

SECURING MULTI-TENANCY AND CLOUD COMPUTING

Security That Ensures Tenants Do Not Pose a Risk to One Another In Terms of Data Loss, Misuse, or Privacy Violation

Table of Contents

Executive Summary	3
Introduction	3
What Is a Tenant?	3
Defining Multi-Tenancy	3
Securing the Multi-Tenant Environment	3
Hypervisor-Based Segmentation	3
Database-Based Segmentation	4
Segmentation Is Needed at All Layers	4
The Role of VM Introspection	4
Automation as an Enabler	4
Conclusion	5
About Juniper Networks	5

Executive Summary

In cloud-based architectures, multi-tenancy means that customers, organizations, and consumers are sharing infrastructure and databases in order to gain price and performance advantages. At its simplest, the “cloud” is an Internet-based environment of computing resources comprised of servers, software, and applications that can be accessed by any individual or business with Internet connectivity. In the case of these “service” offerings, customers (or “tenants”) get a piece of the cloud that contains the resources they need to run their business.

Cloud computing is the basis for infrastructure as a service (IaaS) and software as a service (SaaS). These services offer a pay-as-you-go lease style investment with little or no up front costs versus buying all of the hardware and software outright. Other benefits include the ability to scale easily and tier more services and functionality on an “as needed” basis (e.g., Salesforce.com and its various add-on modules and user-based pricing options). The benefits, in fact, are so compelling that cloud computing is predicted by some to be the replacement for traditional means of obtaining these services and business capabilities by 2014. The big concern is how to ensure that proper security and isolation protects consumers or tenants of these services from the risks they pose to one another.

Introduction

What Is a Tenant?

The notion of a tenant in the context of cloud computing is not as simple as it might first appear. Take Amazon Web Services (AWS), for example. AWS is a cloud service provider with offerings that span application hosting, backup and storage, e-commerce, and media hosting, to name just a few. Companies like Autodesk, Urban Spoon, and Second Life are tenants of AWS, in that they use AWS storage and compute resources to power their customer offerings. Each firm also has customers who store data like personal preferences, credit cards, and information as tenant users of these businesses. In the case of Second Life, for example, if the tenants set up online businesses and services of their own, they, too, will have tenants and so on. In the final analysis, a cloud service tenant is sharing a resource with a community. And similar to a building tenant, the tenant’s space must be separated and isolated from other occupants to achieve a certain degree of security and privacy.

Defining Multi-Tenancy

The idea of multi-tenancy, or many tenants sharing resources, is fundamental to cloud computing. Service providers are able to build network infrastructures and data architectures that are computationally very efficient, highly scalable, and easily incremented to serve the many customers that share them. Multi-tenancy spans the layers at which services are provided. In IaaS, tenants share infrastructure resources like hardware, compute servers, and data storage devices. With SaaS, tenants are sourcing the same application (e.g., Salesforce.com), which means that data of multiple tenants is likely stored in the same database and may even share the same tables. When it comes to security, the risks with multi-tenancy must be addressed at all layers. The next few sections examine how this can be accomplished for shared hardware and application infrastructure.

Securing the Multi-Tenant Environment

Hypervisor-Based Segmentation

Virtualization is quite often the platform that underpins IaaS offerings. Software, such as VMware vSphere, Citrix XenServer, and Microsoft Hyper-V, provides the means of turning a single piece of hardware into a physical host for many VMs. These virtual machines are the databases, file servers, application servers, and Web servers that comprise the typical physical network, and enable the traffic that makes commerce and communication over the Internet possible. They are also the servers offered to customers of IaaS for storing their data or running their web-based businesses. At its core, the virtualization platform includes a very specialized and optimized OS called the hypervisor, which in part serves to map traffic from VMs to the underlying VM host hardware so that it can make its way through the data center and out to the Internet and vice versa. The majority of security concerns in the virtualized infrastructure relate to the co-residency of machines owned by different customers. This places machines in a privileged position relative to one another. And this can elevate the risk for many types of breaches such as unauthorized connection monitoring, unmonitored application login attempts, malware propagation, and various “man in the middle” attacks. VM segmentation and isolation is also an absolute requirement for VMs containing regulation and compliance intense data like employee details, customer information, etc. Most regulatory mandates such as Payment Card Industry Data Security Standard (PCI DSS), Health Insurance Portability and Accountability Act (HIPAA), Sarbanes-Oxley Act (SOX), and Gramm Leach Bliley Act (GLBA) require that access be limited to a business’ need to know, and that control policies be set in place to enforce blocking of unwarranted access. Since the hypervisor intercepts all traffic between VMs and VM hosts, it is the natural place to introduce segmentation for the resources of IaaS tenants where VMs might be housed within the same VM host or VM host cluster.

APIs like VMware VMsafe have enabled an ecosystem of security solutions that embed inside the hypervisor for the purpose of introducing proper segregation, isolation, and protection of tenant resources—thereby enabling secure multi-tenancy. The security solution runs as a service inside the hypervisor and intercepts traffic or packets. In fact, those products supporting VM Introspection, a concept discussed later in this paper, will also have a great deal of information about the VM's state, including installed applications and services. Depending on the vendor of the security software, the solution may provide virtual network visibility to traffic, VM inventories, and VM compliance assessment, as well as application-based access control and malware suppression.

