InfoWorld

SPECIAL REPORT

JANUARY 2011

Cloud Security Deep Dive

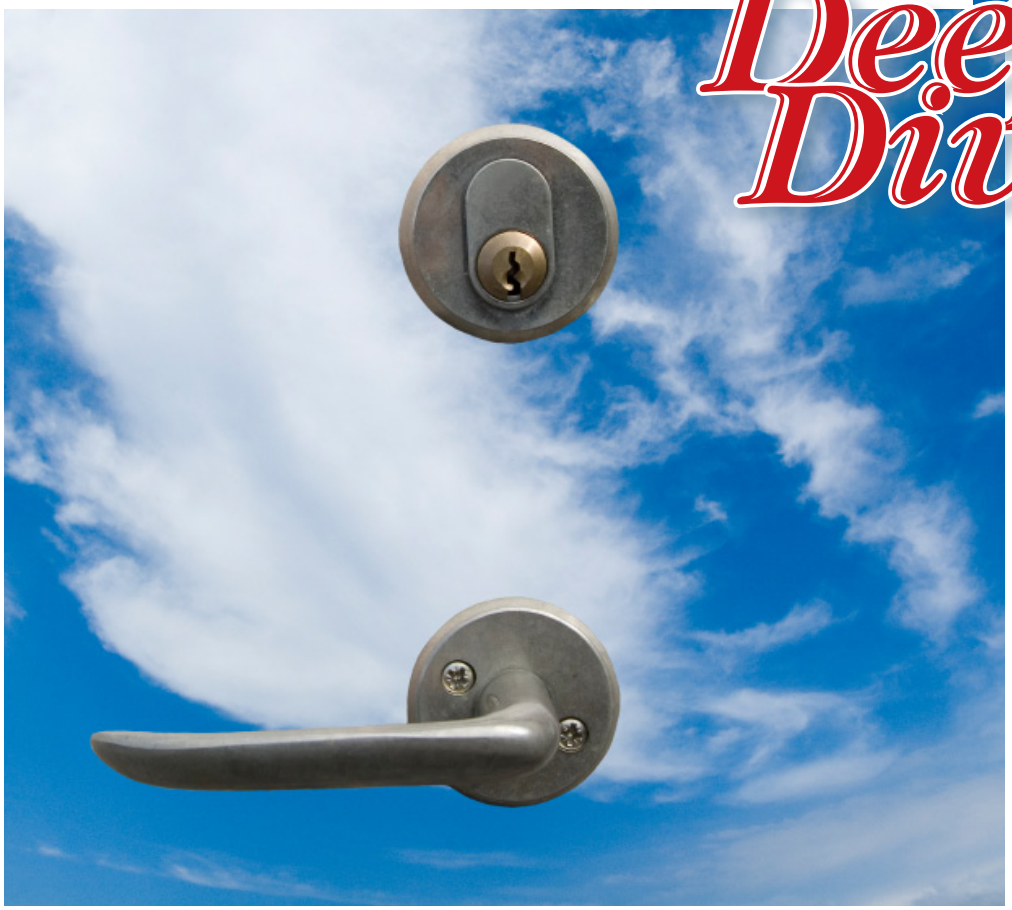

A new security model for a new era

Cloud security changes everything

The scalability of cloud computing depends on sharing resources that were never shared before, demanding a new set of security best practices

MANY COMPUTER SECURITY PRACTITIONERS blow off cloud computing as just a semantic exercise describing conventional applications running across a wide area network.

They are wrong. Cloud computing is a new paradigm that challenges traditional security dogma. Old assumptions are gone forever and the ones that replace them will make the security expert's job harder than ever.

This paper will cover how cloud computing differs from the past and discuss characteristics of the new threat model.

CLOUD COMPUTING MODELS

Before delving into cloud computing exploits and defenses, it helps to get a basic understanding of cloud computing models. Although new models appear to be emerging every day, here are the basic ones nearly everyone agrees on:

- **Infrastructure as a Service (IaaS).** IaaS provides massively scalable, elastic computing resources via the Internet. Some providers offer just a single resource, such as storage space, but most now focus on providing complete computing platforms for customers' VMs, including operating system, memory, storage, and processing power. Clients often pay for only what they use, which fits nicely into any company's computing budget.

- **Software as a Service (SasS).** SaaS providers deliver software functionality through the browser, without the end-user having to install software locally. Typically, SaaS offerings are multitenanted: Customers establish accounts on one huge instance of the software running on a virtualized infrastructure. Common examples include Google's Gmail, Microsoft's Business Online Productivity Suite, and Salesforce.com.

