Defence Research and Development Canada
Recherche et développement pour la défense Canada

DEFENCE R&D DÉFENSE

Data protection in multi-tenant cloud environments

A technical brief

Alan Magar

The scientific or technical validity of this Contract Report is entirely the responsibility of the contractor and the contents do not necessarily have the approval or endorsement of Defence R&D Canada.

Defence R&D Canada – Ottawa

Contract Report
DRDC Ottawa CR 2012-108
December 2012

Canada

Data protection in multi-tenant cloud environments

A technical brief

Alan Magar
TRM Technologies, Inc.

Prepared By:

TRM Technologies Inc.
280 Albert Street, Suite 1000 (10th Floor)
Ottawa, ON, K1P 5G8

Project Manager: Darcy Simmelink (613) 998-1451
Contract Number: W7714-08FE01
Contract Scientific Authority: Kathryn Perrett (613) 993-5132

The scientific or technical validity of this Contract Report is entirely the responsibility of the contractor and the contents do not necessarily have the approval or endorsement of Defence R&D Canada.

Defence R&D Canada – Ottawa

Contract Report
DRDC Ottawa CR 2012-108
December 2012

Contract Scientific Authority

Original signed by Kathryn Perrett

Kathryn Perrett

Defence Scientist, Cyber Operations Section

Approved by

Original signed by Julie Lefebvre

Julie Lefebvre

Head, Cyber Operations Section

Approved for release by

Original signed by Chris McMillan

Chris McMillan

Head, Document Review Panel

© Her Majesty the Queen in Right of Canada, as represented by the Minister of National Defence, 2012

© Sa Majesté la Reine (en droit du Canada), telle que représentée par le ministre de la Défense nationale, 2012

Abstract

This technical brief examines data protection within multi-tenant cloud computing environments. Specifically, this technical brief explores data protection strategies that can be employed within cloud environments to mitigate potential threats from other cloud consumers, as well as from third-party cloud providers. Further research in the areas of federated identity, authorization and encryption are required to address potential implementation issues that could arise in cloud computing environments.

Résumé

Le présent compte rendu technique porte sur la protection des données au sein des environnements en nuage communs à plusieurs secteurs. Il traite plus particulièrement des stratégies que l'on peut employer dans les environnements en nuage afin d'atténuer les menaces potentielles en provenance des autres consommateurs de services infonuagiques, ainsi que des fournisseurs de tels services. Il faut effectuer des recherches approfondies dans les domaines de l'identité fédérale, des autorisations et du chiffrement, afin de pouvoir corriger les éventuels problèmes de mise en œuvre qui peuvent survenir au sein d'environnements de ce type.

This page intentionally left blank.

Executive summary

Data protection in multi-tenant cloud environments: A technical brief

Alan Magar; DRDC Ottawa CR 2012-108; Defence R&D Canada – Ottawa; December 2012.

Organizations are contemplating the use of cloud computing in order to outsource infrastructure management and concentrate on core competencies. It is anticipated that this transition will pay dividends for most organizations in terms of operating efficiency and agility. However, organizations considering a transition to multi-tenant cloud environments should be concerned with the risks inherent in these environments, especially as they pertain to data protection.

This technical brief examines data protection within multi-tenant cloud computing environments. Specifically, this technical brief explores data protection strategies that can be employed within cloud environments to mitigate potential threats from other cloud consumers, as well as from third-party cloud providers.

It is anticipated that organizations can protect sensitive data in cloud computing environments by employing a layered approach consisting of multiple data protection strategies. This approach would ensure data protection at all stages of the data life cycle, including data transfer, data in transit, data at rest, data in use, and data sanitization and remanence.

The challenge, in terms of data protection, is how to extend an organization's existing data protection capabilities into the cloud. Additional research in this area is required, specifically in the areas of identity federation, authorization and encryption.

Sommaire

Data protection in multi-tenant cloud environments: A technical brief

Alan Magar; DRDC Ottawa CR 2012-108; R & D pour la défense Canada – Ottawa; décembre 2012.

De nombreux organismes songent à utiliser l'infonuagique afin d'externaliser la gestion de leur infrastructure en vue de se concentrer sur leurs compétences essentielles. On s'attend à ce que cette transition soit bénéfique pour la plupart des organisations, sur le plan de l'efficience et de la souplesse opérationnelle. Cependant, ces organismes qui projettent d'effectuer la transition vers un environnement en nuage commun à plusieurs secteurs doivent également tenir compte des risques inhérents à ce type d'environnement, particulièrement ceux relatifs à la protection des données.

Le présent compte rendu technique porte sur la protection des données au sein des environnements en nuage communs à plusieurs secteurs. Il traite plus particulièrement des stratégies que l'on peut employer dans les environnements en nuage afin d'atténuer les menaces potentielles en provenance des autres consommateurs de services infonuagiques, ainsi que des fournisseurs de tels services.

On s'attend à ce que les organismes puissent protéger leurs données comprises dans un environnement en nuage en employant une approche multidimensionnelle, composée de multiples stratégies de protection de données. En effet, elle assurerait cette protection à toutes les étapes du cycle de vie des données, y compris le transfert des données, les données en migration, en utilisation, inactives ainsi que le nettoyage et la rémanence des données.

L'enjeu, en matière de protection des données, consiste à accroître la portée des capacités de protection existantes de l'organisme afin d'englober également l'environnement en nuage. Il faut mener des recherches additionnelles en ce sens, particulièrement quant à l'identité fédérale, à l'autorisation et au chiffrement.

Table of contents

List of figures

This page intentionally left blank.

1 Introduction

1.1 Background

Organizations are contemplating the use of cloud computing in order to outsource infrastructure management and concentrate on core competencies. It is anticipated that this transition will pay dividends for most organizations in terms of operating efficiency and agility. However, organizations considering a transition to multi-tenant cloud environments should be concerned with the risks inherent in these environments, especially as they pertain to data protection.

1.2 Purpose

The purpose of this technical brief is to examine data protection within multi-tenant cloud computing environments. Specifically, this technical brief will explore data protection strategies that can be employed within cloud environments to mitigate potential threats from other cloud consumers, as well as from third-party cloud providers.

1.3 Scope

This technical brief will focus exclusively on technical data protection strategies. Other data protection strategies, including those involving governance, legal, compliance and audit, are outside of the scope of this technical brief.

Furthermore, this technical brief will focus on the data protection aspects of multi-tenant cloud environments. There are a number of other security aspects related to these environments, including physical security, network security and system security; these aspects will not be addressed within this technical brief.

Organizations have different assets in the cloud. These assets typically fall into two categories: data and processes. This technical brief will examine both types of assets. However, for the purpose of the technical brief both assets will be referred to simply as data. Furthermore, this report focuses on the protection of sensitive data in the cloud. While organizations will also need to protect non-sensitive data, at least from modification and destruction, this is largely outside of the scope of this report.

2 Cloud Computing

2.1 Overview

This section will provide an overview of cloud computing that will serve as the foundation for the remainder of the technical brief. Specifically, it will examine the following aspects of cloud computing:

- Concepts;
- Threat Environment; and
- Security Objectives.

2.2 Concepts

Cloud computing is 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 technical brief has chosen to adopt the National Institute of Standards and Technologies (NIST) definitions for cloud computing. Cloud computing concepts, which are illustrated in Figure 1, include the following:

- Essential Characteristics;
- Multi-Tenancy;
- Virtualization;
- Service Models; and
- Deployment Models.

[1] *The NIST Definition of Cloud Computing* **[Reference 1]**

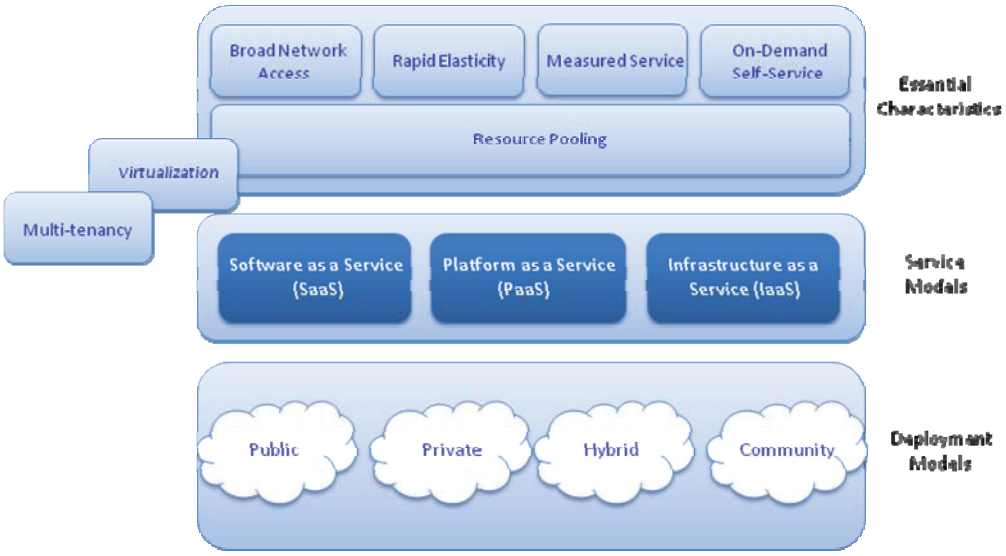

Figure 1 - Visual Model of Cloud Computing [2]

2.2.1 Essential Characteristics

Cloud computing has the following five essential characteristics:

- On-demand Self-Service – On-demand self-service is basically self-service provisioning. A consumer can automatically provision the computing resources that they require without third-party involvement;

- Broad Network Access – Broad network access refers not only to the network but the mechanisms and platforms used to access computing resources as well. The computing resources are network accessible using standard mechanisms from a variety of client platforms;

- Resource Pooling – Resource pooling refers to the sharing of computing resources amongst multiple consumers. The consumer is typically unaware of the location and identity of other consumers of the shared resources;

- Rapid Elasticity – Rapid elasticity refers to the ability to provision and de-provision computing resources commensurate with demand. The consumer can provision as many computing resources as they require for as long as they require; and

- Measured Service – Measured service refers to the fact that computing resources consumed by an organization are monitored, controlled, and reported. This allows the organization to control and optimize its use of computing resources.

