Bakeca.it DDoS

Alessio L.R. Pennasilico

mayhem@alba.st

JUST 4MEETING

Information Security Meeting

Workshops and Training Sessions

25th to 27th.June.2010

Bakeca.it DDoS

How evil forces have been defeated

Alessio L.R. Pennasilico

mayhem@alba.st

JUST 4MEETING

Information Security Meeting

Workshops and Training Sessions

25th to 27th.June.2010

$ whois mayhem

Security Evangelist @ 

$ whois mayhem

Security Evangelist @ 

Board of Directors:

AIP, AIPSI/ISSA, CLUSIT, Italian Linux Society, LUGVR,
Metro Olografix, OpenBeer, Sikurezza.org, Spippolatori

$ whois mayhem

Security Evangelist @ 

Board of Directors:

AIP, AIPSI/ISSA, CLUSIT, Italian Linux Society, LUGVR,
Metro Olografix, OpenBeer, Sikurezza.org, Spippolatori

CrISTAL, Hacker's Profiling Project, Recursiva.org

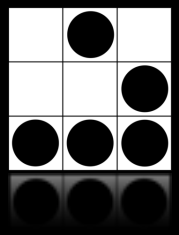

Background

May 9th 2008

I received a phone call…

We have a problem!

Our Goal

To allow people to express themselves!

We want to allow people to exchange ideas and needing, in the simpler and faster way.

Like writing a note on a school dashboard.

We work for the ideas, about work, about private life, about cultures and exchange them between the people of the same city.

Some numbers

180.000 visitors per day

5.000.000 pages per day

45 cities

About 90 employees

On and Off line marketing activities

The problem

Someone is attacking the Bakeca.it WEB farm

The infrastructure

100 Mb/s bandwidth co-located in a Milan ISP webfarm

1 Cisco PIX 525 Firewall

2 Linux Application Load Balancers

About 15 frontend WEB servers

1 Database server as backend

The current situation

High load inbound traffic is hitting the firewall

(*about 100 MB/s*)

The hardware is unable to handle all incoming packets and drops too many connections

Statistics

Before the attack

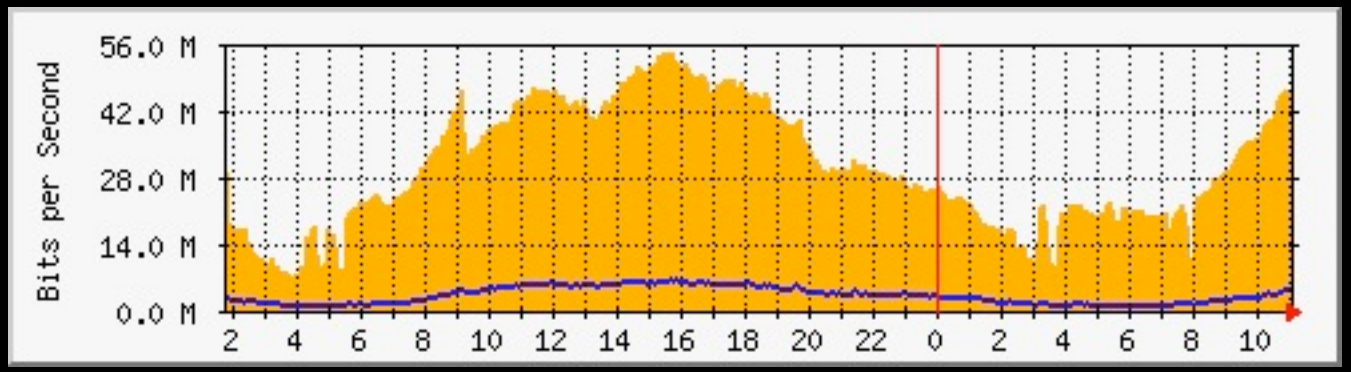

One of the first attacks!