- **Platform as a Service (PaaS).** PaaS providers give developers complete development environments in which to code, host, and deliver applications. The development environment typically includes the underlying infrastructure, development tools, APIs, and other related services. Examples include Google's App Engine, Microsoft's Windows Azure, and Salesforce's Force.com.

Naturally, many cloud service providers mix two or more of these cloud service models or cannot be neatly placed into one type. Additionally, the type of user who can access the cloud further defines each model, as well as what type of cloud is at issue:

- **The public cloud.** Public clouds are created by one vendor and offered to the general public. Public clouds are almost always Internet-accessible and multitenanted.

- **The private cloud.** Private clouds are hosted by the same organization that utilizes the service (which in general does not support multitenancy). The primary value proposition is data accessibility and fault tolerance.

- **The hybrid cloud.** This term typically applies to organizations that have set up private cloud services in combination with external public cloud services. It also refers to service offerings used exclusively by an invited group of private customers (also called a "community cloud").

For more information, see the National Institute for Standards and Technology (NIST) document detailing common cloud terminology.

HOW CLOUD COMPUTING IS DIFFERENT

Cloud computing is a major departure from traditional networks and applications. In general, a service or offering is considered cloud computing if it has at least four of these seven traits:

- Internet (or intranet) accessible
- A massively scalable, user-configurable pool of elastic computing resources (such as network bandwidth, compute power, memory, etc.)
- Multitenancy (one large software instance shared by many customer accounts)
- A broad authentication scheme
- Subscription or usage-based payment
- Self-service
- Lack of location specificity

All of these traits offer new challenges to the computer security professional, but accessibility, multitenancy, broad authentication, and lack of location-specificity are the four items responsible for the biggest technology shift and demand for new security solutions.

ACCESS TO THE INTERNET OR INTRANETS EQUALS HIGH RISK

We already know that any computing resource that is Internet accessible is at a higher risk than one that is not. It's how most of the bad guys break into computers today, whether it involves social-engineered Trojan horse programs, viruses, or human attackers. High-risk environments, like the top secret classified systems of most governments, usually aren't allowed to connect to a network that can connect to the Internet. Too much risk.

Clouds don't use VPN technologies. With cloud computing, it's assumed that all users and application resources are Internet- or intranet-accessible, with all the elevated risk that implies. This means that anonymous attackers can access connection points just as any legitimate user or manager of the system can. In a traditional computing environment, only a small percentage of servers are Internet-accessible. In cloud computing, most servers are Internet-accessible. This must change the security defender's thinking.

Just as most clouds are Internet-accessible, they are also almost entirely Web-based, with browsers as clients and Web servers as the server endpoint connection. Some clouds use simple HTML-based forms and Web pages, but most are an increasingly complex set of Web services and protocols. (Wikipedia has an excellent beginner's tutorial on Web services.)

As the Web matures, most things that were once accomplished using a single computer will be executed by a matrix of Web services, connected together in most cases by XML, SOAP (or REST), and SAML. As Web 2.0 takes over, future security defenders must get to know these services and protocols inside and out, and defend against their deficiencies as they emerge (and there are bound to be plenty).

MULTITENANCY

Multitenancy is a major defining trait for public clouds. Typically, in a traditional environment, only the application's owner and direct employees can access the application data. In a cloud, multitenancy means that multiple, distinct, separate end-user parties share the same service and/or resources. End-users may be aware of this fact and may even be able to directly interact with other end-users. Or they may be unaware that resources are shared and that this is a risk.

In a cloud, risk looms that the parties sharing that cloud will be able to — unintentionally or intentionally — access one another's private data. This has been the case in cloud exploits with major cloud providers during the past few years. In some cases, all it took was modifying a client's unique identifier, sent over in the browser request, to another identifier, and up comes another client's data. Sometimes spillage has occurred when a bug in the cloud service offered up too much data without the client doing anything out of the ordinary.

With IaaS-based clouds in particular, security researchers are discovering a brand new class of vulnerabilities that did not exist in the old world. For example, attackers are finding ways to "cyberjack" another tenant's data and resources by discovering other tenant's IP addresses and computer resources or by searching for other people's data remnants after they release

unneeded resources back to the cloud. As it turns out, some cloud vendors don't erase or format the freed-up storage or memory resources.

