Guidance

Cloud Security Guidance: Separation

Published 14 August 2014

Contents

1. Separation and cloud deployment models
2. Separation and cloud service models

Note: This publication is in BETA. Please send any feedback to the address platform@cesg.gsi.gov.uk.

This section of the Cloud Security Guidance provides more information on Principle 3: Separation between consumers. It examines the importance of cloud deployment models and cloud service models in understanding the separation requirements of a cloud service.

Separation is one of the fundamental security principles of any cloud service, and is required to prevent one consumer of the service interfering with the service (or data) of another.

1. Separation and cloud deployment models

The type of cloud service deployment model adopted will affect your security and assurance requirements. The following sections outline the concept of separation in relation to public cloud, private cloud, and community cloud deployments.

1.1 Public cloud

Public cloud services can be accessed by any public, commercial or government entity in possession of a credit card. For some services, an email address is all that is required to access free trial versions. Consumers using public cloud services must accept that their adversaries can legitimately purchase a service ‘next door’ to theirs. In such instances, services may need a good level of confidence in the ability of the service to protect their data.

1.2 Community cloud

Community cloud services host services for consumers from a specific community, such as the public sector (or commercial partners offering services to the public sector). These communities often have a shared risk appetite and generally expect conformance to an agreed minimum standard or legal agreement.

Community cloud providers can often tailor their offerings to match consumer requirements. For example, a service

provider could choose to meet specific UK government standards for personnel security screening, or conform to the required standard to connect to a government community network. These tailored offerings can reduce the risk relating to one or more principles.

Dedicating a service to a single community (where there is trust between members of that community) reduces the risks associated with separation between consumers. The trust between consumers will depend upon the measures set out in the standards the community members are obliged to conform to. This level of trust, and the types of applications deployed, will determine whether a community cloud service meets the required separation controls.

1.3 Private cloud

Private cloud services are deployed to support a single organisation. They normally offer the ability to tailor the architecture to meet specific security and business requirements. For example, if all consumers of the service are well known and low risk, then the level of assurance in separation required may be low. For processing untrusted or very sensitive data, then the organisation may require higher confidence in the separation controls.

Consumers will need to manage, monitor and maintain the infrastructure, unless an agreement exists with the service provider to do this.

In many situations a private cloud service will operate within a single security domain (for example providing a virtual desktop, or test and development resources within an organisation). In such scenarios, the cloud platform is simply another part of the enterprise IT environment and should be configured, managed and monitored as such. Security controls in private cloud environments normally do not need high levels of assurance, unless consumers have particularly challenging security requirements.

2. Separation and cloud service models

The technical controls required to provide separation will vary depending on the service model of the cloud service. The following sections outline the concepts of separation within the following offerings.

- Infrastructure as a Service (IaaS)
- Platform as a Service (PaaS)
- Software as a Service (SaaS)
- Risks associated with service models

2.1 Infrastructure as a Service (IaaS)

IaaS offerings generally provide compute, networking and storage services. Consumer separation needs to be enforced in all of these elements of the service.

IaaS: compute separation

Within the compute environment, separation between consumers is typically enforced by a hypervisor (though in some circumstances it may also be achieved by allocating physical hardware to consumers). The strength of separation usually depends on the virtualisation technology in use. Using hardware virtualisation and assured virtualisation products should provide higher confidence in the separation provided in the compute environment

(when configured in line with the accompanying security procedures for the product). The administration tools supporting the virtualisation product should be secured too, as they are fundamental to the security provided by the product.

IaaS: network separation

It is important to understand the network separation model in an IaaS offering since consumers will normally be constructing their own virtual networks on multi-tenanted network infrastructure.

A number of network separation technologies could be used by the service provider to enforce network separation. These include the use of virtual LANs (VLANs), virtual routing and forwarding (VRF) technologies or virtual networking capabilities within the compute environment. If consumers cannot gain sufficient confidence in the separation provided by an IaaS service at the networking layer, consumers may enhance their confidentiality protection with their own encryption. This is described in the Consumer Guides.

Consumers may also need to consider whether the service protects or reserves their share of network resources, so that an attack (such as a DDoS) on another consumer does not affect their service provision.