Database-Based Segmentation

Unlike IaaS where multiple tenants share resources, SaaS tenants share a database. Users of Salesforce.com or SmugMug, for instance, pay to use an application that manages their customers and photos respectively. While the value is in the application interfaces that make it easy to manage complex tasks and large data sets, the data itself is stored in a database as rows in tables that the tenants of Salesforce.com and SmugMug databases share. The customer ID is what distinguishes one row from the next. In this area, security concerns run high that misconfigured application code or an error in an access control list may put tenant information at risk of theft and misuse.

For controlling access to database data, there are quite a few tools and technologies available. What is usually implemented is a system for authentication and authorization of the access request so that only certain rows or fields are modifiable based on security policies that ensure that access is warranted. Encryption of data in the database is also common to protect it at rest, so that if it is ever compromised or stolen it would be difficult to decipher the underlying data.

Segmentation Is Needed at All Layers

The types of multi-tenancy security used largely depend on the cloud-based service and how it has been implemented. Most cloud service providers will provision security at all layers—mainly because they will have all types of multi-tenancy in their environments. IaaS users need to understand whether their VMs are being housed in the same host alongside those of other customers and what, if any, provisions the cloud service provider has made to isolate them. In the case where the onus is on the tenant to configure the segmentation, care must be taken to use expert advice in defining and maintaining access control policies that enable warranted access but limit risk. For SaaS tenants, the questions to ask are: How is my data, at rest, protected from prying eyes? And how is access authenticated, authorized, and differentiated so that only the right people are looking at and managing my data? Some questions also should be asked around the cloud service providers' policies for dealing with breaches if they occur (i.e., Are there service-level agreements, mitigation plans, etc.?).

The Role of VM Introspection

Relative to the Internet and network security technologies, virtualization platforms and cloud computing architectures are very new and still evolving. It is important to be aware of innovations that may augment security for multi-tenant environments but may not be broadly known or understood. It is often the case that the standards and reference architectures we rely on for proper implementation lag technological advancement. VM Introspection is a concept that has existed for some time in academic circles and is explained largely as a hypervisor-based service that examines the internal state of a running VM. Technologies have recently been commercialized that leverage VM Introspection in order to provide high levels of segmentation and isolation for guest VMs or cloud service tenants. VM Introspection provides rich detail about the applications and services that are installed on the VM, as well as its configuration. It is possible then to construct security policies on the basis of VM Introspection parameters. An example of such a policy might be: Do not allow a new virtual machine to join a VM group or cluster unless it has a specific OS configuration and hot fix installed. VM Introspection takes security for multi-tenancy to a new level where configuration errors are automatically prevented. This becomes especially important in environments where the onus for configuring security and VM isolation falls on tenants, who may or may not have experience in this area.

Automation as an Enabler

While security for multi-tenant environments might be the overarching concern for adoption, security automation will be the real catalyst for broad use of cloud-based services. Most will agree that the technologies to secure IaaS and SaaS architectures are broadly available and proven. The real challenge is that the tenants aren't always clear on which type of architecture they are using and what, if any, is their role and responsibility for protecting their information. Cloud service providers may implement the technologies, but may not fully control how they are managed and configured, as in the case where tenants themselves have sub-tenants. The key to securing multi-tenancy is for anyone who is a tenant (e.g., a business or consumer of IaaS and SaaS on some level) to ask the cloud provider about existing protections and responsibilities for defining and maintaining policies that ensure isolation from other cloud tenants.

Also key is to ask how much of the process is automated. Cloud computing environments, especially those based on virtualization, are extremely dynamic. Change is constant, and this makes the likelihood of resource and security misconfiguration high. With available technologies that automate VM protection (at least for IaaS), there is no reason to incur the higher risk, especially given the breadth of current and projected cloud service and provider options.

Conclusion

In cloud-based architectures, multi-tenancy means that customers, organizations, and consumers are sharing infrastructure and databases in order to take advantage of price and performance advantages that come with economies of scale. Tenants may share hardware on which their virtual machines or servers run, or they may share database tables where the data of customer A is on one row and that of customer B is on another. Many cloud service customers are comprised of both types of tenants. In either case, security measures are “a must” to ensure that tenants do not pose a risk to one another in terms of data loss, misuse, or privacy violation. Multi-tenancy protections must be offered by cloud service providers for all layers of their offerings (i.e., IaaS and SaaS). Cloud service providers owe it to their customers to have the latest and best approaches as available options. Tenants must ask and be clear about the ways in which they share responsibility for their security and the security of their tenants. Lack of security expertise needn't be a barrier to cloud service adoption, but security automation is key to making experts of would-be novices when it comes to securing a piece of the cloud.

About Juniper Networks

Juniper Networks is in the business of network innovation. From devices to data centers, from consumers to cloud providers, Juniper Networks delivers the software, silicon and systems that transform the experience and economics of networking. The company serves customers and partners worldwide. Additional information can be found at www.juniper.net.

Corporate and Sales Headquarters

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089 USA
Phone: 888.JUNIPER (888.586.4737)
or 408.745.2000
Fax: 408.745.2100
www.juniper.net

APAC Headquarters

Juniper Networks (Hong Kong)
26/F, Cityplaza One
1111 King's Road
Taikoo Shing, Hong Kong
Phone: 852.2332.3636
Fax: 852.2574.7803

EMEA Headquarters

Juniper Networks Ireland
Airside Business Park
Swords, County Dublin, Ireland
Phone: 35.31.8903.600
EMEA Sales: 00800.4586.4737
Fax: 35.31.8903.601

To purchase Juniper Networks solutions, please contact your Juniper Networks representative at 1-866-298-6428 or authorized reseller.

Copyright 2012 Juniper Networks, Inc. All rights reserved. Juniper Networks, the Juniper Networks logo, Junos, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners. Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.