[2] This figure is based on one found in *Security Guidance for Critical Areas of Focus in Cloud Computing* **[Reference 2]**.

2.2.2 Multi-Tenancy

Interestingly, multi-tenancy is not considered an essential characteristic of cloud computing by NIST. However, it has been identified as an important element of cloud computing by the Cloud Security Alliance (CSA). In order to agree upon a definition of multi-tenancy we must first define the concept of a tenant in a cloud computing context. A tenant is a generic term used to denote a customer using a particular cloud computing service to address an IT requirement. The customer can be an organization, a business unit or even an individual. Multi-tenancy in cloud environments refers to the use of the same set of resources by multiple consumers, typically from different organizations. Multi-tenant cloud computing environments provide significant cost advantages over traditional computing environments due primarily to economies of scale.

For example, ACME Corporation is using a public cloud, Cumulus, to host a number of customer-facing applications. ACME is a tenant of Cumulus. The term tenant is sufficiently generic that it can apply to ACME's sales and marketing business units, each with their own customer facing applications, as well. However, in this context ACME, including its sales and marketing business units, constitutes a single tenant. Widget Inc. and Sprocket Inc. also use Cumulus for a variety of cloud computing services. Consequently, Cumulus is a multi-tenant environment.

2.2.3 Virtualization

Virtualization is considered by many to be the foundation for cloud computing. However, as with multi-tenancy, it has not been identified as an essential characteristic of cloud computing by NIST. Virtualization, which introduces an abstraction layer between a physical resource and the service requesting the resource, allows multiple users to share the underlying physical resource. In addition, virtualization provides a degree of isolation so that users cannot interfere with each other's use of the physical resource.

2.2.4 Service Models

This section of the technical brief will examine cloud service models. These service models, which dictate the degree of control the organization has over the computing environment, include the following:

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

IaaS is the foundation of all cloud services, with PaaS building upon IaaS, and SaaS in turn building upon PaaS. All three service models are relevant to discussions in this technical brief.

2.2.4.1 Infrastructure as a Service (IaaS)

The capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary

software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, and deployed applications; and possibly limited control of select networking components (e.g., host firewalls).[3]

IaaS provides the computing infrastructure, along with storage and networking. These computing resources are typically abstracted so that consumers are provided with Virtual Machines (VMs), virtual data storage and virtual network components. Consumers manage this computing infrastructure through management Application Programming Interfaces (APIs). The provider is responsible for securing the underlying infrastructure and abstraction layers, while the consumer is responsible for the remainder of the stack. Google, IBM, VMware and Amazon.com all provide IaaS offerings.

2.2.4.2 Platform as a Service (PaaS)

The capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages, libraries, services, and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly configuration settings for the application-hosting environment.[4]

PaaS, which sits on top of the IaaS layer, provides the computing platform and solution stack. It is intended to facilitate the deployment of applications by concealing the costs and complexity of the underlying platform. Consumers can build and deliver applications on the platform using programming languages and tools that are supported by the stack. The provider is responsible for securing the platform, while the consumer is responsible for securing the applications developed and hosted on the platform. Amazon Elastic Computing Cloud (EC2), Force.com, Google App Engine and Microsoft Azure are examples of PaaS offerings.

2.2.4.3 Software as a Service (SaaS)

The capability provided to the consumer is to use the provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through either a thin client interface, such as a web browser (e.g., web-based email), or a program interface. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user- specific application configuration settings.[5]

SaaS, which sits on top of PaaS, provides the entire user experience, including the content, its presentation, the application(s), and management capabilities. It is sometimes referred to as "on-demand software" because it hosts the software centrally where it is accessed by users, normally using a web browser, over the Internet. In SaaS offerings, the security controls provided by the provider are typically negotiated into the service contract. It is worth mentioning that many SaaS

[3] *The NIST Definition of Cloud Computing* **[Reference 1]**
[4] *The NIST Definition of Cloud Computing* **[Reference 1]**
[5] *The NIST Definition of Cloud Computing* **[Reference 1]**

providers don't use VMs. Instead, these service providers leverage a single logical instance of an application that is capable of handling large numbers of tenants. Google Docs, Salesforce.com and Yahoo Mail are examples of SaaS offerings.

2.2.5 Deployment Models

This section will examine the various cloud computing deployment models. These deployment models *broadly characterize the management and disposition of computational resources for delivery of services to consumers, as well as the differentiation between classes of consumers.*[6] Specifically, this section will examine the following deployment models:

- Private Cloud;
- Community Cloud;
- Public Cloud; and
- Hybrid Cloud.

The principles of multi-tenancy are applicable to all of the deployment models, including private clouds. While private clouds are not usually considered a multi-tenant environment, they can have many of the same issues and concerns as multi-tenant environments due to the presence of a variety of users including employees, contractors, vendors, co-op students, temporary workers, etc. Although this technical brief will focus primarily on typical multi-tenant environments, such as community clouds, the other deployment models are included for the sake of completeness.

2.2.5.1 Private Cloud

The cloud infrastructure is provisioned for exclusive use by a single organization comprising multiple consumers (e.g., business units). It may be owned, managed, and operated by the organization, a third party, or some combination of them, and it may exist on or off premises.[7]

A private cloud is operated solely for a single organization. However, it can be managed by the organization or by a third party. This can be done on-premise or off-premise. Private clouds negate many of the advantages of cloud computing in that the organization is responsible for all of the infrastructure and the operational costs. The principles related to multi-tenancy are applicable to private clouds due to the variety of users, as mentioned previously, as well as the requirement for separation between business units. Many cloud service providers offer private clouds as well. For example, Amazon Virtual Private Cloud (VPC) is an EC2 instance that is completely isolated from their public cloud offering. It is basically an off-premise, virtual private network that an organization can use to host a subset of their IT infrastructure. The organization has complete control over this virtual network environment. Many other vendors, including VMware, Citrix, IBM, Oracle and Red Hat, have their own private cloud offerings.

[6] *NIST Guidelines on Security and Privacy in Public Cloud Computing* **[Reference 3]**
[7] *The NIST Definition of Cloud Computing* **[Reference 1]**

2.2.5.2 Community Cloud

The cloud infrastructure is provisioned for exclusive use by a specific community of consumers from organizations that have shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be owned, managed, and operated by one or more of the organizations in the community, a third party, or some combination of them, and it may exist on or off premises.[8]

A community cloud is a cloud infrastructure that is shared amongst a community of organizations with a common purpose. This common purpose can include a mission, security requirements, policy, or compliance considerations. The community cloud can be managed by the organizations or by a third party and may be located on-premise or off-premise. The infrastructure costs associated with a community cloud are borne by the entire community rather than a single organization. The Google government cloud and the Amazon Web Services (AWS) GovCloud are examples of community clouds. The Google government cloud is open to U.S. federal, state, and local government agencies. The AWS GovCloud is intended to allow U.S. government agencies and contractors to move sensitive workloads into the cloud by addressing their specific regulatory and compliance requirements. It is discussed in more detail in Section 3.2.1.

2.2.5.3 Public Cloud

The cloud infrastructure is provisioned for open use by the general public. It may be owned, managed, and operated by a business, academic, or government organization, or some combination of them. It exists on the premises of the cloud provider.[9]

A public cloud is a cloud infrastructure that is available to a large group or to the general public. It is owned and operated by the cloud service provider selling the cloud services. Amazon EC2, Google App Engine and Windows Azure Services Platform are all examples of public clouds.

2.2.5.4 Hybrid Cloud

The cloud infrastructure is a composition of two or more distinct cloud infrastructures (private, community, or public) that remain unique entities, but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load balancing between clouds).[10]

A hybrid cloud is a cloud infrastructure comprised of two or more clouds (private, community, or public) that are interconnected. This interconnectivity typically permits data and application portability between clouds. For example, an organization may use a private cloud to host a subset of their Information Technology (IT) resources, but integrate it with a security vendor's public cloud providing threat intelligence (e.g., Trend Micro's Smart Protection Network, Cisco IronPort SenderBase Security Network).

[8] *The NIST Definition of Cloud Computing* **[Reference 1]**
[9] *The NIST Definition of Cloud Computing* **[Reference 1]**
[10] *The NIST Definition of Cloud Computing* **[Reference 1]**

2.3 Threat Environment

Multi-tenant cloud environments are somewhat unique in that potentially sensitive data from disparate organizations is collocated or commingled on a shared computing infrastructure. Consequently, this data must be protected not only from other users of the computing infrastructure, but also from the cloud service provider who owns and manages the infrastructure.