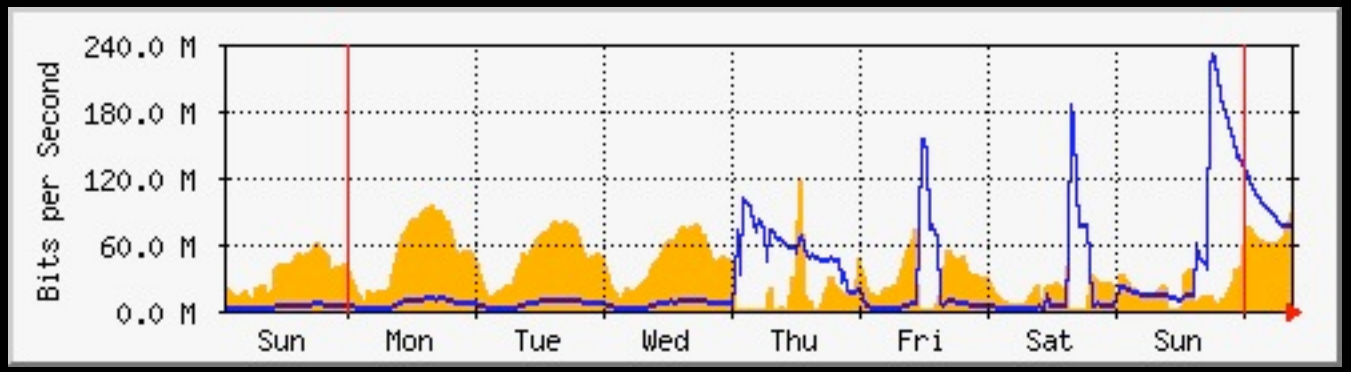

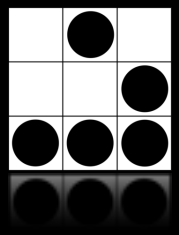

DDoS

DDOS

A distributed denial-of-service attack is an attempt to make a computer resource unavailable to its intended users. Although the means to carry out, motives for, and targets of a DoS attack may vary, it generally consists of the concerted, malevolent efforts of a person to prevent an Internet site from functioning efficiently or at all. Perpetrators of DoS attacks typically target sites or services hosted on high-profile web servers.

Exaust resources

CPU

RAM

Disk Space

Bandwidth

Countermeasures

In a private environment you can easily set quotas about resource usage on your user

but what about Internet connected hosts?

DDoS How-To

Own as many hosts as you can

Make them join your network, to rule them

Tell them what to do, all together!

It's about Money

Owning a botnet can be very remunerative:

you can rent it

you can sell services

(DDoS, SPAM, Phishing, ...)

DDoS for Dummies

Pay Russian Business Network

DDOS Cost: $300 for 24 hours

Month long prices available, no need to plan ahead. Also available for $50 per hour

http://www.washingtonpost.com/wp-dyn/content/article/2007/10/12/AR2007101202461.html?nav=rss_technology

http://www.birmingham-infragard.org/meetings/talks/presentations/DDOS.in.Practice.pdf