It's too early to know whether these types of risks will decrease as cloud security matures or whether they will remain a fixture that defines the new threat model.

BROAD AUTHENTICATION SCHEMES

Internet accessibility and multitenancy pose a challenge when determining how to authenticate large numbers of different clients. In the traditional model, each authenticating user has a full user account located in the application's authentication database (or directory service). But scaling and multitenancy complicates the process, because conventional authentication services tend to offer access to shared common resources by default. In Active Directory, for example, members of the Everyone group can see and list all sorts of resources that a cloud provider would probably not want every client to see.

Initially, many cloud providers tried to solve this problem by using proprietary or private authentication services. But these services rarely have the scalability and functionality needed. First-generation clouds required that all end-users have separate accounts in their databases — similar to the way Web surfers need to log on separately to each website where they have an account. (For example, your Facebook account is not related to your Amazon or iTunes accounts). This is known as "Web Identity 1.0" in identity system circles.

Clearly, asking end-users to create and manage separate log-on accounts for every future Web service they will use doesn't scale — for a multitude of reasons. Using one big SSO (single sign-on) solution that interacted with participating Web sites became the "Web Identity 1.5" way of doing things. With this authentication model, users registered with an "independent," centralized authority to obtain an SSO ID. Then, when visiting participating Web sites, the user could enter their SSO credentials to gain access. An example of this sort of solution was Microsoft's Passport (now morphed into LiveID) and other protocols created by the Liberty Alliance.

But many people balked at the idea of a single entity handling everyone's SSO accounts. What evolved is known as "Web Identity 2.0," or federated identity, as well as identity metasystems. With Web Identity 2.0, there can be a multitude of identity services (made up of both centrally managed and single, stand-alone identities) that can interoperate with a large number of Web sites and services. Popular individual identity services include OpenID, InfoCard, and LiveID. Many of these authentication services are interoperable with each other and use common protocols such as XML, SOAP, SAML, Web services, WS-Federation, and so on.

In the Web 2.0 world, each Web site (or cloud provider) can choose which federated identity service to work with and accept. The Web site or service provider can require particular types of identity assurance (such as a simple password, smartcard, or biometric device) before a user can participate. For example, a cancer survivor Web site may wish to allow anonymous users whereas an online banking Web site may require a site-issued smartcard or other authenticating token to log on.

Conversely, users may be able to submit only the identity data they wish to share with the participating service provider (called claims-based identity). For example, a user purchasing alcohol via the Internet may need only to prove that he or she are over the legal drinking age but not have to show his or her actual identity or birthdate. A central tenet of any good identity metasystem is that users should only have to show the bare minimum of identity information necessary to access the offered service and perform the desired transaction. Submitting (or requesting) too much identity information is considered very "Web 1.0."

In the Identity 2.0 model, authenticated anonymity (called pseudo-anonymity) is possible. In this case, a trusted third party knows the user's real identity and has authenticated him or her, but hands out a different identity credential that is trusted by the Web service provider. Thus, that user can use the Web service without revealing his or her true identity to the Web site.

THE EFFECT OF BROAD AUTHENTICATION ON CLOUD COMPUTING

In the near future, it's likely that both private and public clouds will support Web 2.0 identities. Users of private clouds will likely use their SSO to access public cloud services, while external users may use their SSO to access your company's private or hybrid cloud offering.

The security impact of this is that a cloud attacker is likely to be an authenticated user within the cloud system at the onset of an attack. By contrast, the old assumption is that the attacker didn't start with authenticated access and needed to gain original access to begin high-level system exploitation.

For example, consider two of the earliest and largest cloud services: Google Gmail and Microsoft Hotmail (now called LiveMail). Both contain millions of authenticated users and both now support newer SSO forms such as OpenID and InfoCard. Google and Microsoft have no idea which users are legitimate and which intend to do harm to other users or to the service itself.

Identity is still a work in progress. Solutions will change and morph as people begin to adopt the cloud in great numbers. And as new needs emerge, new solutions and protocols will need to be invented. Whatever trajectory these future developments take, Identity 2.0 is new paradigm that requires a huge mind shift in computer security defenders.

LACK OF LOCATION SPECIFICITY

The term "cloud" implies that a service is available widely, if not globally, with a multitude of origination and destination points. The specific location of the computing resources for a given cloud may not be immediately identifiable by either the client, the cloud vendor, or both. In a traditional network offering, the user or vendor often is aware of where the application or data is being hosted.