Listed below are a number of potential risks to data in these multi-tenant cloud environments:

- Risk of Data Leakage – Individuals, or even business units, within some organizations will move sensitive data into the cloud without proper authorization. Once in the cloud the data is at risk, especially if additional effort is not made to safeguard the data;

- Unauthorized Access to Data – Other tenants of the computing infrastructure can potentially gain access to an organization's data through shared resources, inadequate logical separation, or improper configuration. Alternatively, service provider staff can potentially gain privileged access to organizational data;

- Misuse by Other Organizations – Other organizations sharing the computing resources may consume excessive resources, or inadvertently compromise the security posture of the shared resources through improper configuration or hardening, resulting in loss of data;

- Malicious Activities Targeting the Service Provider – Hackers may purposely target the service provider resulting in disruptions in service and even potential data loss. Malicious code can infect the cloud environment with similar results;

- Improper Data Sanitization – Data sanitization of backup media is complicated in a commingled, multi-tenant environment. Physical destruction of the backup media is not an option as the media likely contains the backup data of other organizations;

- Service Provider Viability – Service providers, as with any business, can go out of business. Depending on the speed with which they are forced to shut down their operations, organization data may be at risk. For example, in February 2009 cloud service provider Coghead suddenly shut down. Customers had nine weeks to remove their data from its servers or lose the data altogether; and

- Unforeseen Events – Unforeseen events can affect services for extended periods of time. A good example of an unforeseen event happened in Texas in April 2009. As part of a fraud investigation the Federal Bureau of Investigation (FBI) raided computing centers and seized hundreds of servers. Unfortunately, many organizations unrelated to the investigation had the misfortune of having their services collocated on these servers.[11]

[11] This event is described in detail at http://www.wired.com/threatlevel/2009/04/data-centers-ra/. A similar event happened more recently in June 2011 in Virginia. This event is described in detail at http://www.informationweek.com/news/security/management/231000897

2.4 Security Objectives

While the security objectives in cloud computing are identical to those in traditional IT environments, the threat environment and technologies employed are quite different as a result of the cloud service models employed. The primary concern is to ensure data protection in light of the risk posed by the various tenants and staff in the cloud computing environment. Specifically, this section will examine the following security objectives:

- Confidentiality – Maintaining the confidentiality of data in a multi-tenant cloud environment takes on an added importance. Not only is the data potentially collocated or commingled with data from other cloud consumers, but some cloud service provider staff members have privileged access to cloud computing resources hosting organizational data. Care needs to be taken to ensure that data confidentiality can be maintained in such an environment;

- Integrity – Maintaining the integrity of data in a multi-tenant environment is equally important. Care needs to be taken to ensure that data integrity[12] can be maintained in such an environment. Otherwise, the organization may lose confidence in data hosted in the cloud; and

- Availability – Cloud service providers usually offer up to 99.9% availability for all of an organization's computing resources. Availability can be affected temporarily or permanently, and a loss can be partial or complete. Denial-of-service attacks, equipment outages, and natural disasters are all threats to availability.

All three of these security objectives (confidentiality, integrity, and availability) should be encapsulated in a service provider's Service Level Agreement (SLA) with its customers. The SLA clearly denotes the expected level of security to be provided and, in the event that the provider fails to provide the security as specified, the compensation due to the cloud consumer. However, service and security levels can be difficult to measure and enforce. The CSA and the Federal Risk and Authorization Management Program (FedRAMP), discussed in notes below, can aid in the encapsulation of security objectives in SLAs. Furthermore, the Trusted Computing Group (TCG), also discussed below, formed a Trusted Multi-tenant Infrastructure (TMI) Work Group to develop a standards framework for realizing security objectives in the cloud.

Note - Cloud Security Alliance (CSA)[13]

The CSA is a not-for-profit organization with a mission to promote the use of best practices for providing security assurance within cloud computing, and to provide education on the uses of cloud computing to help secure all other forms of computing. The CSA has done a considerable amount of research in the area of cloud computing security. For example, it has produced the *Security Guidance for Critical Areas of Focus in Cloud Computing* **[Reference 2]**. This document is a set of best security practices that the CSA has put together for 14 domains involved in governing or operating the cloud.

[12] Readers are reminded that data is a generic term used to denote processes as well. The integrity of processes should also be maintained.
[13] Additional information on the CSA can be found at https://cloudsecurityalliance.org/

Note – Federal Risk and Authorization Management Program (FedRAMP)[14]

The FedRAMP is a program that was developed in conjunction with local governments, academia and private industry to facilitate the process of assessing and acquiring cloud services by federal agencies. Specifically, it establishes a standard approach for security assessments, authorization and continuous monitoring for cloud products and services.

Note – TCG TMI Work Group[15]

In 2010 the TCG formed the Trusted Multi-tenant Infrastructure (TMI) Work Group with the purpose of developing a standards framework for implementing:

- Shared infrastructures;
- Multi-provider infrastructures;
- Reference models and implementation guidance; and to
- Identify and address gaps in existing standards.

Specifically, the work group intends to develop an open framework, leveraging existing TCG specifications, for the practical deployment of trusted cloud or shared infrastructures. Existing TCG specifications that could be leveraged include the Trusted Network Connect (TNC) architecture and the Trusted Platform Module (TPM), as well as work being conducted within the TCG Infrastructure Work Group and the Virtualization Work Group. Readers interested in learning more about this initiative are encouraged to read *Cloud Computing and Security – A Natural Match* **[Reference 4]**.

[14] Additional information on FedRAMP can be found at http://www.gsa.gov/portal/category/102371
[15] Additional information on the TCG TMI Work Group can be found at http://www.trustedcomputinggroup.org/solutions/cloud_security

3 Data Protection

3.1 Overview

Data protection in a multi-tenant cloud environment is a challenge for a number of reasons. Not only is the data no longer located inside the corporate perimeter, but it is situated in an environment that is significantly different from traditional computing infrastructures. This environment employs new physical and logical architectures, is elastic and, perhaps most importantly, is multi-tenant. Consequently, a new data protection strategy must be adopted for this environment. This new approach, which is illustrated in Figure 2 and Figure 3, must protect sensitive data at every step. Specifically, it must address data protection at the following stages:

- Data Transfer;
- Data in Transit;
- Data at Rest;
- Data in Use; and
- Data Sanitization & Remanence.

It should be noted that the requirements for data protection will vary on the service model adopted, the deployment model selected, as well as an organization's tolerance for risk. Furthermore, the implementation of these data protection strategies may affect the usability of the data. During the discussion of data protection strategies, this section will highlight many of these differences and trade-offs.

Note – Security as a Service (SecaaS)

SecaaS is basically a cloud-based security service that can be leveraged by cloud consumers to secure their computing infrastructure, including data, in the cloud. While SecaaS offers some of the same advantages to large organizations as cloud computing itself, namely agility and cost savings, it is perhaps most beneficial to small and medium organizations. These smaller organizations rarely have a full security department and consequently do not typically have the necessary in-house expertise or the technology capable of providing the full suite of security services. SecaaS allows these organizations to implement security uniformly across their entire computing infrastructure. Rather than having to implement a complete security infrastructure themselves, they can merely pay for what they require.

Note – Defence-in-Depth

Defence-in-depth, sometimes call layered security, is an information security principle whereby multiple safeguards are employed to protect an asset. In theory, if an attacker were able to bypass one of the safeguards, the asset would still be protected by others. This principle is applicable to data protection in multi-tenant cloud environments. Specifically, data should be protected using multiple safeguards in order to ensure that a weakness in one safeguard does not compromise the organization's data protection strategy.