Targets

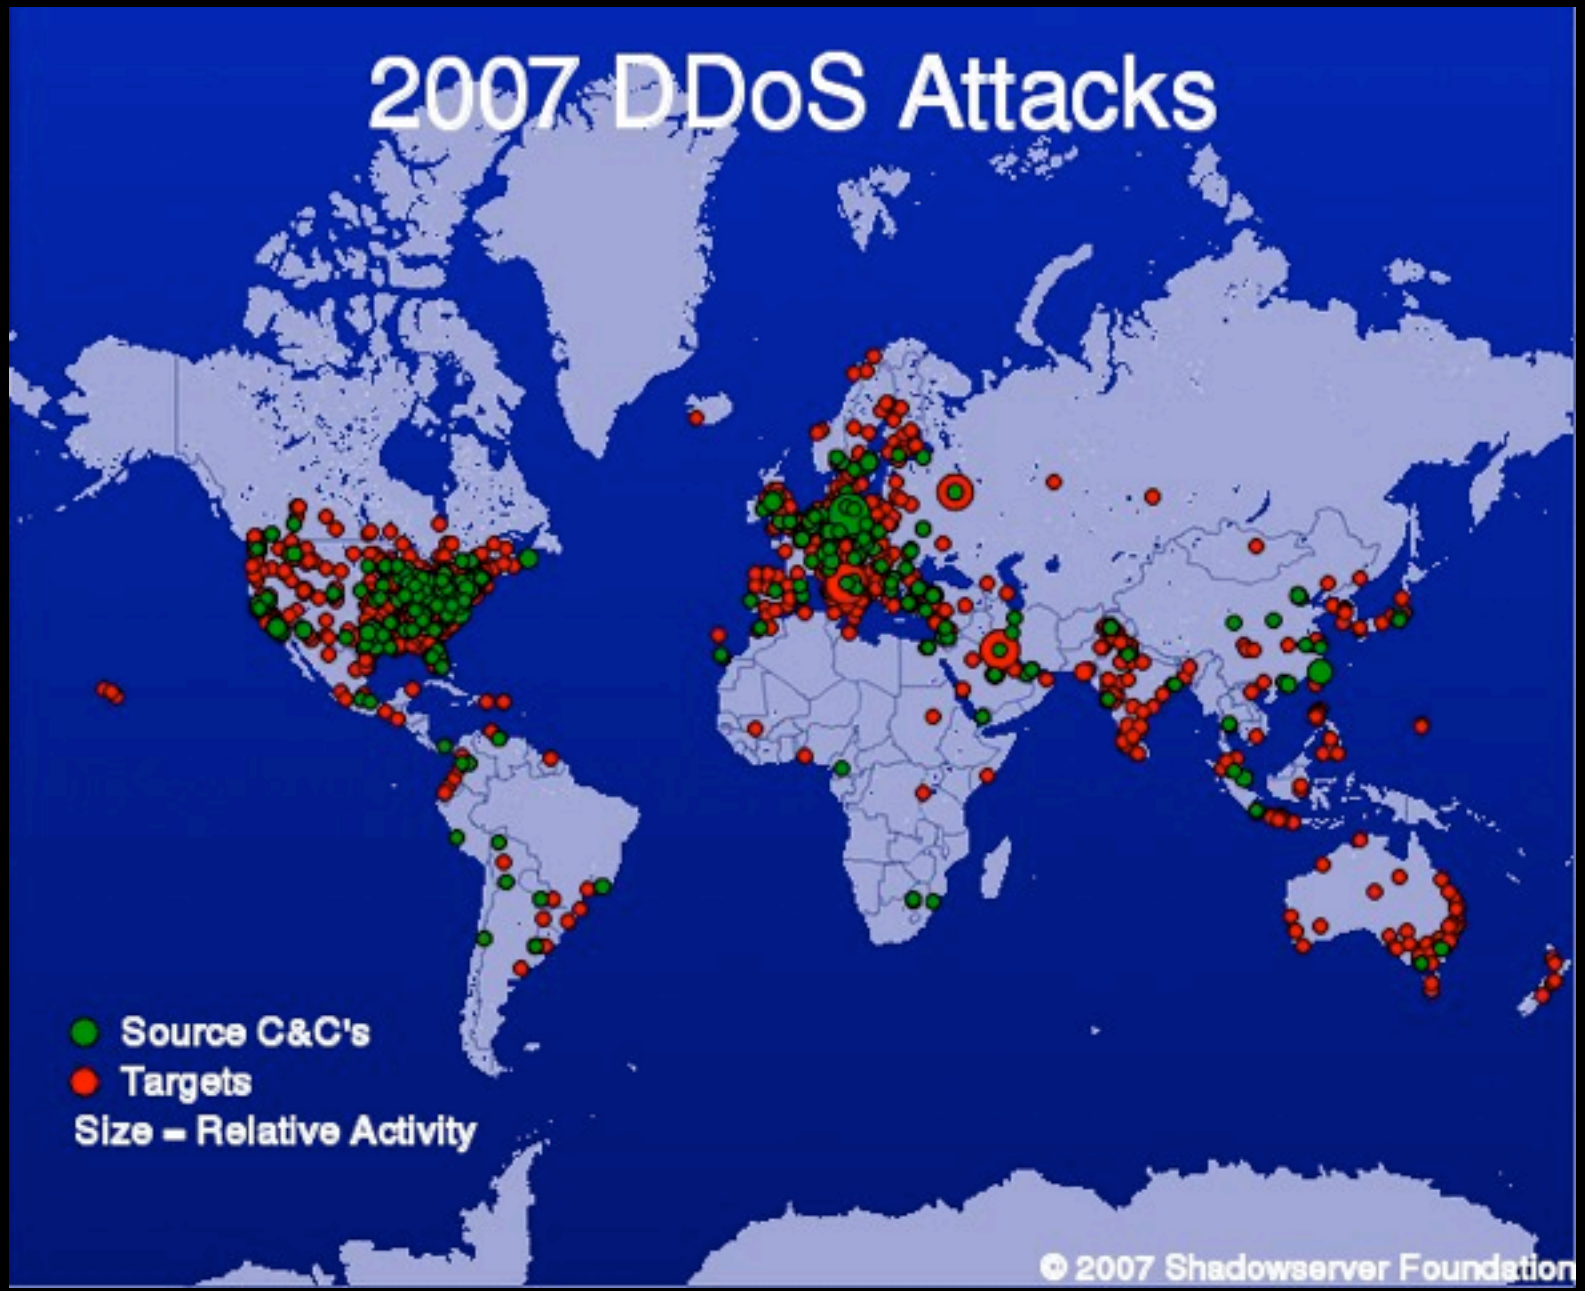
2007 DDoS Attacks
Source C&C's
Targets
Size = Relative Activity
© 2007 Shadowserver Foundation

