denial of service attacks 2 over a coffee

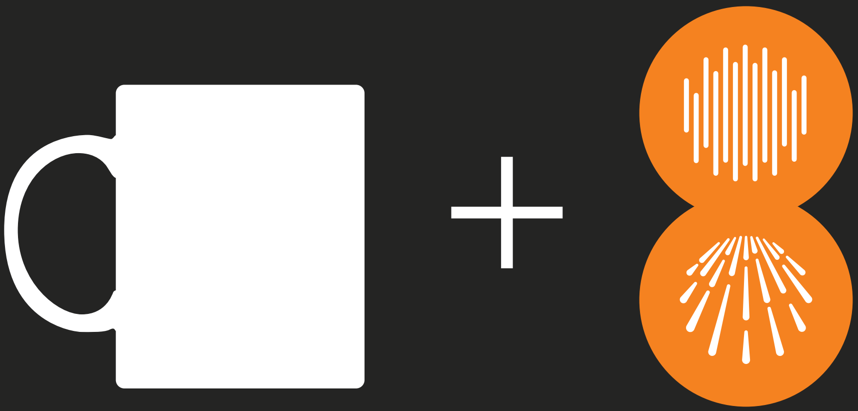

Business Services

orange™

editorial

Whether they are geopolitical, cybercriminal, financial, or commercial in nature, many different reasons can motivate a distributed denial of service (DDoS) attack. They are numerous and adapted to different contexts: cyber war, hacktivism, extortion, blocking the competition, etc. This has led to the development of a new market for attack services that are worthy of the best professional security packages.

In this second blog book dedicated to denial of service attacks, you will learn all about "commercial" DDoS services and one technique used for amplification attacks. You will also discover the likelihood that you too will take part in illicit operations—without your knowledge. Finally, you will see that the DDoS phenomenon has only just begun, thanks in part to cloud computing services.

Understand, anticipate, protect. These are the three keywords you need to know when dealing with denial of service attacks.

Vincent Maurin

content

real-time info from the IMDDOS helpdesk

by Jean François Audenard

As you've likely already heard, a commercial service for launching distributed denial of service (DDoS) attacks, called IMDDOS, was discovered by Damballa, a company that identifies zombie computers (botnets).

an innovative commercial endeavor in this sector

In itself, this service is not new. What's striking is that this service is now **all over the web** (no need to scour for some obscure forum or Internet relay chat) and intended to appeal to everyone. Well, everyone who speaks Chinese…
and wants to attack a website!

a direct line to IMDDOS customer service

To learn a little more, we contacted Damballa's customer service team. Communicating with the team is done through QQ chat, a hugely popular chat service in China. We immediately received a polite response. Our contact (for simplicity's sake, let's say "he") introduced himself as a member of the customer service team.

no questions about our identity or our plans

Not once were we asked about who we were, where we were, or why we were interested in this kind of service. When you head to the gas station with a gas can, no one asks you if it’s for your lawnmower or if it’s to burn your neighbor’s house down. Well, the same goes here! Just remember to stop at the checkout counter on your way out.

service levels

Three service levels are offered, including one free “trial” version: prices are based on the length of your service subscription. We were quoted the following prices:

- € 35 for 1 month
- € 130 for one year
- € 200 for permanent “unlimited” access

The free trial version lasts seven days.

classic attack features

Each pack lets you launch a variety of different attacks, including classic attacks like SYN Flood, TCP/UDP Flood, RAW packets, and other more obscure kinds, such as DK and NB.

An **online dashboard is also available** to help you manage your attack by tracking its progress. You just have to install a bit of software on your computer to configure the attack (target IP address, port number, etc).

attack strength: how full is your henhouse?

As for the rate of the attack (bits per second or packets per second), the customer service guy told us it would depend on how many chickens we have. At first, we thought, “we didn’t quite hear you right.”

But yes, you did read correctly: they use the term “chicken” to refer to each machine used to launch an attack. And, we have to admit, it sounds a heck of a lot friendlier than “zombies” and “bots.” These guys didn’t forget to do their marketing research!