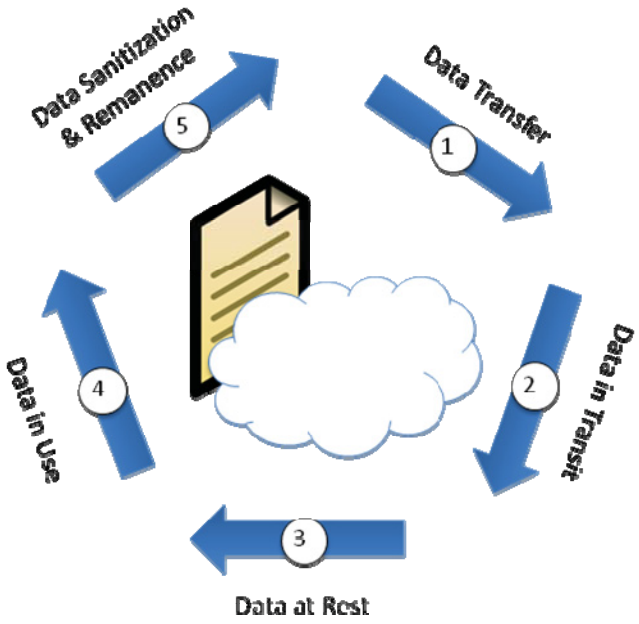

Figure 2 - Data Protection Life Cycle in the Cloud

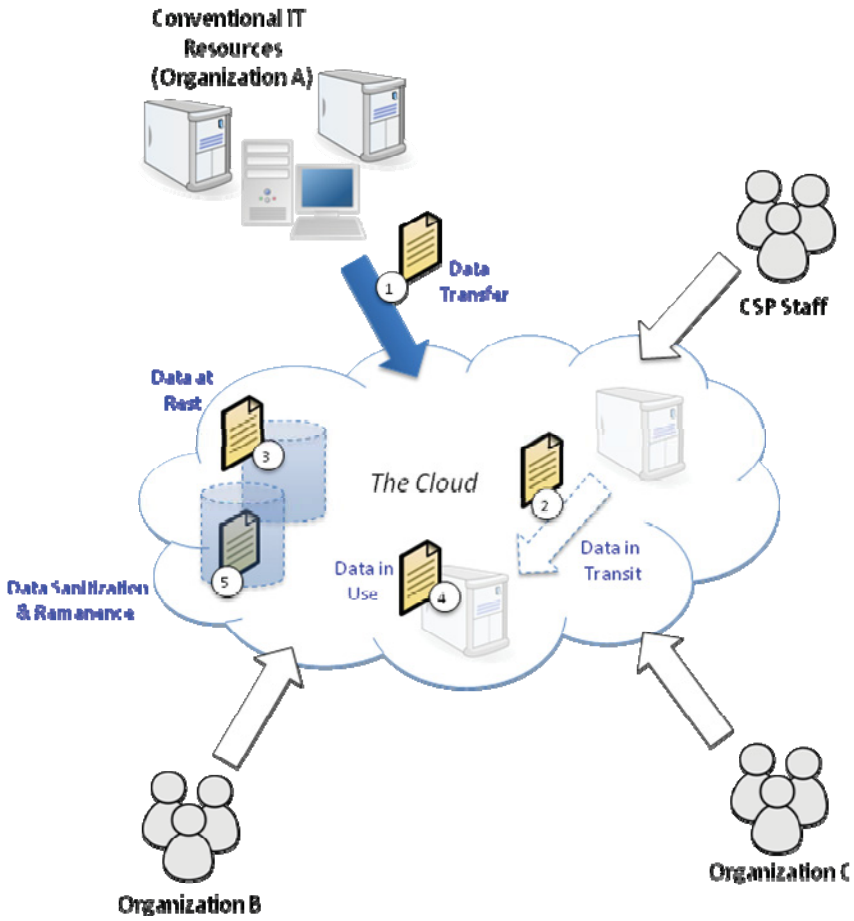

Figure 3 - Data Protection in a Multi-Tenant Cloud

3.2 Data Transfer

Data transfer refers to the transfer of data from the organization's private network to the cloud environment. There are two aspects to consider with respect to data transfer: authorized data transfer and unauthorized data transfer.

3.2.1 Authorized Data

The most effective strategy for data protection in public cloud computing environments is to ensure that no sensitive or regulated data makes it into this environment in the first place. Some data is so sensitive, and the ramifications of disclosure so severe, that organizations should not even consider moving this information to the cloud. This likely applies to national security information, sensitive intellectual property and even some information subject to regulatory controls. While this approach will significantly simplify data protection, it may not be possible, or even warranted (see note below), for private or community clouds. However, organizations should make every effort to control what data is authorized for transfer to the cloud. Other data

protection strategies that can be adopted for data being transferred to the cloud are data anonymization and tokenization.

Note – Amazon Web Services (AWS) GovCloud

AWS GovCloud is a cloud offering targeted at U.S. government agencies and contractors with specific regulatory and compliance requirements. AWS GovCloud has been specifically designed to support data subject to compliance regulations such as the International Traffic in Arms Regulations (ITAR). It supports sensitive data by limiting the cloud to U.S. citizens and employing additional security safeguards.

3.2.1.1 Data Anonymization

Data anonymization is the process whereby information that links data to a specific identity is removed. Personally Identifiable Information (PII) and Sensitive Personal Information (SPI) are candidates for data anonymization. The theory is that by anonymizing data it is then releasable to the cloud. However, if not done properly inferences can be made about individuals, and in some cases, the personal identity can be determined. A good example of this is the release of "anonymized" hospital data by the Massachusetts Group Insurance Commission in the mid-1990s. This information was released, with all names, addresses, and social security numbers removed, in order to help researchers. However, a graduate student in computer science (Latanya Sweeney) was able to extract the Massachusetts Governor's health records, including diagnoses and prescriptions, from the "anonymized" data. In 2000 she demonstrated that 87% of Americans could be uniquely identified using only their ZIP code, birth date, and gender.[16]

While an exhaustive examination of data anonymization research is outside of the scope of this paper, some data anonymization techniques, and corresponding research papers, are listed below:

- *k*-Anonymity – This anonymization technique involves altering the dataset so that it maps to a number of individuals, thereby making any link ambiguous. The greater the number of individuals that it maps to, the more ambiguous the link, and as a result, the more anonymous the data. A dataset is considered to provide *k*-anonymity protection if the information for each individual contained in the dataset cannot be distinguished from at least *k*-1 individuals whose information also appears in the dataset. *k*-anonymity is discussed in further detail in *k-Anonymity: A Model for Protecting Privacy* **[Reference 5]** and *k-Anonymity* **[Reference 6]**;

- Generalization & Suppression – These are anonymization techniques that can be used to provide k-anonymity. Generalization involves replacing a value with a less specific but semantically consistent value (e.g., replacing specific ZIP codes with generalized ZIP codes indicative of a larger geographical area). Suppression involves not releasing a value at all. There are a number of advantages to these techniques. First, recipients of the data can be told what was done to the data. Second, it guarantees anonymity. Third, it minimally distorts the data. Generalization and suppression are discussed in further detail in *Achieving k-Anonymity Privacy Protection Using Generalization and Suppression* **[Reference 7]**; and

[16] http://arstechnica.com/tech-policy/news/2009/09/your-secrets-live-online-in-databases-of-ruin.ars

- Greedy Partitioning Algorithm – The greedy partitioning algorithm employs multidimensional partitioning to achieve *k*-anonymization. This algorithm is discussed in further detail in *Achieving Multidimensional k-Anonymity by a Greedy Approach* **[Reference 8]**.

3.2.1.2 Tokenization

Tokenization is a data protection technique that replaces sensitive information in data with non-sensitive tokens. The tokens are in fact a reference to the encrypted text which is stored in a central repository, or data vault. There is no mathematical relationship between a token and the data value. The relationship is purely referential. Tokens can preserve the original format of the data in terms of length and data type. They can also be used to preserve a number of leading and trailing characters (e.g., a credit card). There are two models of tokenization solutions. They are as follows:

- In-house – An in-house tokenization model is typically adopted by medium to large organizations, that have significant annual revenues and/or sensitive data to secure, and wish to maintain complete control over their data; and

- Outsourced – An outsourced tokenization model is typically adopted by small to medium organizations, that have moderate annual revenues and/or sensitive data to secure, and do not require complete control over their data.

In terms of usability, tokens enable applications such as marketing analysis, order confirmation, loss prevention and fraud detection to work without risk of compromise of the data. Furthermore, since tokens can preserve the length and format of the original data, they are transparent to databases and applications. Tokenization is usually provided in conjunction with encryption. Encryption is discussed in Section 3.4.3.

Note – Tokenization Standards

The Payment Card Industry (PCI) Security Standards Council (SSC) Scoping Special Interest Group (SIG) is working on guidelines for tokenization as they relate to PCI Data Security Standard (DSS). In PCI terminology tokenization refers to the process by which the Primary Account Number (PAN) is replaced with a non-sensitive value called a "token". The use of tokenization in the PCI is discussed in more details in *PCI DSS Tokenization Guidelines* **[Reference 9]**.

The Accredited Standards Committee X9 is working on a standard to define tokenization requirements related to credit card data in the financial services industry. The standard has the working title of X9.119, Retail Financial Services – Requirements for Protection of Sensitive Payment Data – Part 1: Using Encryption/Tokenization Methods.

3.2.2 Unauthorized Data

Individuals, and even business units, may unilaterally decide to move data to the cloud without the organization's knowledge or approval. Once in the cloud the data will be at an increased risk of compromise due to the environment and the lack of additional data protection. Consequently, organizations should make every attempt to prevent unauthorized data from leaking to the cloud. There are a few technologies that can be used by organizations to help prevent data leakage to the cloud. These technologies, which are the same ones used to prevent data leakage to public networks (such as the Internet) in many organizations, include the following:

- Security Labelling;
- Rights Management;
- Data Loss Prevention (DLP); and
- Integrated Solution.

Security labelling and rights management are also applicable to data at rest (Section 3.4). All of the technologies, including the integrated solution, can also be used to prevent cloud data from leaking into the public domain.

3.2.2.1 Security Labelling

One of the challenges in preventing sensitive data from leaking into the cloud is the difficulty in distinguishing sensitive data from non-sensitive data. Affixing a security label denoting the sensitivity of the data will facilitate this identification process. Ideally, the security label should be applied by the owner of the data and it should be cryptographically bound to the data through the use of a digital signature. The digital signature also ensures the integrity of the data.

We can safely assume that some organizations will store data in the cloud without full awareness as to the sensitivity of this data. The challenge then becomes how to discover and label sensitive data in the cloud. The discovery of sensitive data within the cloud can be accomplished by scanning cloud data stores using DLP solutions (Section 3.2.2.3). Once potentially sensitive data has been identified, the appropriate security label will need to be determined and applied. While this can be a manual process, the fact that this is extremely time consuming will limit the amount of data that can be supported. A better approach would be an automatic classification system capable of performing content analysis on the data in order to determine, and apply, the appropriate security label. Unfortunately, these solutions are currently in the research phase.

