THE DATACENTER OF THE FUTURE

Chapter 1: Security, Privacy and Risk Management

Author Jon Greaves,
Vice President and CTO,
Carpathia Hosting

Q/A with Ron Gula,
co-founder and CEO of
Tenable Network Security

November 2, 1988—a key day in the history of computer systems security. On this day, we realized how vulnerable networked computer systems could be when Cornell University graduate student, Robert Morris, released what is widely acknowledged as the first computer worm. Morris' code exploited vulnerabilities in two UNIX programs, Sendmail and Fingerd. After realizing his code had spread much more widely than expected, he attempted to alert everyone of the problem, but it was already too late, and the Internet—still in its nascent form—had become crippled.

Fast forward to 2001—Code Red, Code Red 2 and Nimda exploited vulnerabilities in software, but the difference this time, was the shear number of infected hosts fueled by the dramatic growth of the Internet, and the intent of the attacks being purely malicious. Over 2.2 million computers - 65% of which were servers - were infected by Nimda, making it the most destructive computer program ever seen to date. It's estimated the worldwide cost of these three worms was $3.2B.

Before 1988, the Internet was considered "innocent" with very few security controls in place. In fact, open sharing of computing resources and data became commonplace. Morris' worm was a wakeup call to Internet users causing them to take security solutions more seriously, which over time has come to include things like firewall and the intrusion detection systems we are familiar with today.

Enter in the era of perimeter protection where one considers the local network "trusted", and the remote/Internet as "untrusted". A firewall marshals traffic between these two networks with an intermediary network known as a demilitarized zone (DMZ), allowing local applications to publish themselves to the Internet. Perimeter protection was a fairly successful strategy to repel borders with the big assumption that your internal network was completely trustworthy.

Nimda proved this to be a false sense of security when laptops - which connected to the Internet from homes, hotels, airports, etc., - became infected. These laptops were then taken to the office, connected up to the trusted network, and then transported the worm in this case, onto the network. The worm had free reign to attack the local network and with the local network speed, it was able to replicate much faster than it did on the Internet.

When acknowledged that perimeter protection was insufficient, the search for a new methodology began. As with many computer security inventions, military strategy formed the basis: Defense in Depth. This new approach was conceived by the National Security Agency (NSA) to increase the survivability of an environment in the event of an attack by using a layering tactic.

DEFENSE IN DEPTH—frequently referred to as "security in depth"—is often visualized as the petals of an onion, with perimeter protection being the outermost layers. The goal is to delay a hacker or an attack to allow a busy organization time to detect and respond before the attack gets close to critical systems. As you peel back layers, additional controls are in place such as hardening of the operating system and cryptographic trust of connections between applications or machines.

This model has served us well. Today, unified security management systems, which combine physical and IT security systems, correlate all of these controls to inform the administrators of new threats as they become active. True Unified Security Management Systems consider all aspects of network security: People, Processes and Technologies. You want to know that your analyst who just walked in the building in Washington DC cannot possibly also be entering a server room in Marino Del Ray, CA. The security teams also now have the ability to lock a terminated employee not only from accessing the network once terminated, but also from the physical building.

Defense in Depth—frequently referred to as "security in depth"—is often visualized as the petals of an onion, with perimeter protection being the outermost layers. The goal is to delay a hacker or an attack to allow a busy organization time to detect and respond before the attack gets close to critical systems.

Now, the industry is moving to providing IT capabilities as a service—via the cloud or Internet—allowing users to access technology without having the hardware or software in the same physical location as themselves or even knowing where the infrastructure is.

As enterprises embrace the next evolution in cloud computing, we reveal many new challenges for information security. An IDC recent survey showed that "the number one concern about cloud services is security" (http://blogs.idc.com/ie/?p=210). But first, it's not just the security of the system that's key, it is also ensuring proper stewardship of sensitive information stored on systems such as customer or employee records. Second, when applications are running on 10's if not 100's of servers, being shared with a number of other businesses all of which are outside your "perimeter", this can create some very interesting challenges to solve.

We now need to engineer "survivability" layers into our virtual machines, which run in cloud computing environments. This starts with:

1. Ensuring your virtual machine is patched to the latest revisions of the operating system.
2. Hardening operating system images to verify all unused services are removed and proper controls are in place so that only trusted parties may establish a connection to the operating system.

