CERT'S PODCASTS: SECURITY FOR BUSINESS LEADERS: SHOW NOTES

Controls for Monitoring the Security of Cloud Services

Key Message: Depending on the service model, cloud providers and customers can monitor and implement controls to better protect their sensitive information.

Executive Summary

The general characteristics of cloud computing's three service models include on-demand self service, broad network access, pooling of resources, rapid elasticity of provisioning resources, and service or resource monitoring. A cloud can be modeled in seven layers with controls at each layer. The responsibility for monitoring and implementing controls falls to either the cloud service provider or the cloud customer based on the service model. Such controls can help ensure, through auditing and monitoring, that sensitive information in the cloud is adequately protected. The contract or service level agreement between the customer and the provider is the keystone for building a strong trust relationship and ensuring reasonable interactions.

In this podcast, Jonathan Spring, a member of CERT's Network Situational Awareness team, discusses how cloud service users can better ensure that their sensitive information is as secure as possible.

PART 1: RAPID ELASTICITY OF RESOURCES

A Short Definition of Cloud Computing

According to the U. S. National Institute of Standards and Technology report DRAFT *Guidelines on Security and Privacy in Public Cloud Computing*, cloud computing is composed of five essential characteristics:

- on-demand self service
- broad network access
- pooling of resources
- rapid elasticity of provisioning resources
- service or resource monitoring

and three service models:

- software as a service (SaaS)
- platform as a service (PaaS)
- infrastructure as a service (IaaS)

Rapid Elasticity of Provisioning Resources

This is one of the key characteristics of using services in the cloud. For example, if a user wants to increase their CPU cycles from 1 billion to 10 billion, this can be done in a matter of seconds. This is possible due to the available resources owned and offered by the cloud provider. In other words, such resources are elastic given they are shared across many customers.

Customer demand for resources can grow or shrink as needed. Being able to access this capability via a cloud service can be much more desirable than an organization having to acquire and maintain these resources in house.

Considerations When Using Any Cloud Service

Regardless of service model, customers of cloud services need to consider the following before contracting for such services:

- Customers, as information owners, are traditionally responsible for monitoring the hardware and software that are used to store, process, and transmit their information. This is not the case when using public cloud services.
- Both customers and providers are at somewhat of a security disadvantage due to shared control, which results in a visibility gap (and who has control changes).
- The customer cannot technologically guarantee that there is no malicious process mediating hardware access by their software.

Some of the controls described below will help mitigate this risk.

Cloud computing comes down to a strong trust relationship between the customer and the service provider, supported by contracts and service level agreements (SLAs).

Customers need to ensure that they include adequate auditing and monitoring powers in their SLAs.

PART 2: CONTROLS FOR FACILITY, NETWORK, HARDWARE, VIRTUALIZATION, AND MIDDLEWARE LAYERS

Seven Layers of a Cloud Service

Based on the Cloud Security Alliance's work, a cloud is modeled in seven layers:

1. the facility
2. the network
3. the hardware
4. virtualization and the operating system
5. platform architecture including middleware, APIs, utilities, and libraries,
6. the applications
7. the users

The party responsible for implementing the controls listed below (service provider or customer) depends on the service model. All service provider controls should be explicitly called out in SLAs.

Based on providers dominating the market, customers likely have limited ability to assess or audit a provider's controls.

All of these controls (and others) are described in more detail in resources [1] and [2].

Controls at the Facility Layer

These are implemented by the service provider as the provider is responsible for the facility layer in all service models. Controls include:

- physical security of the facility
- robust policies for surveillance of the facility
- access by unauthorized personnel prohibited
- all personnel vetted to minimize insider threat
- continuity of operations plan for the facility developed, regularly tested, and addresses environmental disasters

Controls at the Network Layer

These are implemented by the service provider as the provider is responsible for the network layer in all service

models. Controls include:

- firewalls
- intrusion detection and prevention systems at the network layer border
- active IP (Internet protocol) and domain name black list policies (to keep bad things out)
- network monitoring to detect anomalies
- ability to analyze network flow data when an intrusion occurs

Controls at the Hardware Layer

These are implemented by the service provider as the provider is responsible for the hardware layer in all service models. Controls include:

- hardware monitoring for elastic rapid provisioning
- hardware monitored to ensure it is free from tampering (switched wires, etc.)

Controls at the Virtualization and Operating System Layer

These may be implemented by either the service provider or the customer based on the service model. Controls include:

- the central authority responsible for managing all virtual machines is secure and controlled
- non-essential functionality is removed from the operating systemintrusion detection and prevention systems at the network layer border
- an ability to return compromised virtual instances to a known, good state when they are compromised. This process should be documented.

Controls at the Middleware Layer

Middleware is hard to pin down. It may range from virtualization management tools and data format conversion to enforcing access controls. Middleware does include security functions, which are often a target.

These may be implemented by either the service provider or the customer based on the service model. For example, if the customer is only purchasing infrastructure as a service, the customer is responsible for these controls. Controls include:

- middleware that is securely coded and fuzz tested
- encrypted communication in an application programming interface (API)
- vetting the developer of the middleware if it is not the service provider

---

PART 3: CONTROLS FOR THE APPLICATION AND USER LAYERS

Controls at the Application Layer

These may be implemented by either the service provider or the customer based on the service models. Controls include:

- applications that are securely coded and fuzz tested
- for web-based applications, use of SSL (Secure Sockets Layer) and DNSSEC (Domain Name System Security Extensions)
- monitoring applications for compromise, particularly given that they are multi-tenant and memory persistent
- periodically refreshing application images from a known, good, read-only copy
- performing white listapplication monitoring (only execute permitted applications)

Controls at the User Layer

These may be implemented by either the service provider or the customer based on the service model. Controls include:

- periodic security awareness and training for all users (always a customer responsibility)
- monitoring of user access patterns, locking out accounts that behave in a suspicious manner. This is similar to credit card monitoring for fraudulent transactions.
- making sure users are trained in cloud service acceptable use, including handling of sensitive information in the cloud

The SLA is the keystone for ensuring reasonable interactions with your service provider.

Resources

[1] Spring, Jonathan. “Monitoring Cloud Computing by Layer, Part 1.” IEEE Security & Privacy Magazine, IEEE Computer and Reliability Societies, IEEE, March/April 2011.

[2] Spring, Jonathan. “Monitoring Cloud Computing by Layer, Part 2.” IEEE Security & Privacy Magazine, IEEE Computer and Reliability Societies, IEEE, May/June 2011.

Cloud Security Alliance guidelines

Jansen, Wayne & Grance, Timothy. DRAFT *Guidelines on Security and Privacy in Public Cloud Computing*. Draft Special Publication 800-144. National Institute of Standards and Technology, January 2011.

ENISA (European Network and Information Security Agency) Cloud Computing Risk Assessment

Google Apps Messaging and Collaboration Products, Google Security Whitepaper, 2011.

CERT Podcast: The Upside and Downside of Security in the Cloud, June 16, 2009.

Copyright 2011 Carnegie Mellon University