This brings up all sorts of interesting dilemmas. For example: How are security defenders supposed to protect data when they don't even know where it is? How can a cloud provider identify a client's data (for legal and other purposes)? How does the cloud provider securely erase a client's data if the client exits the cloud solution? How can a particular client's data be prevented from leaving the host country of origin, if even the cloud operators don't know where the data is?

THE ROLE OF VIRTUALIZATION

Virtualization tends to play a big role in cloud services, either as the underpinning for the service or, in the case of IaaS, as part of the cloud service offering itself. And virtualization has every security risk that a physical computer environment has — plus guest-to-host and guest-to-guest vulnerabilities.

Clearly, cloud computing amounts to more than just semantic games. It presents a unique set of challenges that security defenders must rise up to meet.

CLOUD SECURITY DEFENSES

Cloud security is an evolving field. To begin with, cloud solutions are subject to all the conventional attacks — buffer overflows, password attacks, physical attacks, exploitation of application vulnerabilities, session contamination, network attacks, man-in-the-middle attacks, social engineering, and so on. But the unique characteristics of cloud computing present a new set of challenges as well.

Start by assuming attackers are logged-in, authenticated users, and begin your defenses from there — and there may be many attackers, thanks to the cloud's global reach. This level of attacker access means that many conventional defenses (such as separated security zones, firewalls, and so on) will have little relevance.

CLOUD SECURITY DEFENSE CLASSIFICATIONS

Managing cloud security is different than maintaining ordinary enterprise security. Security professionals should analyze all cloud offerings (including their own) within a holistic security framework to make sure all angles are covered.

Keep in mind that each combination of cloud services has its own unique set of risks and countermeasures. Table 1 organizes these variables into a set of classifications and subtopics intended to generate further discussion and analysis.

CLOUD DEFENSE BEST PRACTICES

Enterprise security is a vast discipline and each of its many aspects must be reexamined in light of the cloud. Take user deprovisioning as an example. Normally, when an employee leaves a company, access to company applications and data is removed. But if that employee has subscribed to cloud services on behalf of the company, from the beginning the company should have had the technology in place to track those subscriptions, or the former employee may still be able access company data. You can't deprovision if you don't know what was provisioned in the first place.

TABLE 1: CLOUD SECURITY CLASSIFICATIONS AND SUBTOPICS

A security defender responsible for cloud security should consider a wide range of parameters when developing a cloud security defense plan.

SECURITY CLASSIFICATION	RELATED SUBTOPICS
Infrastructure security	Physical security, environmental controls, business continuity/disaster recovery, network infrastructure, firewalls, proxies, routers, access control lists, staffing/employee background checks, availability (performance and anti-DoS), security policies (including what can be made available to customers), remote access, mobile access and platforms, identity/authentication/federation, billing systems, virtualization issues, high availability
Resource provisioning	Provisioning; modification; ownership and control, access; deprovisioning; reuse/reassignment of: users, computing resources, computer systems, or IP address space; domain name services; directory services; self-service configuration management
Storage and data security	Privacy/privacy controls, data tagging, data storage zoning, data retention policies, data permanence/deletion, encryption (at-rest, in-transit, key management, Federal Information Processing Standards/Federal Information Security Management Act), digital signing/integrity attestation, multitenancy issues, archiving, backup, recovery, data classification, locality requirements, malicious data aggregation prevention
Application security (if applicable)	Security design lifecycle, identity/authentication/federation, session management, data input validation, error handling, vulnerability testing, patching, authentication, data integration/exchange, APIs, proxies, application sandboxing, versioning, bug/issue tracking
Audit/compliance	Logging, monitoring, auditing, compliance, accreditation, legal issues, regulations, locality requirements, discovery, forensics, SLAs, public communication plans, fraud detection
General security	Anti-malware, anti-spam, patching, incident response, data leak prevention

Implications of this type apply across a whole range of security issues. Here are some key ones to consider. For better or worse, the degree to which you can apply best security practices often depends on the provider.

SAY GOODBYE TO THE DMZ

One of the major changes from traditional computing is the pervasiveness of the cloud. By its very definition, it is meant to be everywhere. If the DMZ wasn't dead before, it certainly is now. The DMZ was always porous, with many more holes than any defender wanted to admit. The cloud, with authenticated attackers, just puts the nail in the coffin.

What is a security defender to do? Well, for one, think in terms of data classification and ownership and marry that with strong security domain isolation. Cloud providers should have ways for defenders to mark or tag data with ownership and security classification and to enable defenses based upon those attributes. Data should be protected in such a way that unauthorized (but authenticated, multitenanted) viewers can be prevented from seeing another's data. If a client needs to prevent its data from leaving its home country, the cloud provider should make sure the data never does.