3.2.2.2 Rights Management

Rights management is a persistent form of security that travels with the data in much the same way as encryption. However, it takes encryption one step further. Rather than simply providing access to authorized users and preventing access by unauthorized users, rights management dictates, through a usage policy, what the user is capable of doing with the data. In terms of data leakage to cloud environments, rights management can help mitigate this risk. Furthermore, in the event that sensitive data does leak to the cloud, rights management will ensure that the data

continues to be protected. In the event of leakage of sensitive data, organizations can even prevent further access to the data by changing the usage policy.

3.2.2.3 Data Loss Prevention (DLP)

DLP, sometimes called extrusion prevention, is a technology targeted at preventing sensitive information from leaving an organization in an unauthorized manner. It accomplishes this in one of two ways: by discovering and protecting sensitive information located on servers and desktops within the organization, or by monitoring outbound network communications for sensitive information.

Although DLP predates cloud computing, it is well suited for preventing data from leaking into the cloud. While Uniform Resource Locator (URL) filtering can be used to monitor, and even prevent, users connecting to cloud services, DLP provides more granular controls. DLP can actually examine the data being transferred, rather than just the destination.

Once sensitive data is in the cloud, DLP is still of use. DLP can be used to identify sensitive data stored in the cloud. It does this by scanning systems, including storage systems, for sensitive data. Sensitive data identified in this manner can be labelled, fingerprinted[17], or included in a report. DLP can also be used to prevent data leakage from the cloud. This could include embedding DLP agents into machine images as part of an IaaS offering.

Note - DLP Limitations

There is this story about an old man who saw a boy looking down at the ground while circling under a lone street light on a dark street. The old man asked the child "what happened?" and the child replied "I lost a coin and I am looking for it". The old man joined the child in looking for the coin and after quite a while of not finding a coin asked the child "where have you lost the coin?" to which the child answered "over there" pointing into the dark. "So why are you looking for the coin here?" asked the old man, and the child answered "because it's dark over there." [18]

DLP is effective at scanning unencrypted network traffic for data leakage, but it is completely ineffective when scanning encrypted traffic. In other words, to draw a parallel with the story above, it is effective at looking where the light is and less effective at looking in the dark where the data is likely to be. So, unless the network-based DLP can decrypt the network traffic, it is going to be unable to scan it for concealed content. It is worth mentioning that a web gateway can be located alongside the network-based DLP. It could basically terminate and decrypt encrypted communications and pass the decrypted content to the network-based DLP for scanning.

[17] Fingerprinting sensitive data consists of calculating a cryptographic hash of the data. This cryptographic hash, which uniquely identifies the data, can be used to identify the sensitive data in the event that someone attempts to leak it out of the organization.
[18] http://securitynirvanablog.wordpress.com/tag/ssl/

3.2.2.4 Integrated Solution

The technologies discussed in this section of the technical brief should not be used in isolation but rather as part of an integrated solution. Specifically, the following areas of integration can increase the effectiveness of the individual technologies in preventing data leakage to the cloud:

- Security Labelling and Rights Management – The security labelling and rights management components should be integrated so that when a user affixes a security label to a piece of data, the corresponding rights management policy is automatically applied to the data;

- Security Labelling and DLP – The security labelling and DLP components should be integrated so that the DLP solution is preconfigured to recognize the security labels and, based on those labels, perform certain actions. For example, the DLP solution could be configured to prevent data with a specific security label from going to the cloud; and

- Rights Management and DLP – The rights management and DLP components should be integrated so that the DLP solution is preconfigured to enforce the rights management policy.

3.3 Data in Transit

Data in transit covers data travelling to the cloud, data moving within the cloud, and data travelling from the cloud. This data is usually secured using a secure network security protocol. These protocols can be applied at various levels of the Open Systems Interconnection (OSI) stack, including the application layer (e.g., Secure Shell (SSH)), the transport layer (e.g., Transport Layer Security (TLS)) and the network layer (e.g., Internet Protocol Security (e.g., IPsec)). These protocols can be implemented for SaaS, PaaS and IaaS environments. Organizations need to ensure that the cryptographic algorithm selected has been appropriately vetted and that the cryptographic keys are of sufficient strength.

The challenge in solely protecting data while it is in transit is that it makes the assumption that the endpoints are secure. In the case of cloud computing, this assumption may not be true due to the fact that the endpoints are shared and that service provider staff likely have privileged access to the endpoints. Data can also be protected in transit using persistent security mechanisms. These persistent security mechanisms are discussed in some detail in Section 3.4.

3.4 Data at Rest

There are a number of mechanisms for protecting data at rest within the cloud. Not all mechanisms are standardized and not all are feasible for all cloud environments. Data at rest mechanisms include the following:

- Separation/Isolation;
- Access Management;
- Encryption;

- Integrity Protection;
- Data Dispersion; and
- Monitoring & Audit.

It is worth mentioning that security labelling and rights management can also be used to protect data at rest. These capabilities were discussed in Sections 3.2.2.1 and 3.2.2.2, respectively.

3.4.1 Separation/Isolation

Perhaps the most obvious data protection mechanism for multi-tenant cloud environments is to isolate the data from other consumers. This is somewhat challenging due to the fact that the data tends to be collocated or commingled in this type of environment. The method of accomplishing separation/isolation is dependent on the service model. This section will examine separation/isolation for each of the three service models.

3.4.1.1 Infrastructure as a Service (IaaS)

Network zoning should be implemented by the service provider to ensure that sensitive back-end services are appropriately separated from the publically available front-end services. Networks within large organizations are typically comprised of a Public Access Zone (PAZ), an Operations Zone (OZ) and a Restricted Zone (RZ). The Internet is considered a Public Zone (PZ). All zones are separated by security controls, including a firewall, in order to better protect the services located in a particular zone. Service providers should never collocate services with different security requirements on the same server. In other words, a physical server should not span security zones. The reason for this is that an improper configuration could enable communications to travel directly between security zones, effectively bypassing inter-zone security controls. Even within the same security zone, Virtual Local Area Networks (VLANs) and firewalls can be used to provide additional levels of separation.

In addition to understanding the network zoning implemented by the service provider, the consumer also needs to understand whether their VMs will be hosted on the same physical system as the VMs of other consumers and if so, what safeguards the service provider has implemented to isolate them from one another. Obviously, there are a number of concerns with co-located VMs. First and foremost, a VM from one consumer can be used to affect the operation, either intentionally or accidentally, of other VMs hosted on the same physical platform. Likewise, if another consumer fails to adequately secure their VM, it could easily be compromised and serve as a beachhead from which to launch attacks and potentially compromise other VMs.

Organizations will also need to understand how their data will be stored within the cloud, such as whether it will be located on dedicated or shared storage systems. The advantages to dedicated storage include greater flexibility and control, and a higher level of security. The advantages of shared storage, and the disadvantages of dedicated storage, all relate to cost.

3.4.1.2 Platform as a Service (PaaS)

In terms of PaaS separation/isolation there are four levels of multi-tenancy that provide varying levels of separation/isolation, mainly through the use of logical separation mechanisms. These four levels are as follows:

- Level 1 – Separate Machine: Each user is provided with a separate physical machine;
- Level 2 – Configurable Instance: Each user is provided with a configurable instance on the same machine;
- Level 3 – Single Instance: A single instance is shared amongst multiple users; and
- Level 4 – Multiple Instances: Multiple instances are shared amongst thousands of users.

3.4.1.3 Software as a Service (SaaS)

Separation/isolation in SaaS environments is a challenge due to the fact that multiple tenants are sourcing the same application, and likely storing their data in the same database. There are three approaches offering varying degrees of separation/isolation in SaaS environments. These approaches, which are listed in descending order of separation/isolation, include the following:

- Separate Databases – The use of separate databases to store consumer data provides the highest level of separation/isolation in SaaS environments. Unfortunately, this approach is also the most expensive. Consequently, this approach negates many of the advantages of cloud computing;
- Shared Database with Separate Schemas – This approach supports multiple tenants in the same database, each with their own set of tables that are grouped in a schema created specifically for the consumer. The use of a shared database with separate schemas balances separation/isolation and cost. It provides a moderate degree of logical separation/isolation while allowing service providers to support a relatively large number of tenants on each database server; and
- Shared Database with Shared Schema – This approach uses the same database and the same set of tables to store data from multiple tenants. The use of a shared database with shared schema tilts solidly in favour of cost at the expense of separation/isolation.

The use of shared databases also necessitates the use of access controls to ensure separation/isolation. Specifically, access controls are leveraged within the database to ensure that access, including modifications, to rows and fields are controlled according to policy. Access management is discussed in Section 3.4.2.

3.4.2 Access Management

Access management, which includes authentication and authorization, can be used to provide data protection in cloud computing environments. Authentication refers to the establishment of a valid identity. Authorization refers to the granting of access to resources according to policy.

The problem with access management in cloud computing environments is twofold. First and foremost, service providers tend to provide limited access management with weak authentication (e.g., username and password) and coarse authorization (e.g., administrator and user). Second, it is a challenge for most organizations to extend their access management capabilities to the cloud environment. Organizations that do not extend their access management capabilities to the cloud will run into the federation problem, whereby the domain hosting the services, in this case the cloud, does not recognize the users attempting to access the services. Consequently, the cloud is unable to authenticate the users and determine what resources they are authorized to access. Some organizations configure a Virtual Private Network (VPN) tunnel from the cloud to their enterprise network in order to allow users to authenticate using the enterprise Lightweight Directory Access Protocol (LDAP) directory. However, this approach has issues in terms of connectivity, latency, disaster recovery, etc.