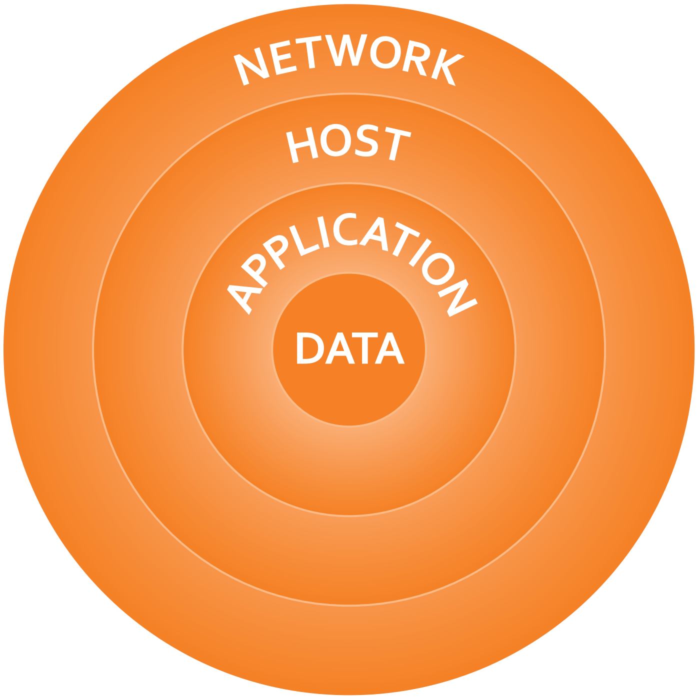

Source: Wikipedia

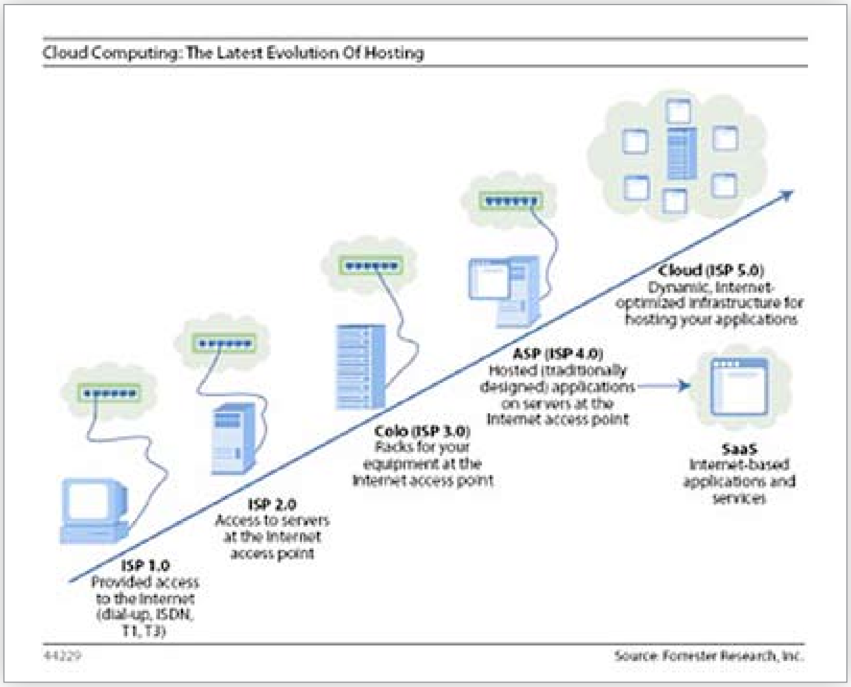

Cloud Computing: The Latest Evolution Of Hosting

ISP 1.0
Provided access to the Internet (dial-up, ISDN, T1, T3)
ISP 2.0
Access to servers at the Internet access point
Colo (ISP 3.0)
Racks for your equipment at the Internet access point
ASP (ISP 4.0)
Hosted (traditionally designed) applications on servers at the Internet access point
Cloud (ISP 5.0)
Dynamic, Internet-optimized infrastructure for hosting your applications
SaaS
Internet-based applications and services
44229

Source: Forrester Research, Inc.

Our perimeter is now a single machine, so perimeter controls must be installed on the virtual machine itself. This takes the form of a host-based firewall and potentially host based intrusion detection (and prevention). As you can tell, by distributing security to this level, it creates orders of magnitude, more complexity, and will require an evolution in unified management consoles to cope with the volume and dynamic nature of these new computing platforms. Along with the challenges come many new opportunities to take security to the next level. Work is going on now in the virtualization layer to build in new forms of protection to mitigate against attacks by looking for anomalies in the usage of the resources as a trigger that something dangerous may be occurring.

Since a virtual machine image is really a collection of files that when not under the control of virtualization, software packages are “inert”. This creates an opportunity for vulnerability scanning tools that today are used to verify configuration against live systems to take advantage of the image format and perform analysis of threats offline. The same with patching, a process that today is performed on a live running machine and which typically causes disruption to the operations of the application can now be considered an offline activity allowing verification/testing to take place in a safe environment before promotion to production. This is a best practice for patching in any environment, but rarely carried out in the real world due to the cost of creating a separate environment for testing purposes.