Graphical representation

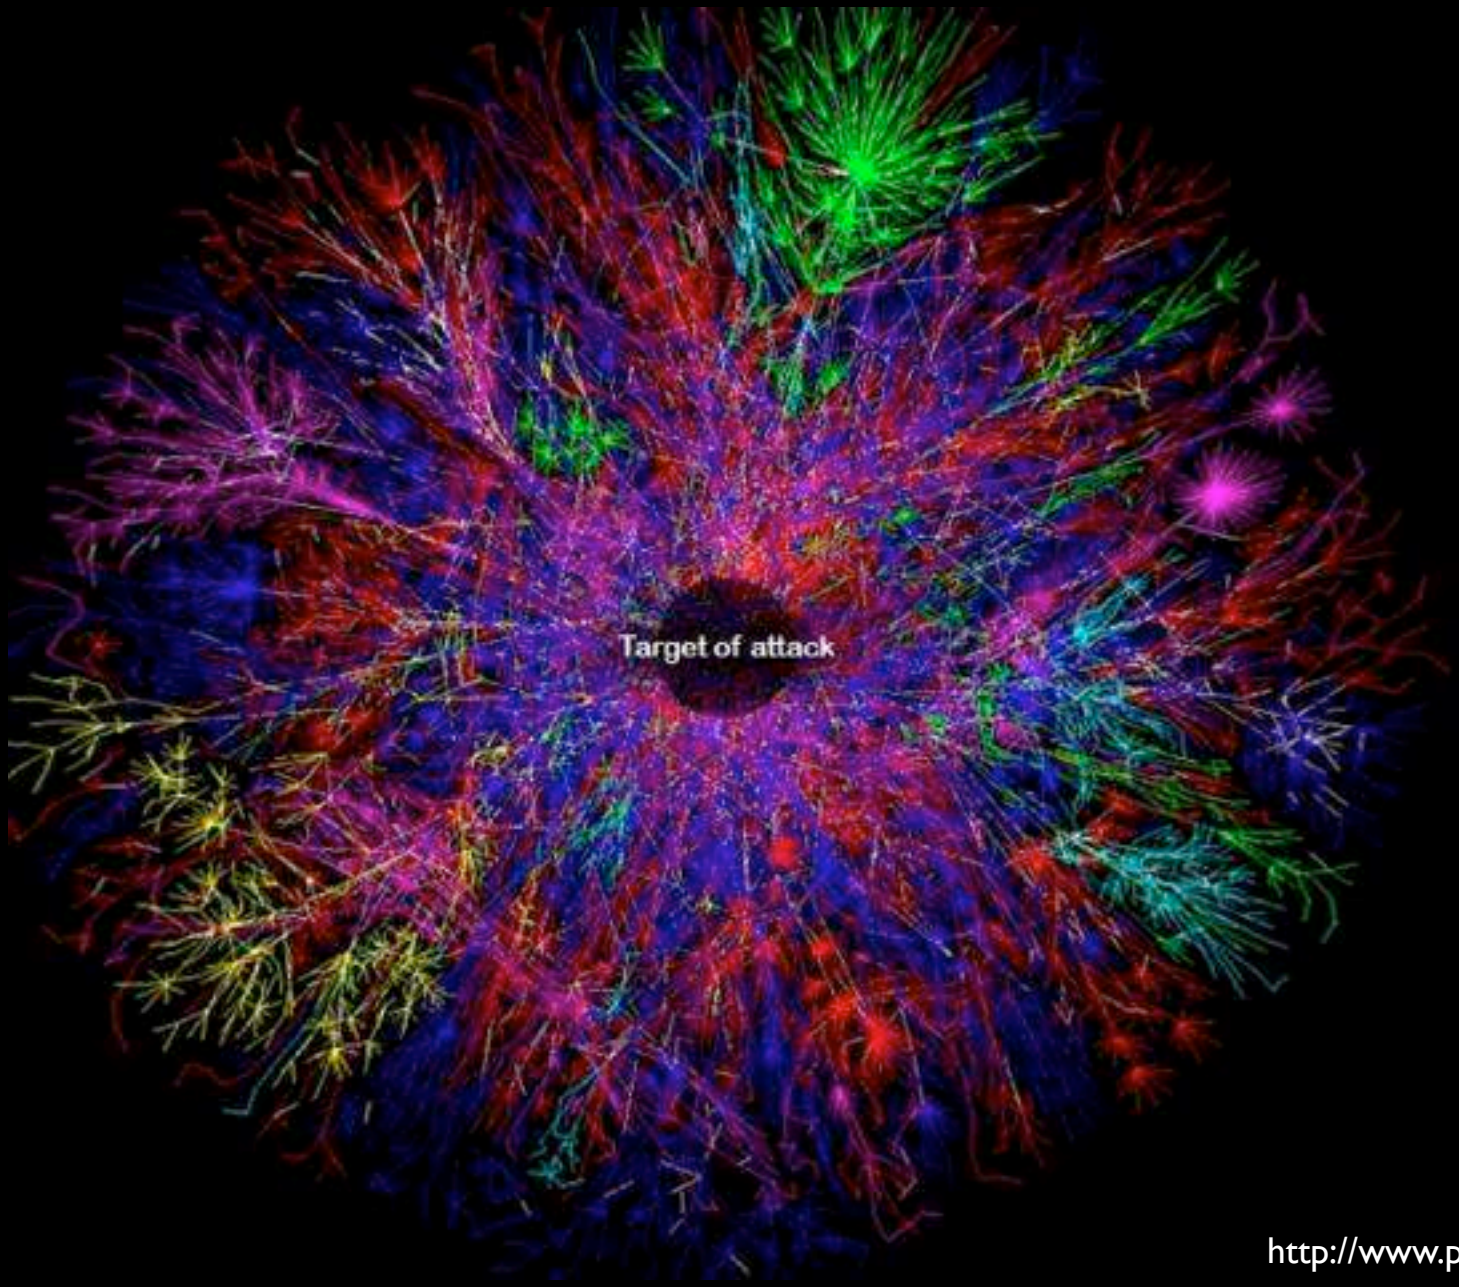

http://www.prolexic.com/zr/

Victims

2000: Amazon, Yahoo, CNN, eTrade

2002: Root DNS server

2007: Estonia

2008: Bakeca.it

It's not about Hackers!