A better approach to providing access management in cloud computing environments is identity federation. Federation provides a standards-based means with which identity information can be shared across domains and between organizations in order to facilitate Single Sign-On (SSO). It is a critical component of cloud computing as it allows an organization to extend their authentication infrastructure into the cloud using open standards (e.g., Security Assertions Markup Language (SAML), Web Services (WS)-Federation) while enabling SSO for their users.

3.4.3 Encryption

Due to the multi-tenant nature of cloud environments it makes sense to employ cryptography, and specifically encryption, to secure sensitive data. In some cases the consumer will rely on the service provider to encrypt consumer data and manage the cryptographic keys. However, consumers will need to ensure that they are comfortable with service provider staff having potential access to data. Consumers also need to ensure that key management is being performed properly. Many service providers only provide basic encryption key schemes. In fact, some service providers use a single key to encrypt all of a customer's data.

In most cases, the consumer will be responsible for key management. This is referred to as tenant-owned key management. However, there are challenges to implementing encryption in cloud environments, including the following:

- Secure Key Stores – Key stores, as with other sensitive data, need to be protected at all times as the compromise of the key will result in unauthorized access to all data encrypted with that key;

- Access to Key Stores – Access to key stores should be carefully controlled. Specifically, it should be restricted to trusted users and implemented in a manner consistent with the principles of role separation and least privilege. For example, the role that is capable of

using a key should not be assigned to the same person who is assigned the role tasked with storing the key; and

- Key Backup and Recovery – Secure key backup and recovery must be implemented to prevent the loss of a key from resulting in the permanent loss of data. Key backup and recovery allow for the controlled decryption of data in the event that the keys used to encrypt the data are lost.

In terms of usability, encryption has a number of drawbacks including the consumption of additional computing resources (memory and processor utilization), a larger storage footprint, and key management issues. Furthermore, not all applications and databases are capable of storing and processing encrypted data. Perhaps most importantly from a usability perspective, encrypting data at rest will prevent, or at least inhibit, indexing or searching of that data. This is discussed in more detail in Section 3.5.

Note – Organization for the Advancement of Structured Information Standards (OASIS) Key Management Interoperability Protocol (KMIP)[19]

The purpose behind OASIS KMIP is *to define a single, comprehensive protocol for communication between encryption systems and a broad range of new and legacy enterprise applications, including email, databases, and storage devices.* In other words, KMIP provides a means of communication with which any encryption system can communicate with any key management system. This will allow organizations to deploy a single enterprise key management infrastructure with which to manage keys for all of their encryption systems, including those in the cloud environment. Alternatively, a service provider could utilize this standard to leverage a single key management system for all consumers in the cloud, regardless of their specific encryption systems.

3.4.3.1 IaaS

IaaS typically uses encryption to protect volumes or VM images. Volume storage encryption protects volumes from snapshot cloning, from being searched by the service provider, and from being compromised in the event of the physical loss of drives. Similarly, VM images are susceptible to theft or modification when they are dormant and when they are running. VM images can be encrypted at all times in order to mitigate this threat; however, this approach detrimentally affects performance.

3.4.3.2 PaaS

Encrypting data at rest within PaaS environments is generally more complex, involving any of the following options:

[19] Additional information on OASIS KMIP can be found at http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=kmip

- Client/Application Encryption – Data is encrypted in the PaaS application or the client accessing the platform;

- Database Encryption – Data is encrypted in the database using encryption built in and supported by the database platform;

- Proxy Encryption – Data must pass through an encryption proxy before being sent to the platform; and

- Other – Additional options may include APIs built into the platform, external encryption services, and other variations.

3.4.3.3 Software as a Service (SaaS)

Encryption of data at rest within SaaS environments basically involves encrypting all of the tenant's data within the database so that the stored data cannot be compromised in the event that the database is stolen. Given that this type of encryption is usually built into the application (the database), it is usually quite difficult for consumers to implement encryption of data at rest in a SaaS environment. Consequently, they will need to request this capability from their service provider. However, in order to better enforce multi-tenancy isolation, consumers should insist that per-customer keys be used.

3.4.4 Integrity Protection

For data in transit most security protocols automatically provide integrity protection. However, for data at rest integrity protection is typically provided independently of data encryption. There are a number of ways to provide integrity protection for data at rest. These include cryptographic hashes, also known as cryptographic checksums, as well as digital signatures. A cryptographic hash is a one-way function that generates a fixed-size datum for a given piece of data. The integrity of the data can be verified by recomputing the hash value and comparing it to the original value. If the hash values are the same then the data has not been modified. A cryptographic hash is considered one-way in that it should be computationally impossible to determine the original data from the hash value. Cryptographic hashes should also be collision resistant. In other words, it should be computationally impossible to find two pieces of data that have the same hash value. A digital signature uses a private signing key to encrypt the hash value for a given piece of data. Consequently, anyone using the public verification key to decrypt the hash value can be confident that the data was signed by a particular user, and, assuming the hash values match, that the data has not been modified since it was signed.

3.4.5 Data Dispersion

Data dispersion is a data protection technique that does not rely on encryption. Instead, it breaks data into fragments and stores them on distributed servers. For example, a file is divided into n fragments. Each of these n fragments is signed and then distributed to n remote servers. An authorized user can then reconstruct the file f by retrieving m arbitrarily chosen fragments. Data dispersion can be used in conjunction with encryption to further enhance data protection.

Information Dispersal Algorithms (IDAs) were first proposed by Michael Rabin in 1989.[20] These algorithms basically parse data into matrices using matrix multiplication. IDAs are applicable to both data at rest (storage arrays) as well as data in motion. In storage arrays, the data is parsed and stored in different arrays in order to mitigate the possibility of compromise. For data in motion, IDAs are used to parse data prior to transmission and reassemble data at the receiving device.

3.4.6 Monitoring & Audit

The protection of data at rest can also be enhanced through monitoring and audit. There are two types of monitoring directly relevant to cloud computing. They are as follows:

- Database Activity Monitoring (DAM) – DAM is real time, or near real time, monitoring of database activity. Specifically, it captures and records all Structured Query Language (SQL) activity, including privileged user actions, across a variety of database platforms. Policy violations detected, such as SQL injection attacks or database replication, can result in automatic alerts. DAM typically necessitates the use of agents located on database servers. All database activity is analyzed on a central collection server; and

- File Activity Monitoring (FAM) – FAM is real time, or near real time, monitoring of file activity. Specifically, it monitors and records all file activity across file servers located in the cloud. Policy violations, such as unauthorized file access, can result in automatic alerts. FAM typically necessitates the use of agents located on file servers or the placement of a physical appliance between the cloud storage and the cloud consumers.

The organization will need to make a decision on whether the monitoring and audit function should be maintained in-house, outsourced to the cloud service provider, or outsourced to a third-party. The primary concern with outsourcing this function to the cloud service provider is that in all likelihood the provider's staff will have privileged access to the organization's resources located in the cloud. Consequently, in keeping with the principle of separation of duties, the monitoring and audit function should either be kept in-house or outsourced to a trusted third-party.

3.5 Data in Use

It is all well and good to protect data in transit and at rest, but unless the data is also protected while it is in use it will be susceptible to attack. For example, encrypted data is protected in transit and at rest. However, it is typically decrypted during processing, rendering the data vulnerable to attack. Nearly all data use, other than storage, necessitates that the data be decrypted to be used. Consequently, the challenge is in protecting data during processing while maximizing its utility. Up until quite recently an organization could have one or the other (either data protection or usability). This limitation on protecting data in use is starting to change with the advent of homomorphic encryption and predicate encryption.

[20] *Efficient Dispersal of Information for Security, Load Balancing, and Fault Tolerance* **[Reference 10]**

3.5.1 Homomorphic Encryption

In 2009 IBM researchers, in conjunction with a graduate student (Craig Gentry) from Stanford University, developed a homomorphic encryption scheme that enables the processing of, and computations on, encrypted data. This research is documented in *Fully Homomorphic Encryption Using Ideal Lattices* **[Reference 11]** and *A Fully Homomorphic Encryption Scheme* **[Reference 12]**. At Eurocrypt 2010 a working implementation of full homomorphic encryption was demonstrated. While additional effort is needed to make the research practical, it is anticipated that this is just a matter of time. Specifically, additional research needs to be undertaken in order to reduce the significant computational effort required for homomorphic encryption.

3.5.2 Predicate Encryption

Other cryptographic research efforts are underway to limit the amount of data that would need to be decrypted for processing in the cloud. Predicate encryption provides the key owner with greater control over access to encrypted data. It does this through the creation of subkeys that can be used to decrypt a subset of the data. For example, a key owner could create subkeys that would allow a mail server to partially decrypt messages in order to appropriately route them. Relevant predicate encryption research includes Predicate Encryption Supporting Disjunctions, Polynomial Equation, and Inner Products **[Reference 13]**, Predicate Privacy in Encryption Schemes **[Reference 14]**, and Functional Encryption for Inner Product Predicates from Learning with Errors **[Reference 15]**.

3.6 Data Sanitization & Remanence

At some point in time, sensitive data stored in the cloud will need to be deleted. This may occur due to the fact that the organization has terminated their relationship with the service provider, the organization has changed their business model and no longer requires cloud services, or the data is no longer required. Regardless of the reason, the process whereby data is permanently deleted from storage media is referred to as sanitization. Sanitization is typically accomplished by overwriting the data, degaussing the storage media, physically destroying the storage media, or by other means.