With the ever increasing numbers of regulatory standards, another challenge exists for those who wish to embrace cloud computing. For example, lets assume you are running an ecommerce site that accepts credit cards to process customer orders. As part of your architecture you decide to embrace a public cloud solution to provide capacity during your busy season from Cyber Monday through to the new year. This offers you many advantages such as not purchasing additional servers used only for three months a year but you still are required to comply with Payment Card Industry (PCI) standards. PCI consists of six control objectives. These objectives were created with a very traditional IT infrastructure in mind. So for example the first objective states “build and maintain a secure network” with the first requirement “install and maintain a firewall configuration to protect cardholder data”. With your ecommerce system now running virtually in a cloud computing environment gathering evidence and providing you are meeting these items becomes extremely difficult. ■

Q & A

A brief interview with Ron Gula - co-founder and CEO of Tenable Network Security - shows how cloud computing could impact the current views on security.

Q. ***How do you see the adoption of cloud computing impacting the way we think about security today?***

A. As with any new technology, there are advantages and disadvantages. I got a good start in my career working for US Internetworking in the late 90's where they were able to get customers to outsource their critical applications like Peoplesoft and SAP. I would love to make the argument that we were more secure than the customer and that this was the main reason they wanted to give us their business. However, each customer was different. Sometimes, they were outsourcing because of the cost model (rent vs. own), sometimes it was a manpower issue (USi was 24x7, and they were not), sometimes it was Internet bandwidth and sometimes it was security. Of course in the 90s, some customers were still impressed that we had firewalls.

Today, as we look into the 21st century, I feel cloud computing will be something used in every organization mostly as a method to save on costs. My concern is the users of cloud computing will wash their hands of the security issues surrounding cloud computing. Who runs these servers? How secure are they? How reliable are they? Some organizations that are interested in these types of technologies might not even know to ask these sorts of questions.

Q. ***One of the challenges of any security system is the sheer volume of data that needs to be processed and interpreted. Unified Threat Management was the solution to this in today's infrastructure solutions. What role do you see UTM providing in cloud computing environments?***

A. I think one of the big problems with security today is computers are too flexible. They can have a variety of purposes, uses and configurations. This gives the rise to complexity, which is often said to be the enemy of security.

My hope is that in a cloud computing environment, customers will make use of single purpose applications. For example, consider a web farm that runs 1000s of web servers. I would expect they are all configured, secured, patched and hardened the same way. This save you money and time and also makes it easy to spot when something isn't configured correctly. If you have single purpose servers that are used a certain way, when they break, become compromised or have some sort of error, they behave differently. And lastly, when you go to harden these single purpose applications, it is much easier to know how they will work so you can put appropriate security measures like firewalls and system security settings in place.

My point here is that if done right, outsourcing a single application to a cloud computing service can be very efficient and secure. If you were to compare this with an organization which simple provided Linux operating systems to you, and it was up to you to configure and run these your selves, you might still be "in the cloud" but you don't have any of the benefits of the single purpose applications.

And finally, to get back to UTM, if you have a cloud computing environment which is single purpose (like a bunch of similar configured web servers) your UTM should be looking for behavior indicative of a compromise or error. These are deviations from "known good" behaviors. In a random or mixed environment, the UTM will be looking for "known bad" behaviors such as virus outbreaks, attacks detected via intrusion detection rules and so on. There has been much written on looking for known good and known bad behaviors. I am very much in favor of looking for "known good" but I also understand enterprise networks can be complex, even if there is an attempt to keep things simple. Either way, you need a UTM (SIM, Firewalls, logging, IDS, anti-virus, .etc) to watch your network. I just feel you are much more effective when monitoring for "known good" than "known bad".

Q. ***All good security solutions blend proactive and reactive security systems as a way provide a holistic picture of an environment. How do you see these tools adapting in highly virtualized and dynamic computing environments?***

A. There are some very, very cool reactive network security technologies that have been produced over the past decade. Unfortunately, I see very few of these being deployed operationally. The issue is reliability.

For example, if you want to reconfigure a firewall after a network IDS sees an attack, the IDS better be right more than %99.99999 of the time. The first time it is wrong and legitimate traffic is blocked, you have both a technical issue of needing to fix this detection rule, as well as a political issue of impacting legitimate traffic.

What I do see is that anytime an organization can combine hardening of their network to only allow authorized services with automation, they usually have a well-run network. Hardening a network means different things to different people, but through the use of firewalls, running minimal configurations per host and having minimal user accounts a network can reduce the amount of potential attack space that can be exploited by an insider or outsider. Automation makes things happen regardless if a user is there to run the test as well. For example, patch and configuration auditing can detect a vast majority of missing patches and configurations, which are against policy.