IaaS: storage

In an IaaS model, consumers may have direct control over an area of a multi-tenanted storage environment. Providing consumers with this direct control gives the ability for a single malicious or compromised consumer to directly launch attacks at the storage components of the service in order to gain access to another consumer's data. Therefore it is important that separation be enforced within the storage of the service.

IaaS consumers can reduce their reliance on storage separation by encrypting their own data. This presents its own challenges, and will only be effective if the encryption keys for the data can be securely stored in such a way that they would not be accessible to an attacker who has access to the storage.

2.2 Platform as a Service (PaaS)

PaaS offerings can provide rich and complex interfaces to consumers. They cover a wide range of implementation technologies, and are likely to be at different levels of security maturity. Depending on the technologies involved, it may be difficult, time-consuming and expensive to gain high levels of assurance in the separation provided by a PaaS offering. The CPA scheme does not currently have Security Characteristics for PaaS components, meaning that any formal assurance would need to be bespoke.

Applications wanting more assurance in the separation provided by a PaaS offering may want to build upon an IaaS offering which has sufficient assurance in such a way that separation will be enforced by the underlying IaaS separation controls. Alternatively, it may be appropriate to run the PaaS solution within a private or community cloud service.

Two common PaaS approaches, with different associated risks, are described below.

Shared application hosting

A common PaaS model, particularly for the delivery of web application hosting, is the hosting of consumers' applications on top of a shared operating system. In this model, the operating system and application host (eg the

web server and the scripting host or runtime) are responsible for preventing consumers from affecting each other. If attackers can legitimately run an application on the same host, they have access to a large attack surface to attempt to escalate privileges and gain unauthorised access.

Managed host services

Another common PaaS model is the provision of managed OS services. In this model, the consumer has a dedicated physical or virtual machine, but rather than manage the operating system themselves, they purchase this as a service from the service provider. The risks in this model are very similar to that of an IaaS service, since the separation enforcement is likely to be based on the same underlying technology (typically a hypervisor). There are two key additional risks (compared to an equivalent IaaS service) to bear in mind when considering the suitability of this model for an application:

- The service provider's administrators will have privileged credentials and access to the operating system. This means it is easier for them to access consumer data than if they only had access to disk images.
- In providing the management service, the provider will likely need to create additional connections between consumer machines and their management infrastructure. This infrastructure is an additional attack vector compared to the simple IaaS case.

2.3 Software as a Service (SaaS)

In SaaS offerings the separation between consumers is often enforced by software controls running within a single instance of an application. The strength of the separation is dependent on the application architecture and implementation. The underlying platform and infrastructure is not usually relied upon to enforce separation, which makes it difficult to gain high degrees of confidence in the strength of separation. In SaaS offerings it may be difficult to understand where and how consumer data is protected.

Some SaaS offerings may be able to be dedicated to a single consumer, leveraging the controls of an IaaS or PaaS solution which gives the consumer confidence in separation of their data.

In SaaS offerings the service provider will typically rely on application level controls rather than controls in the infrastructure or platform - meaning that if a component in the service is compromised then the data of many consumers may be visible to that component.

Applications wanting more assurance in the separation provided by a SaaS offering may want to build upon an IaaS or PaaS offering which has sufficient assurance in the separation in such a way that separation will be enforced by the underlying separation controls. Alternatively, it may be appropriate to run the SaaS solution within a private or community cloud service.

2.4 Risks associated with service models

The following table summarises the risks associated with each of the service models.

Service model	Associated risks
Infrastructure as a Service	Offerings implemented using hardware virtualisation and assured virtualisation products can provide a good level of separation assurance suitable for most OFFICIAL data in community and public cloud

(IaaS)	platforms. However, like all complex software, IaaS offerings will never be free from vulnerabilities and the risks that these bring.
Platform as a Service (PaaS)	PaaS offerings that share infrastructure between consumers (particularly where these may include malicious consumers) expose additional attack surface and may present a potentially higher risk than assured IaaS offerings. It may not be practical to gain robust assurance in a specific PaaS offering because not enough is known about the robustness of the underlying technology. These technologies are also evolving rapidly and consumers should regularly verify that their platform choice meets their business and security needs.
Software as a Service (SaaS)	SaaS offerings that share platforms or infrastructure between consumers (particularly where these may include malicious consumers) expose additional attack service and unless architected well, will often present a potentially higher risk than the same software installed within assured IaaS or PaaS offerings for a single consumer.

Legal information