Cloud providers should physically prevent (using physical network dividers, routers, switches, IPSec, access control lists, and so on) server and data commingling. If a particular server never needs to talk to most other servers, it should be prevented from doing so. If a client computer shouldn't be able to talk to other client computers, make sure there is no way for an authenticated users to leverage cloud access to gain access to the other.

ENCRYPT YOUR DATA

Data should always be encrypted when stored (using separate symmetric encryption keys) and transmitted. If this is implemented appropriately, even if another tenant can access the data, all that will appear is gibberish. Shared symmetric keys for data encryption should be discouraged, yet tenants should be able to access their own encryption keys and change them when necessary. Cloud providers should not have ready access to tenant encryption keys.

USE STRONG AUTHENTICATION

Make sure any Identity 2.0 system you participate in has a strong history of good security and uses open protocols. Proprietary, single-site authentications systems may seem to present lower risk than shared systems do, but the information surrounding proprietary systems is rarely shared. Systems that use and support open standards usually have the added protection of community analysis. Weaknesses are frequently found early and coded out. Newfound vulnerabilities are usually shared and fixed faster.

PREPARE TO PREVENT DDoS ATTACKS

Attackers are often content with simply denying legitimate users access to their services using DDoS attacks. Luckily cloud systems are usually very resilient against simple flood attacks and excel at ramping up more bandwidth and resources in the face of gigabytes of malicious traffic.

Be aware, however, that attackers may attempt to take down upstream or downstream nodes that are not under the control of the cloud provider. More than one Internet access provider has been forced to cut off access for a victimized site simply to preserve access for everyone else. You shouldn't be forced to come up with creative defenses the first time you're hit with a DDoS, so make sure you have adequate DDoS defenses and response plans ready to go. All the hard work and peering agreements should be established ahead of time.

USE DNS SECURITY

If your cloud provider's DNS services support it, consider implementing DNSSEC (DNS Security) between your DNS servers and the provider's DNS servers. Once enabled, DNSSEC ensures that dependent clients always get verified, authenticated DNS resolution entries from authoritative DNS servers. Unfortunately, DNSSEC is enabled only on a small percentage of DNS providers. Ask your cloud provider if it will consider creating DNSSEC "trust anchors" between your site(s) and the provider's — or consider implementing static DNS records on your side, to prevent malicious DNS redirection attacks.

USE RED HERRINGS AND AN EARLY WARNING SYSTEM

Some cloud providers and customers create "red herring" data as an early warning system. Red herring data is fake data that is injected into a database and then monitored to see if it "leaks out." For instance, suppose the cloud provider creates complete client records for Fred and Wilma Flintstone. Everything about the fake record would be unlikely to exist in the real world. Then the cloud vendor (or client) uses data leak detection systems and procedures to monitor for that specific data. If the data is found outside the cloud provider's system, then it could be an early warning that malicious data theft has occurred and should be investigated.

KILL ALL OLD DATA

Cloud providers should ensure that all data no longer needed is permanently erased from computer memory and storage. Shared resources shouldn't mean permanently shared storage. Clients should contact cloud providers to make sure that all submitted data is solely owned by the client and learn what measures providers take to ensure the permanent deletion of unneeded data.

ADDITIONAL RESOURCES

Cloud computing is a new paradigm that requires new security defenses. The material included here covers only a small portion of the considerations you need to weigh when preparing a holistic cloud defense. The following Web assets are excellent sources of additional information:

NIST's cloud section
http://csrc.nist.gov/groups/SNS/cloud-computing
A great place to quickly get up to speed on cloud terminology without reading a book.

Cloud Security Alliance
http://www.cloudsecurityalliance.org
Site represents a good collection point for enterprise-level cloud-related security. Look under the New Research section.

Cloud Security Alliance IT Certification
http://www.cloudsecurityalliance.org/certifyme.html

Cloud Threats and Security Measures, MSDN, J. Meier
http://blogs.msdn.com/b/jmeier/archive/2010/07/08/cloud-security-threats-and-countermeasures-at-a-glance.aspx

Black Hat Webcast: Chewing the Cloud: Attacking Cloud-Based Services
http://www.blackhat.com/html/Webcast/Webcast-2010_cloudsec.html
An interesting Webcast on cloud attacks.