Failure to completely erase the data will result in data remanence. Data remanence is the residual representation of data that has been in some way nominally erased or removed. Data remanence can result in the unauthorized recovery, and disclosure, of sensitive data. In cloud environments data remanence is applicable to all three cloud services (SaaS, PaaS, and IaaS). Many service providers make no mention of data remanence. When pressed some service providers cite compliance with U.S. Department of Defense (DoD) 5220.22-M (the National Industrial Security Program Operating Manual (NISPOM)) **[Reference 16]**. Unfortunately, while this standard does state two approved methods for data destruction, it does not provide any specifics on how these methods are to be accomplished.

Readers interested in data sanitization and remanence are encouraged to consult the following references:

- National Security Agency (NSA)/Central Security Service (CSS) Storage Device Classification Manual **[Reference 17]**;
- Secure Deletion of Data from Magnetic and Solid-State Memory **[Reference 18]**;
- RCMP IT Media Overwrite and Secure Erase Products **[Reference 19]**;
- Data Remanence in Flash Memory Devices **[Reference 20]**;
- Reliably Erasing Data from Flash-based Solid State Drives (SSDs) **[Reference 21]**; and
- Data Remanence: Secure Deletion of Data in SSDs **[Reference 22]**.

4 Further Research: Extending Data Protection to the Cloud

4.1 Overview

This report makes the recommendation that further research should be conducted into extending data protection to the cloud. In order to better understand the issues involved we will contrast traditional data access with cloud data access, and identify some of the gaps requiring further research.

4.2 Traditional Data Access

Traditional data access, which is illustrated in Figure 4, consists of the following steps:

1) Authentication – The identity of the user is authenticated. This is usually accomplished using a username and password (e.g., Windows domain);

2) Authorization – Based on the authenticated identity of the user, the sensitivity of the resource, and the security policy, the system makes a determination as to whether the user is entitled to access the resource or not; and

3) Data Access – Assuming that both the authentication and authorization processes were successful, the user is provided with access to the resource.

There is usually not a requirement to encrypt sensitive data within most organizations. Not only is the data protected using access control, but the data remains in the confines of the organization's private network. This private network is typically managed by trusted employees.

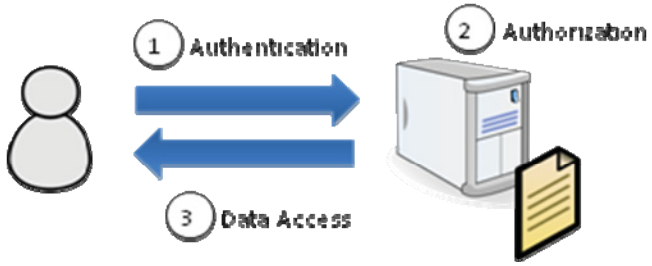

Figure 4 - Traditional Data Access

4.3 Cloud Data Access

The challenge facing many organizations is to take a relatively simple data access scenario and extend it to the multi-tenant cloud environment. This cloud data access, which is illustrated in Figure 5, consists of the following steps:

1) Authentication – In a cloud data access scenario the user typically establishes his identity in much the same way as in a traditional data access scenario. However, there are two key differentiators. First and foremost, the user authenticates to an Identity Provider (IdP) within the private network. Second, upon successful authentication the user is issued with an identity token (e.g., SAML assertion). It is worth mentioning that authentication within the cloud may necessitate the use of stronger authentication, including multi-factor authentication (e.g., smart cards, biometrics). Fortunately, federation standards support multiple Levels of Assurance (LoA). LoA refers to the level of confidence in the authenticity and veracity of the authentication credentials being used. LoA is typically used to assess the thoroughness of the identity-proofing process as well as the strength of the authentication process;

2) Identity Federation – Identity federation provides a standards-based means with which identity information can be shared across domains and between organizations. In the cloud data access scenario, identity federation enables the organization to extend its authentication infrastructure into the cloud while enabling SSO for its users. Specifically, the user presents his identity token to the Service Provider (SP). The SP relies on the IdP to authenticate the user's identity rather than having to do it itself;

3) Authorization – As with traditional data access, authorization in the cloud leverages the authenticated identity of the user, the sensitivity of the resource, and the security policy to determine whether the user is entitled to access the resource. However, rather than providing authorization on an application-by-application basis within the cloud, it typically leverages the organization's authorization infrastructure. Specifically, the SP also acts as a Policy Enforcement Point (PEP), preventing access to the resource until an access determination has been made. This access determination is made by the Policy Decision Point (PDP), which is administered by the Policy Administration Point (PAP). Both of these components would be located in the private network due to their sensitivity. Furthermore, the authorization infrastructure should leverage a flexible policy language (e.g., eXtensible Access Control Markup Language (XACML)) capable of supporting the fine-grained policies required in the cloud environment;

4) Decryption/Re-encryption – This technical brief outlined the importance of encrypting sensitive resources due to the presence of potential threat agents with privileged access in the multi-tenant cloud environment. The challenge then becomes how to decrypt/re-encrypt a resource for a particular user once the user has been granted access to it. And just as importantly, how to accomplish this in a manner that does not expose the sensitive data or cryptographic keys within the cloud environment. In terms of providing access to data without exposing it, or the cryptographic keys, both homomorphic and predicate encryption seem to have potential. Alternatively, decryption of the sensitive data, and processing of the unencrypted data, could occur within the confines of the private network; and

5) Data Access – Assuming that the previous steps were successful, the user is provided with access to the resource. Access to the unencrypted data should only occur within the private network in order to mitigate the threat of attack within the cloud computing environment.

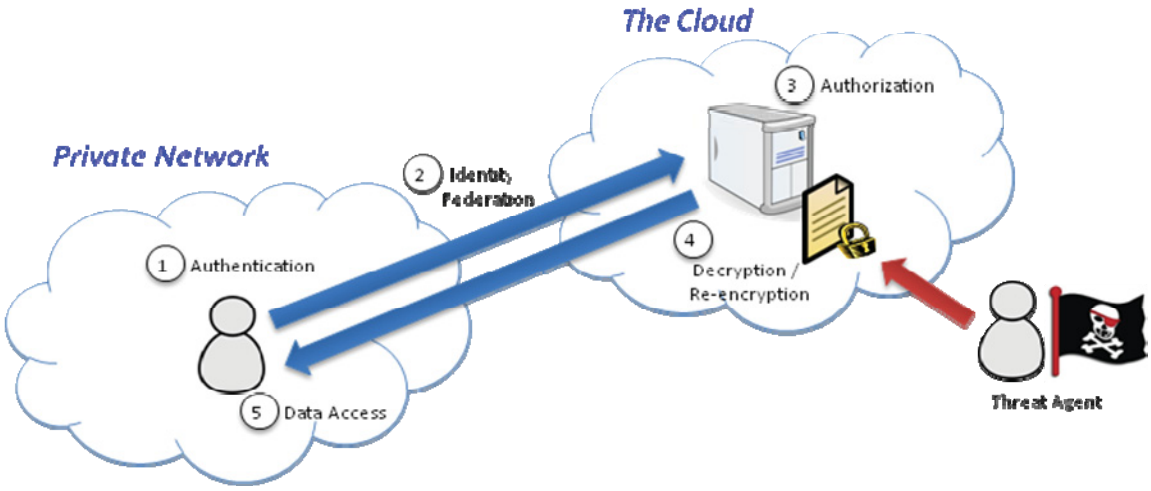

Figure 5 - Cloud Data Access

4.4 Research Gaps

From the two data access scenarios it is apparent that there are gaps when extending data protection to the cloud that require further research and prototyping. Particular emphasis should be placed on the following topics:

- Identity Federation – While federation is fairly straightforward, at least in theory, there are a number of aspects that require additional investigation, especially as they pertain to cloud computing. These include LoAs, IdP discovery protocols, and support for multiple federation protocols;

- Authorization – Most organizations do not have a centralized authorization infrastructure (PEP, PDP and PAP), let alone one that can be extended to the cloud. Issues that will need to be explored further include security policy definition and distribution within the cloud, potential latency issues due to the distributed nature of the authorization infrastructure, and applicability to the three cloud computing service models; and

- Encryption – Encryption within the cloud, and specifically for data in use, is the area that requires the most additional research. Homomorphic encryption and predicate encryption seem to have a great deal of potential and their use should be explored further. In addition, decryption of the sensitive data, and processing of the unencrypted data, within the confines of the private network is a potential solution for data in use. This approach, which complicates key management, is also a candidate for further research.

5 Conclusion

Data and processes are at risk in multi-tenant cloud computing environments due to potential threats from other cloud consumers, as well as from third-party cloud providers. In order to ensure data protection in this environment, the data must be protected at all stages of its life cycle, including data transfer, data in transit, data at rest, data in use, and data sanitization and remanence. This technical brief proposed a data protection strategy whereby sensitive data is protected using a layered approach consisting of multiple data protection techniques. However, rather than introducing entirely new technologies for protecting data in the cloud, most organizations would be better off extending their existing data protection capabilities into the cloud. Additional research in this area is required, specifically in the areas of identity federation, authorization and encryption.

References

[1] P. Mell and T. Grance, NIST Definition of Cloud Computing, NIST, September 2011.