If you sign up for one of their packs, then you automatically receive 2,000 of these chickens to launch your attack. As for the free trial, you’ll have to settle for only 100 chickens. Anyway, a 2,000-machine botnet might seem relatively small, but remember, it is still enough to do a lot of damage. The botnet I mapped out was composed of 2,300 machines and I assure you they pulled their weight and then some.

complete with farming tools

What’s really “original” about this service is that, in addition to the guaranteed number of “chickens,” its standard features also include everything you need to put together **your own network of machines!** And indeed, among the products delivered are:

- a tool for creating a virus (called a “small Trojan”) to get “chickens”
- a tool (updated every week) to make your small Trojan invisible to antiviruses
- another tool to hide the virus in a file (jpg, txt, etc.)

With these tools, all you need to do is send out your virus to as many people as possible to infect their machines and create a regular poultry farm. All in all, it’s a good turnkey kit for building and managing your own botnet from home.

Chinese websites excluded from the service contract

Our contact let us know that we could attack any website we wanted, except for Chinese government websites and some other Chinese sites. Are we witnessing an underlying passive attitude by local authorities? Or well-placed “alliances”? **You be the judge.**
Another possibility: companies that sign up for a special protection service are put on a list of sites that cannot be attacked.

no references, but advice available

Among other questions, we asked the agent how many “chickens” were needed to launch an attack.
He told us 100 “chickens” were enough for a small website, 1,000 for a medium-sized site, and 10,000 for a really big site.

When we asked for “references” (from past attacks), he said no. **You don’t mess around with the privacy of commercial exchanges**. ;-)

“companies that sign up for a special protection service are put on a list of sites that cannot be attacked”

warning

the last word

The DDoS market is undergoing massive changes: after years on the black market, it now seems to be moving, at least in part, to a point where **commerce makes the law**. It seems like a new chapter has begun with the arrival of IMDDOS.

Fortunately, commercial offers are still rather scarce at the moment, but they could increase very quickly over the next few years. The cards are changing hands and anyone involved in Web security will have to **keep track of the game and play along!**

the blog post online
http://oran.ge/YW5TR7

DDoS attacks over spam

by Vincent Maurin

In early May 2011, the Czech Republic hosted the "Security and Protection of Information" conference, featuring an opening speech on "Denial-of-service (DoS) attacks using white horse systems: new proof-of-concept DoS against the domain name system (DNS) servers".

The presentation covered **a new attack method that denies service on DNS servers using spam campaigns**.

appetizer: DNS toast

This tasty bite is standard because attackers modify the DNS zone with a domain name in order to register the target server as a name server (e.g., foobar.com NS 1.2.3.4). Any Simple Mail Transfer Protocol (SMTP) server attempting to send an email to the domain name has to collect all necessary information (as outlined by Request for Comments 2821 by asking the responsible DNS server (in this case, 1.2.3.4).

first course: spam salad

The first course of attacks features typical techniques that send spam using the address(es) belonging to a domain previously declared on the DNS server (e.g., user.123@foobar.com, user.456@foobar.com). It is also possible to use sub-domain names (e.g., user.789@dummy.foobar.com).

Note: researchers have identified more than 14,000 unique IP addresses (apparently issued by the same botnet) for spam operations.

main course: white horse steak

The chef's secret here is using what researchers call "**white horses**" which are servers with high bandwidths. Big SMTP server hosts such as Yahoo Mail, Microsoft and Google are favorite targets.

Since each one has the strength to hit hard, these servers can effectively fight spam attacks, although many cases require an analysis of the sender's legitimacy. In compliance with the RFC in force, SMTP servers verify all information associated with the domain name of each message (in this case, a "MX RR" or "A RR" DNS request for the foobar.com zone).

Researchers note that **it's possible to use a botnet of 50,000 machines**, with each machine sending messages to 100 different white horse systems.

good to know

The trick lies in using what researchers call "White Horses," which are servers equipped with high bandwidth capacity.

dessert: DNS medley with distributed DoS sauce

To perform their task, SMTP servers directly or indirectly ask the responsible DNS server about the domain name, **directing a large volume of** "blacklisting domains is not very effective, because an attacker can use several domain names during an attack"

warning

DNS requests to the target server (in this case, 1.2.3.4).

As the server does not necessarily have the capacity to handle these requests, it becomes a **victim of service denial**. As a reminder, servers requesting information are among the white horses that have access to significant resources.

Working with the figures listed above, we quickly arrive at 50,000 machines x 100 white horses x 1 message = **5 million messages or MX RR requests** to the targeted DNS server.

anti-indigestion remedies

Protecting DNS zones is obviously the first line of defense. However, nothing can guarantee that a cyber-squatted zone will not choose your DNS server as a target.