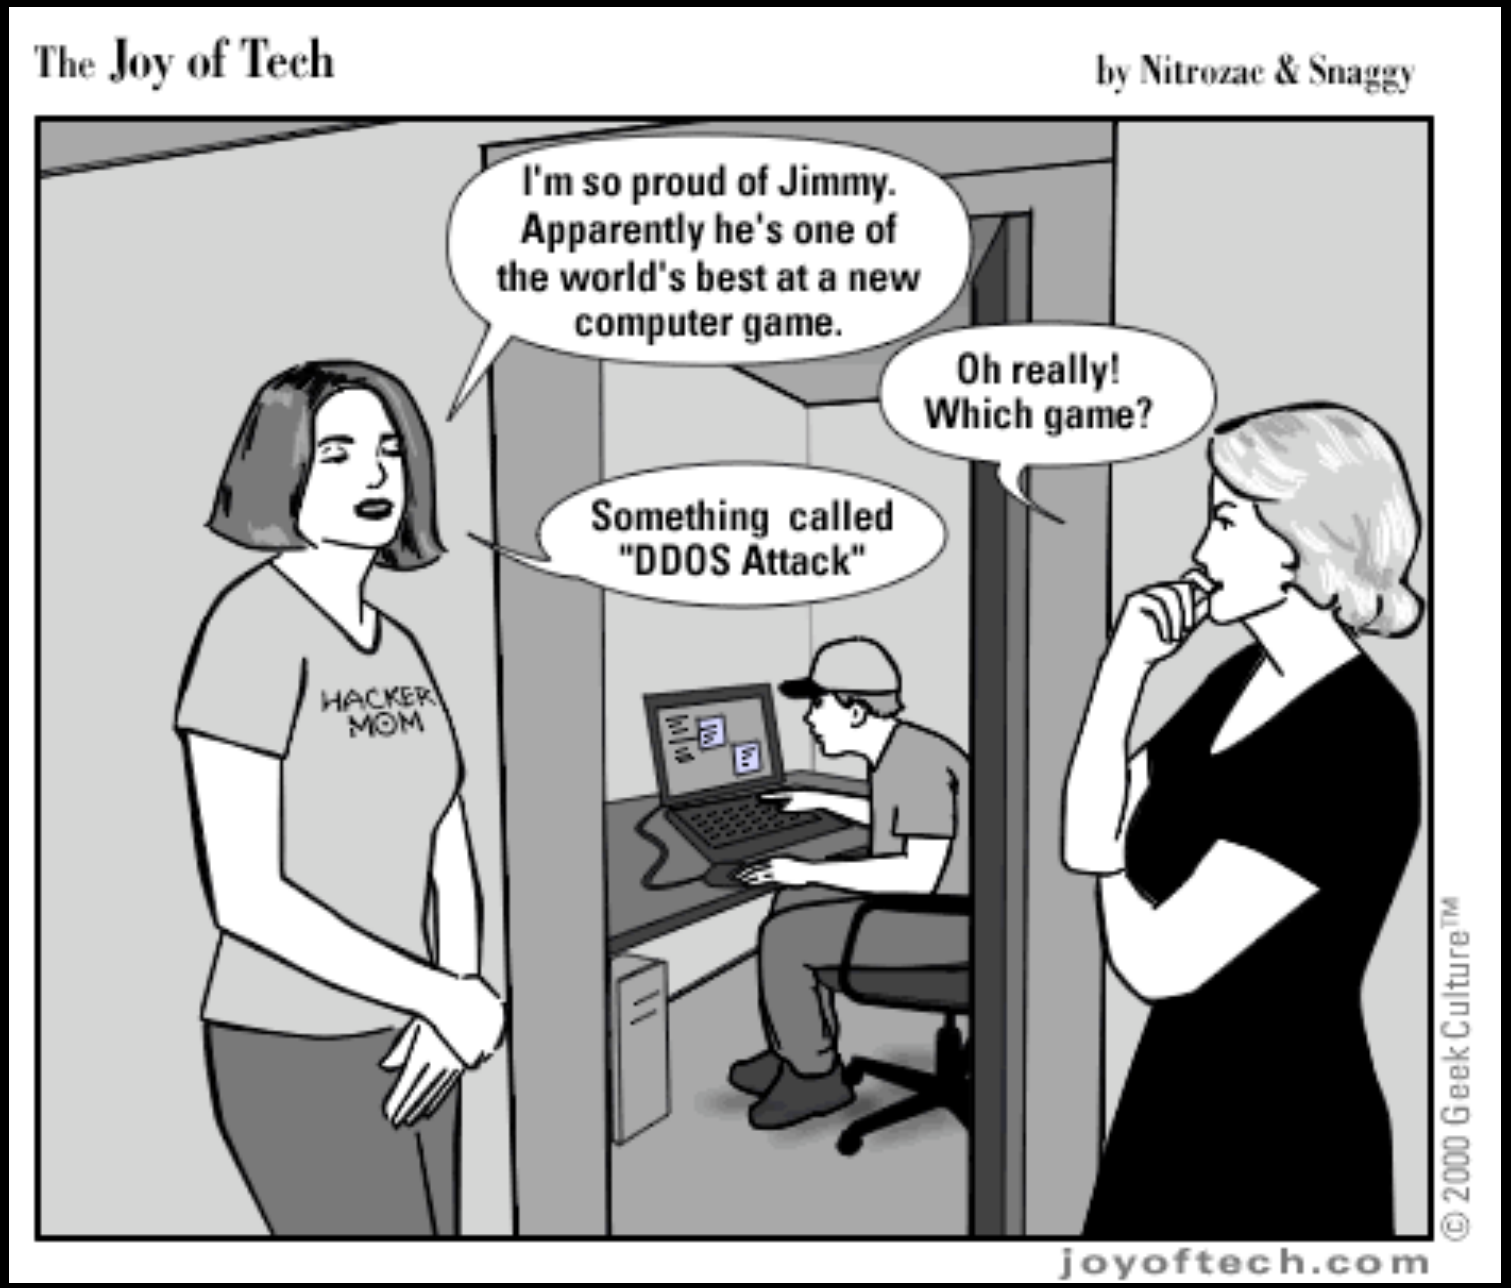

Managing an attack

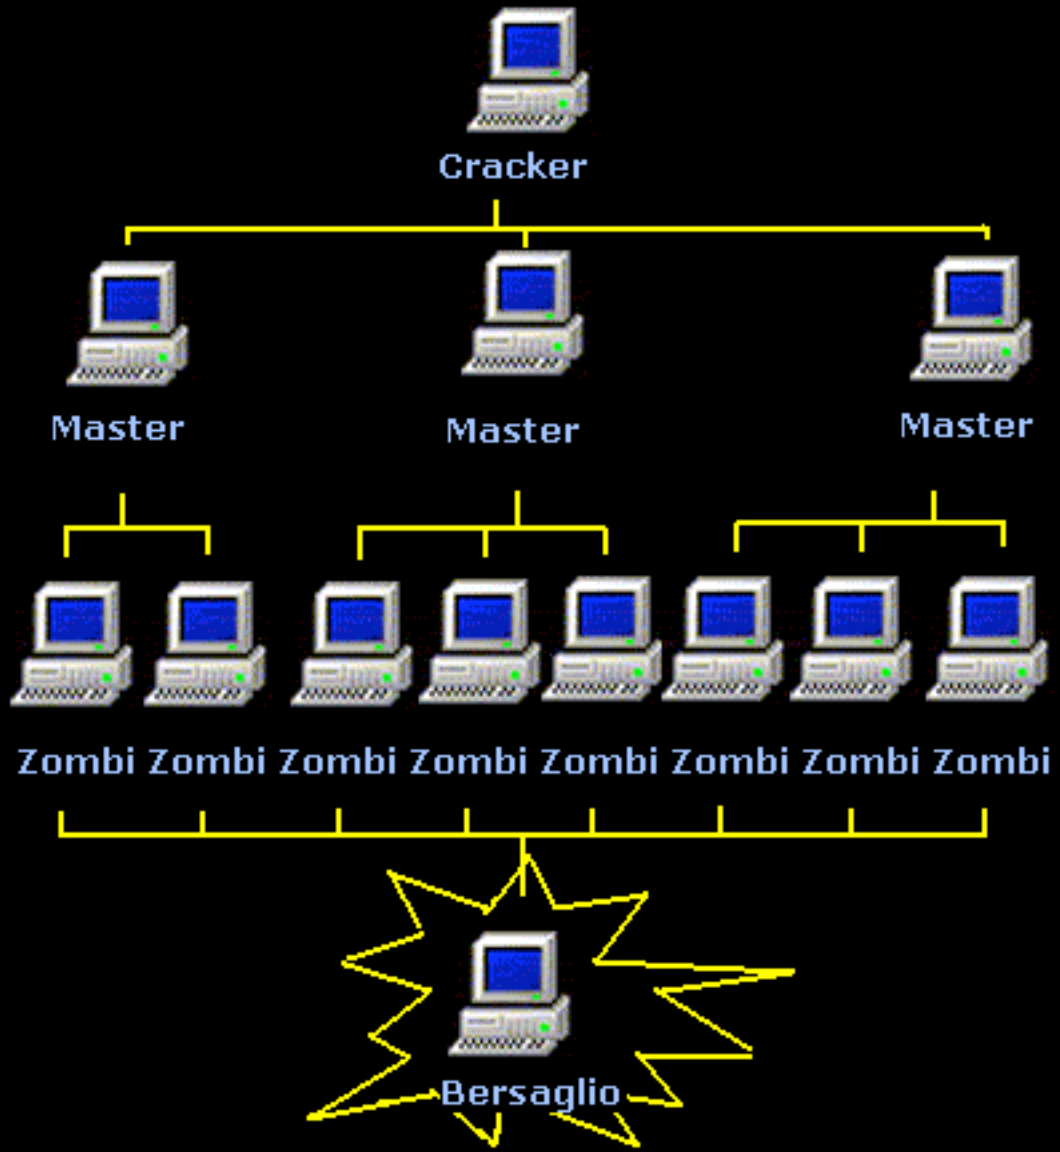

Spot the attacker

It's really difficult because of

the command and conquer strategy

It's difficult to spot the real attacker machine

It's difficult to build a list of the attacking hosts

Difficult to mitigate

Cannot use blacklists, too many dynamic hosts

There's no main attack player, every host manages a very small part of the attack

It's always very easy to cut-off real users :(

DDoS Techniques

ICMP Amplifications - SMURF

DNS Amplifications

Traffic Flood

Service Congestion

NetStrike

Empty connection Flood

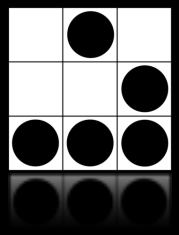

The Attack

SYN Flood

The traffic aggregated about 100 Mb/s of TCP SYN flagged packets

We were in charge of mitigating it

We tried to filter out embryonic connections

Three way handshake

Client – SYN=1 -> Server

Client <- SYN=1,ACK=1 – Server

Client – ACK=1 -> Server

Meanwhile...

I was giving a lecture

at Camerino's Univeristy

Discussing about the problem with OpenBeer friends, we had an **idea**...

Changing technology

The PIX was not able to handle all those packets