> "Hardening a network means different things to different people, but through the use of firewalls, running minimal configurations per host and having minimal user accounts a network can reduce the amount of potential attack space that can be exploited by an insider or outsider. "

In virtualized environments, this is no different. The fact that a system is virtualized does not make it any less immune to an attack. If an organization does not have the proper approach to looking for unauthorized activity, configurations and changes to their network systems, be they virtualized or real servers, they will likely have many servers that are vulnerable to exploitation of some sort.

Q. ***What opportunities does cloud computing provide to security companies? Do you see the management of security itself becoming a cloud service?***

A. As we move further into the 21st century, we will see the emergence of new types of business models as well as new types of technologies that enable new types of services. In the late 90s, the state of the art MSP could watch your firewall and do some automated vulnerability scanning. Today, you can get an MSP to run your SIM, gather all of your logs, perform brand protection and certify that your ecommerce system meets the standards of the credit card industry. You can also get services for almost every type of function that occurs in your network including authentication, secure email and SPAM filtering, secure web hosting, secure chat hosting, secure DNS, secure data storage, secure SQL databases and so on. Many of these service companies offer combinations of various types of services as well.

What this means for a security company is that they have options. If are running a security company and want to deliver a service to your customers, you now need to calculate if running your own infrastructure truly gives you any advantage over running your own. The advantage could be a cost savings, a time to market savings, or even some sort of scalability that would be hard to do alone.

Lastly, if you are funding a security company and concerned about cash flows, sometimes it is difficult to decide how much money you should invest in your infrastructure before going live or making any profit at all. With cloud computing, you can focus on getting your service offering correct and purchase what you need from a cloud computing vendor as you go. ■

Ron Gula *is the co-founder and CEO of Tenable Network Security which produces a variety of vulnerability configuration auditing, log analysis and network monitoring tools including the Nessus vulnerabiltiy scanner. He was also the original author of the Dragon IDS and CTO of Network Security Wizards, which was acquired by Enterasys Networks. At Enterasys, Mr. Gula was Vice President of IDS Products and worked with many top financial, government, security service providers and commercial companies to help deploy and monitor large IDS installations. Mr. Gula was also the Director of Risk Mitigation for US Internetworking and was responsible for intrusion detection and vulnerability detection for one of the first application service providers. Mr. Gula worked for BBN and GTE Internetworking where he conducted security assessments as a consultant, helped to develop one of the first commercial network honeypots and helped develop security policies for large carrier-class networks. Mr. Gula began his career in information security while working at the National Security Agency conducting penetration tests of government networks and performing advanced vulnerability research. Mr. Gula has a BS from Clarkson University and an MSEE from the University of Southern Illinois. Ron Gula was the recipient of the 2004 Techno Security Conference "Industry Professional of the Year" award.*

Jon Greaves*, vice president and chief technology officer at Carpathia Hosting, Inc., is a recognized leader in the information technology services industry. In his current role, Greaves will lead Carpathia's commitment to use technology to deliver innovative services and to ensure the delivery of repeatable and scalable solutions to Carpathia customers. He will also focus on continuing the evolution of the strategy and products and services portfolio at Carpathia. Greaves has spent a significant part of his career in managed services and operations with a particular emphasis on remote services delivery. Prior to joining Carpathia Hosting, Greaves was a Distinguished Engineer and held the role of CTO of Sun Microsystems' Services business. Prior to Sun, Greaves was instrumental in leading the architecture and development of SevenSpace's industry-leading remote management and delivery platform. He is recognized as an expert on systems security and privacy, playing a strategic role in representing the telecommunications industry in developing international standards in response to Critical Infrastructure Protection and U.S. Presidential Decision Directive 62 and 63. Greaves has also held positions at British Telecom, MCI and Concert.*

CARPATHIA™
HOSTING, INC.

Carpathia Hosting, is a leading provider of enterprise managed hosting services for government agencies and businesses that require Colocation, Managed Services, Data Center Management, and Cloud Computing. Employing dynamic technologies that remove hardware dependencies and improve efficiencies, Carpathia Hosting solutions strive to reduce operational costs while surpassing SLA requirements. As a datacenter neutral company, Carpathia Hosting is quickly becoming the hosting company of choice for companies that demand security, quality and high performance.

CORPORATE
43480 Yukon Drive, Suite 200 Ashburn, Virginia 20147
Voice: 1.703.740.1730 Toll Free: 1.888.200.9494 Fax: 1.703.997.5577

DATA CENTERS
Ashburn, VA | Harrisonburg, VA | Phoenix, AZ | Los Angeles, CA

References to other products are made to show compatibility. All companies and/or products mentioned in this document are registered or trademarked by their respective organizations. The inclusion of third party products does not infer endorsement by these parties, unless otherwise noted.