Blacklisting the domains in question is also an easy form of protection, but it is not very effective. Attackers are free to use several domain names during operations (thus adapting their spam campaigns to blacklists).

Only protection mechanisms on the white horse side will be effective, such as anti-spam protection and setting up rules to limit DNS request streams.

Another solution is **anti-distributed DoS protection**, which can be set up with or before your infrastructure.

the blog post online
http://oran.ge/XN9AXL

Second Life: players unknowingly recruited for a DDoS attack?

by Jean-François Audenard

A group of Second Life players were taken for a ride as part of a digital revenge plot. Without their consent or knowledge, their machines were used to launch a series of denial of service attacks against another website.

To open Second Life, players have to install some software on their computers: in addition to the viewers distributed by Linden Lab, the company that developed Second Life, players also have the option of using viewers developed by third parties.

It was precisely through one of these third-party clients that the problem originated.

the Emerald Viewer client

The alternative client in question, called Emerald Viewer, is especially popular among the Second Life crowd. It's developed by Modular Systems, which seems somehow tied to Linden Lab.

According to information available on The Alphaville Herald, a website dedicated to analyzing behaviors in virtual worlds like Second Life, everything started with a suspected leak of personal data concerning players who use Emerald Viewer.

Following these allegations, an individual from Modular Systems decided to change the homepage displayed when its users logged on to Emerald Viewer so as to generate a huge volume of HTTP requests aimed at the person who issued the charges.

modified HTML code attack

According to screenshots (here and here) published in this article on The Alphaville Herald, it's pretty clear that the homepage HTML code was altered so that every player who logged on **automatically generated 32 web requests** to the site http://iheartanime.com/. In three days, nearly 16 million requests were generated by the altered code.

Obviously, everything was hidden from the players, who thus became the unknowing agents of this attack. It was all hidden using standard HTML techniques (iframes tag integrated with an invisible div tag of 1 pixel by 1 pixel).

good to know

You should consider setting up a monitoring system for all of your web pages.

vague explanations from Modular Systems

The explanations offered by Modular Systems concerning these events are rather vague. All we know is that one of the main developers (known as Fractured Crystal) admitted, on the Modular Systems blog, to being the author of the homepage alterations. He decided to step off the project and handed it over to others. He also said that he never intended to create a DDoS attack.

But for me, **this story smells fishy**.

> "any company whose website receives a lot of hits can be an attractive target"

warning

what lessons should we take away from this story?

Any company whose website receives a lot of hits can be an attractive target. If attackers are able to alter page code, they will have a natural amplification system at their disposal.

Of course, attackers can alter page code for other purposes, such as changing content or sending out attack codes or any other attack method of choice. You should therefore consider **setting up a monitoring system** for all of your web pages. Some online services allow for remote monitoring, or you can easily do it using a little script.

how can you detect an attack like this one?

The first detection mechanism is rather simple: **an abnormally high number of requests** is a good sign that something is up. But false alarms are always possible.

A more reliable system would be to analyze the web server logs to detect too many hits coming from one and the same referrer.

the blog post online
http://oran.ge/XNa7ch

cloud computing: weapon of choice for DDoS?

by Jean-François Audenard

Is the Cloud a weapon of choice for launching dedicated denial of service (DDoS) attacks?
Denial of service attacks are seeing a **huge boom**.
One flagrant example of this is the recent WikiLeaks news.
But what does cloud computing have to do with the resources used to launch DDoS attacks? Will the Cloud make botnets and hacktivist groups just two more things of the past?

Here's how I would answer these questions:

botnets and cyber-hacktivists

In a denial of service attack, the goal of the attacker is to render a website completely nonresponsive and unusable for its typical users. One of the keys to a DDoS attack is the bot (or zombie). Most of the time this is a **machine infected with a virus** that is **remotely controlled by the attacker**, all while the real owner of the machine remains completely oblivious.

Less frequently, a group of hackers who share a common goal work together to attack a target. Online protests like this are generally known as cyber-hacktivism. The Anonymous group's assault on banking Websites in connection with the Wikileaks affair is just one recent example.

a network of infected machines isn’t ideal

In the case of zombies/bots, infecting a sufficient number of machines in a short timeframe remains fairly **complex:** you have to set traps and make sure people will fall for them to infect their machines. This takes skill, tools, and enough time to recruit a large enough pool of machines.