We decided to use an **OpenBSD** server as the firewall

We enabled the **PF** SYN-Proxy feature

Null Route

Yeah, we know… black-holing some AS would've been simpler and faster…

However, the customer wasn't in charge about the routing. He doesn't own his AS and the ISP would had not allowed him to request such settings on their routers…

Manage everything with PF

For this reason we continued to implement **OpenBSD** features to mitigate any further attack...

Bingo!

The new firewall was able to handle

over 100 Mb/s of SYN flood

The whole infrastructure was up and running again in the “right OpenBSD”™ :) way

PF SYN-Proxy configuration

```
pass in on $outside proto tcp from any \
                to $balancers port 80 synproxy state
```

Saturday

FabioFVZ, OpenBeer founder, returned back in Venice

mayhem went to Florence to give a speech at e-privacy conference

The attacker

He didn't appreciate our new filtering techniques and hacks :)

For this reason he started using some more resources...

Bang!

The ISP upgraded our bandwidth to 200 Mb/S

OpenBSD was managing about 100 Mb/s of TCP SYN flood

Then the SYN flood bandwidth started growing up … and growing up …

First limit

At 185 Mb/s the OpenBSD console was unresponsive

The IRQ rate was too high

No traffic was routed towards the balancers

The international issue

First instance: the ISP temporary filtered out all the international connections to our infrastructure

This caused some users to be filtered, but the bandwidth used was drastically reduced

(*about 90 Mb/s of total traffic*)

Idea

The problem was too complex

We tried to split it in simpler parts

Clustering

We put a second firewall to manage the traffic

No PF-Sync, no CARP were implemented

This was to improve performances and reduce packets to manage

Our idea was to create two different, independent, fast systems, both able to handle any traffic by themselves

It was Saturday

No specific hardware was available

No expensive hardware black box available

We were able to use "only" generic x86 hosts, already present in the server farm

10 DELL rack servers were available there to be installed as new WEB servers for the HTTP frontend cluster

First proposal to the ISP

Please route all traffic directed to our infrastructure to those two IP, in Round-Robin

