessment Handbook – June 2007
http://privacy.org.nz/assets/Uploads/Privacy-Impact-Assessment-Handbook-June2007.pdf

State Services Commission (SSC)
Government Use of Offshore Information and Communication Technologies (ICT) Service Providers - Advice on Risk Management - April 2009
http://ict.govt.nz/assets/Uploads/Drupal/offshore-ICT-service-providers-april-2007.pdf

Cloud Security Alliance (CSA)
Security Guidance for Critical Areas of Focus in Cloud Computing V3.0
https://cloudsecurityalliance.org/guidance/csaguide.v3.0.pdf

Top Threats to Cloud Computing V1.0
https://cloudsecurityalliance.org/topthreats/csathreats.v1.0.pdf

European Network and Information Security agency (ENISA)
Cloud Computing Information Assurance Framework
http://www.enisa.europa.eu/activities/risk-management/files/deliverables/cloud-computing-information-assurance-framework/at_download/fullReport

Cloud Computing Risk Assessment
http://www.enisa.europa.eu/activities/risk-management/files/deliverables/cloud-computing-risk-assessment/at_download/fullReport

Critical Cloud Computing – A CIIP perspective on cloud computing services
http://www.enisa.europa.eu/activities/Resilience-and-CIIP/cloud-computing/critical-cloud-computing/at_download/fullReport

Security and Resilience in Governmental Clouds
http://www.enisa.europa.eu/activities/risk-management/emerging-and-future-risk/deliverables/security-and-resilience-in-governmental-clouds/at_download/fullReport

National Institute of Science and Technology (NIST)

NIST Cloud Computing Reference Architecture

http://www.nist.gov/customcf/get_pdf.cfm?pub_id=909505

NIST Cloud Computing Security Reference Architecture (DRAFT)

http://collaborate.nist.gov/twiki-cloud-computing/pub/CloudComputing/CloudSecurity/NIST_Security_Reference_Architecture_2013.05.15_v1.0.pdf

The NIST Definition of Cloud Computing

http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf

The Committee of Sponsoring Organizations of the Treadway Commission (COSO)

Enterprise Risk Management for Cloud Computing

http://www.coso.org/documents/Cloud%20Computing%20Thought%20Paper.pdf