[2] Cloud Security Alliance, Security Guidance for Critical Areas of Focus in Cloud Computing, Version 3.0, 2011.

[3] W. Jansen and T. Grance, NIST Guidelines on Security and Privacy in Public Cloud Computing, NIST, December 2011.

[4] Trusted Computing Group, Cloud Computing and Security – A Natural Match, April 2010.

[5] L. Sweeney, k-Anonymity: A Model for Protecting Privacy, International Journal on Uncertainty, Fuzziness and Knowledge-based Systems, Vol. 10, pp. 557-570, 2002.

[6] V. Ciriani, et al., k-Anonymity, Advances in Information Security, Springer US, 2007.

[7] L. Sweeney, Achieving k-Anonymity Privacy Protection Using Generalization and Suppression, Carnegie Mellon University, International Journal on Uncertainty, Fuzziness and Knowledge-based Systems, Vol. 10, pp. 571-588, 2002.

[8] G. Murthy and R. Srinivas, Achieving Multidimensional k-Anonymity by a Greedy Approach, International Conference on Web Services Computing (ICWSC), 2011.

[9] PCI Security Standards Council, PCI DSS Tokenization Guidelines, Version 2.0, August 2011.

[10] M. Rabin, Efficient Dispersal of Information for Security, Load Balancing, and Fault Tolerance, Journal of the ACM, Vol. 36, April 1989.

[11] C. Gentry, Fully Homomorphic Encryption Using Ideal Lattices, STOC ’09, June 2009.

[12] C. Gentry, A Fully Homomorphic Encryption Scheme, PhD Thesis, Stanford University, September 2009.

[13] J. Katz, A. Sahai and B. Waters, Predicate Encryption Supporting Disjunctions, Polynomial Equations, and Inner Products, EUROCRYPT ’08, 2008.

[14] E. Shen, E. Shi and B. Waters, Predicate Privacy in Encryption Systems, TCC '09 Proceedings of the 6th Theory of Cryptography Conference on Theory of Cryptography, 2009.

[15] S. Agrawal, D. Freeman and V. Vaikuntanathan, Functional Encryption for Inner Product Predicates from Learning with Errors, ASIACRYPT, 2011.

[16] U.S. Department of Defense, DoD 5220.22-M National Industrial Security Program Operating Manual (NISPOM), July 1997.

[17] NSA/CSS, NSA/CSS Storage Device Classification Manual (SDDM), 2007.

[18] P. Guntmann, Secure Deletion of Data from Magnetic and Solid-State Memory, University of Auckland, July 1996.

[19] Royal Canadian Mounted Police, IT Media Overwrite and Secure Erase Products, IT Security Bulletin B2-002, May 2009.

[20] S. Skorobogatov, Data Remanence in Flash Memory Devices, CHES 2005 Workshop, August 2005.

[21] M. Wei, et al., Reliably Erasing Data from Flash-based SSDs, USENIX Conference on File and Storage Technologies, 2011.

[22] O. Al Homaidi, Data Remanence: Secure Deletion of Data in SSDs, MSc Thesis, Blekinge Institute of Technology, February 2009.

List of symbols/abbreviations/acronyms/initialisms

API	Application Programming Interface
AWS	Amazon Web Services
CSA	Cloud Security Alliance
CSS	Central Security Service
DAM	Database Activity Monitoring
DLP	Data Loss Prevention
DoD	Department of Defense
DSS	Data Security Standards
EC2	Elastic Computing Cloud
FAM	File Activity Monitoring
FBI	Federal Bureau of Investigation
FedRAMP	Federal Risk and Authorization Management Program
IaaS	Identity as a Service
IDA	Information Dispersal Algorithm
IdP	Identity Provider
IPsec	Internet Protocol Security
IT	Information Technology
ITAR	International Traffic in Arms Regulations
KMIP	Key Management Interoperability Protocol
LDAP	Lightweight Directory Access Protocol
LoA	Level of Assurance
NISPOM	National Industrial Security Program Operating Manual
NIST	National Institute of Standards and Technologies
NSA	National Security Agency
OASIS	Organization for the Advancement of Structured Information Standards
OSI	Open Systems Interconnection
OZ	Operations Zone
PaaS	Platform as a Service
PAN	Primary Account Number
PAP	Policy Administration Point

PAZ	Public Access Zone
PCI	Payment Card Industry
PDP	Policy Decision Point
PII	Personally Identifiable Information
PZ	Public Zone
RCMP	Royal Canadian Mounted Police
RZ	Restricted Zone
SaaS	Software as a Service
SAML	Security Assertions Markup Language
SecaaS	Security as a Service
SIG	Special Interest Group
SLA	Service Level Agreement
SP	Service Provider
SPI	Sensitive Personal Information
SQL	Structured Query Language
SSC	Security Standards Council
SSD	Solid State Devices
SSH	Secure Shell
SSO	Single Sign-On
TCG	Trusted Computing Group
TLS	Transport Layer Security
TMI	Trusted Multi-Tenant Infrastructure
TNC	Trusted Network Connect
TPM	Trusted Platform Module
URL	Uniform Resource Locator
VLAN	Virtual Local Area Network
VM	Virtual Machine
VPC	Virtual Private Cloud
VPN	Virtual Private Network
WS	Web Services
XACML	eXtensible Access Control Markup Language

DOCUMENT CONTROL DATA

(Security classification of title, body of abstract and indexing annotation must be entered when the overall document is classified)

1. ORIGINATOR (The name and address of the organization preparing the document. Organizations for whom the document was prepared, e.g. Centre sponsoring a contractor's report, or tasking agency, are entered in section 8.)

TRM Technologies Inc.
280 Albert Street, Suite 1000, Ottawa, ON, K1P 5G8

2. SECURITY CLASSIFICATION (Overall security classification of the document including special warning terms if applicable.)

UNCLASSIFIED
(NON-CONTROLLED GOODS)
DMC A
REVIEW: GCEC JUNE 2010

3. TITLE (The complete document title as indicated on the title page. Its classification should be indicated by the appropriate abbreviation (S, C or U) in parentheses after the title.)

Data protection in multi-tenant cloud environments: A technical brief

4. AUTHORS (last name, followed by initials – ranks, titles, etc. not to be used)

Magar, A.

5. DATE OF PUBLICATION (Month and year of publication of document.)	6a. NO. OF PAGES (Total containing information, including Annexes, Appendices, etc.)	6b. NO. OF REFS (Total cited in document.)
December 2012	46	22

7. DESCRIPTIVE NOTES (The category of the document, e.g. technical report, technical note or memorandum. If appropriate, enter the type of report, e.g. interim, progress, summary, annual or final. Give the inclusive dates when a specific reporting period is covered.)

Contract Report

8. SPONSORING ACTIVITY (The name of the department project office or laboratory sponsoring the research and development – include address.)

Defence R&D Canada – Ottawa
3701 Carling Avenue
Ottawa, Ontario K1A 0Z4

9a. PROJECT OR GRANT NO. (If appropriate, the applicable research and development project or grant number under which the document was written. Please specify whether project or grant.)	9b. CONTRACT NO. (If appropriate, the applicable number under which the document was written.)
15bs	W7714-08FE01

10a. ORIGINATOR'S DOCUMENT NUMBER (The official document number by which the document is identified by the originating activity. This number must be unique to this document.)	10b. OTHER DOCUMENT NO(s). (Any other numbers which may be assigned this document either by the originator or by the sponsor.)
DRDC Ottawa CR 2012-108	

11. DOCUMENT AVAILABILITY (Any limitations on further dissemination of the document, other than those imposed by security classification.)

Unlimited

12. DOCUMENT ANNOUNCEMENT (Any limitation to the bibliographic announcement of this document. This will normally correspond to the Document Availability (11). However, where further distribution (beyond the audience specified in (11) is possible, a wider announcement audience may be selected.))

Unlimited

13. ABSTRACT (A brief and factual summary of the document. It may also appear elsewhere in the body of the document itself. It is highly desirable that the abstract of classified documents be unclassified. Each paragraph of the abstract shall begin with an indication of the security classification of the information in the paragraph (unless the document itself is unclassified) represented as (S), (C), (R), or (U). It is not necessary to include here abstracts in both official languages unless the text is bilingual.)

This technical brief examines data protection within multi-tenant cloud computing environments. Specifically, this technical brief explores data protection strategies that can be employed within cloud environments to mitigate potential threats from other cloud consumers, as well as from third-party cloud providers. Further research in the areas of federated identity, authorization and encryption are required to address potential implementation issues that could arise in cloud computing environments.

14. KEYWORDS, DESCRIPTORS or IDENTIFIERS (Technically meaningful terms or short phrases that characterize a document and could be helpful in cataloguing the document. They should be selected so that no security classification is required. Identifiers, such as equipment model designation, trade name, military project code name, geographic location may also be included. If possible keywords should be selected from a published thesaurus, e.g. Thesaurus of Engineering and Scientific Terms (TEST) and that thesaurus identified. If it is not possible to select indexing terms which are Unclassified, the classification of each should be indicated as with the title.)

Cloud computing
Data protection
Information protection
Encryption
Identity federation

Defence R&D Canada

Canada's leader in Defence
and National Security
Science and Technology

R & D pour la défense Canada

Chef de file au Canada en matière
de science et de technologie pour
la défense et la sécurité nationale

DEFENCE DÉFENSE

www.drdc-rddc.gc.ca