In any case, you will have to stay **under the radar** so as not to be detected by an antivirus or another detection system. In addition, network access points are typically asymmetrical, the upload bandwidth available for each bot is slow, etc. In short, you’ll have to recruit a lot of machines.

the limits of cyber-hacktivism

Recruiting a large number of people with the same goals is a formidable attack strategy. When workstation defenses are inefficient (owners agree to infect their own machines or use attack tools), it is possible to rally enough people around a common cause.

But the main problem is finding that common cause: it is difficult if not impossible to motivate people to attack a competitor’s website or extort funds. In addition, the person in charge of the attack remains **dependent on the cooperation and availability of the recruits**. Launching repeated attacks over several days or weeks is unrealistic since motivation will quickly wear off.

creating a botnet in the Cloud

The Cloud has several really attractive features:

- access to **high-performance** networks with a lot of available bandwidth
- ability to quickly and **remotely activate** new resources
- **payment per use** and systems spread over different parts of the globe

In addition to these, we should also add another important feature: the commitment of cloud service suppliers to respect and protect the privacy of customer data in the Cloud against any attacks.

Anyone looking to put together a network of zombie machines could thus use the Cloud to their advantage.

anonymity thanks to “all remote” and “all automatic”

Using information from one or more previously **stolen debit cards** (or prepaid cards), attackers can open accounts with several cloud platforms provided by different suppliers.

Of course, any attacker will want to use proxies and other **anonymous networks**, such as Tor, so as to not leave a trail when creating a cloud account or interacting with a cloud management platform.

good to know

It’s up to security professionals, service providers, and service operators to make sure the benefits of the Cloud outweigh the drawbacks.

> "anyone looking to put together a network of zombie machines can use the Cloud to their advantage"

warning

VMs preset for attack

In each Cloud, attackers can deploy an ISO image (or a virtual machine, ie VM, image) set up to connect immediately upon startup to one or more meeting points to receive marching orders. Meeting points can either be machines located in the Cloud or social network platforms like Twitter.

VMs built for this purpose will of course be optimized so as to use very little live memory or disk space: this way the attacker can streamline costs and launch even more VMs with the extra cash.

elastic botnets using remote API

The cool thing is that more and more clouds are configuring application programming interfaces (API) so you can manage them remotely. In a perfect world (this isn't the case yet but it's the trend), APIs would be standardized.

For an attacker, this is perfect. They can develop a script to **remotely start or stop as many VMs as they need** to launch their attack. Did someone say flexibility and remote management?

huge bandwidth

All these machines are on **high-performance platforms** and have powerful network connections with high bandwidth. What's more, they are located all over the world. So as far as network access goes, there's a strong chance they will hit the nail hard on the head.

impunity ensured by privacy protection?

Like I said above, cloud service providers have to take every measure available to make sure that all data they receive remains securely in the Cloud and that no one can access it without the owner's permission. This rule of course applies to everyone, including the cloud provider's administrators.

So it's **impossible to check out a suspicious VM**. Which is great for attackers: their VMs will stay safely tucked away inside a cloud.

next time: security techniques

Yes, clouds can actually work to encourage denial of service attacks. And I didn't even touch on all the possible scenarios: compromising vulnerable VMs, using PaaS platforms, etc.

It's up to security professionals, service providers, and service operators to make sure the benefits of the cloud outweigh the drawbacks.

In future posts, I'll talk about what kinds of measures cloud computing providers can take to combat the "botnet-ification" of their Cloud. Stay tuned!

the blog post online
http://oran.ge/TQMUlS

about the authors

Jean-François Audenard

Within Orange Business Services, I'm in charge of securing our cloud computing solutions and services. I'm the passionate kind and only look at things this way: no 50/50 for me, I'm an engaged and engaging blogger, I like to go off the beaten track. Sincerity is my tone and optimism and voluntarism my two engines.

Vincent Maurin

I work for Orange Business Services as a security leader within Products and Services Development. My previous jobs as a technical "worker bee" lead me to pay specific attention to the difficulties of implementing companies' security strategies and policies. Security, efficiency and pragmatism are my main pillars.

Our blog :
http://blogs.orange-business.com/connecting-technology/

Document available for download at :
http://knowledge-center.orange-business.com/

Edited by Orange Business Services
20.12.2012

France Télécom – 78, rue Olivier de Serres, 75015 Paris, France – SA au capital de 10 595 541 532 – 380 129 866 RCS Paris – Document non contractuel – juin 2012

Business
Services

orange™