Sorry, it's not possible :(

DNS Balancing

For this reason we decided to reconfigure the DNS A records to point the two IP addresses

In this way the traffic was forwarded with Round Robin algoritm to both firewalls

The states problem

Both firewall were maintaining their own connection state table

New need: all traffic should be routed back to the same firewall that forwarded it

Traffic flow

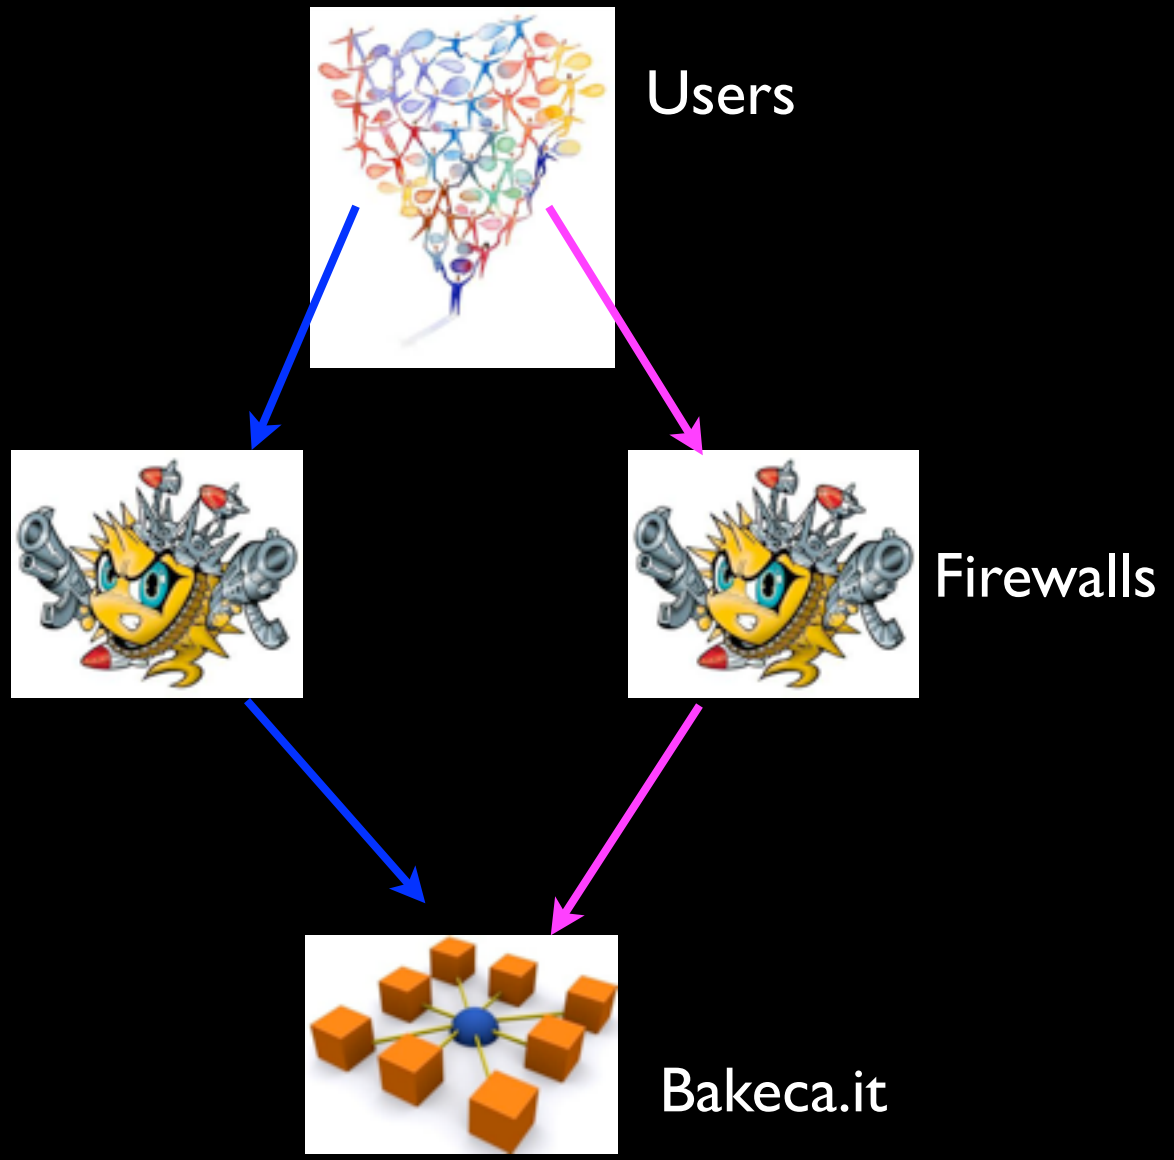

Traffic flow

Asymmetric Routing
Dropped Connection

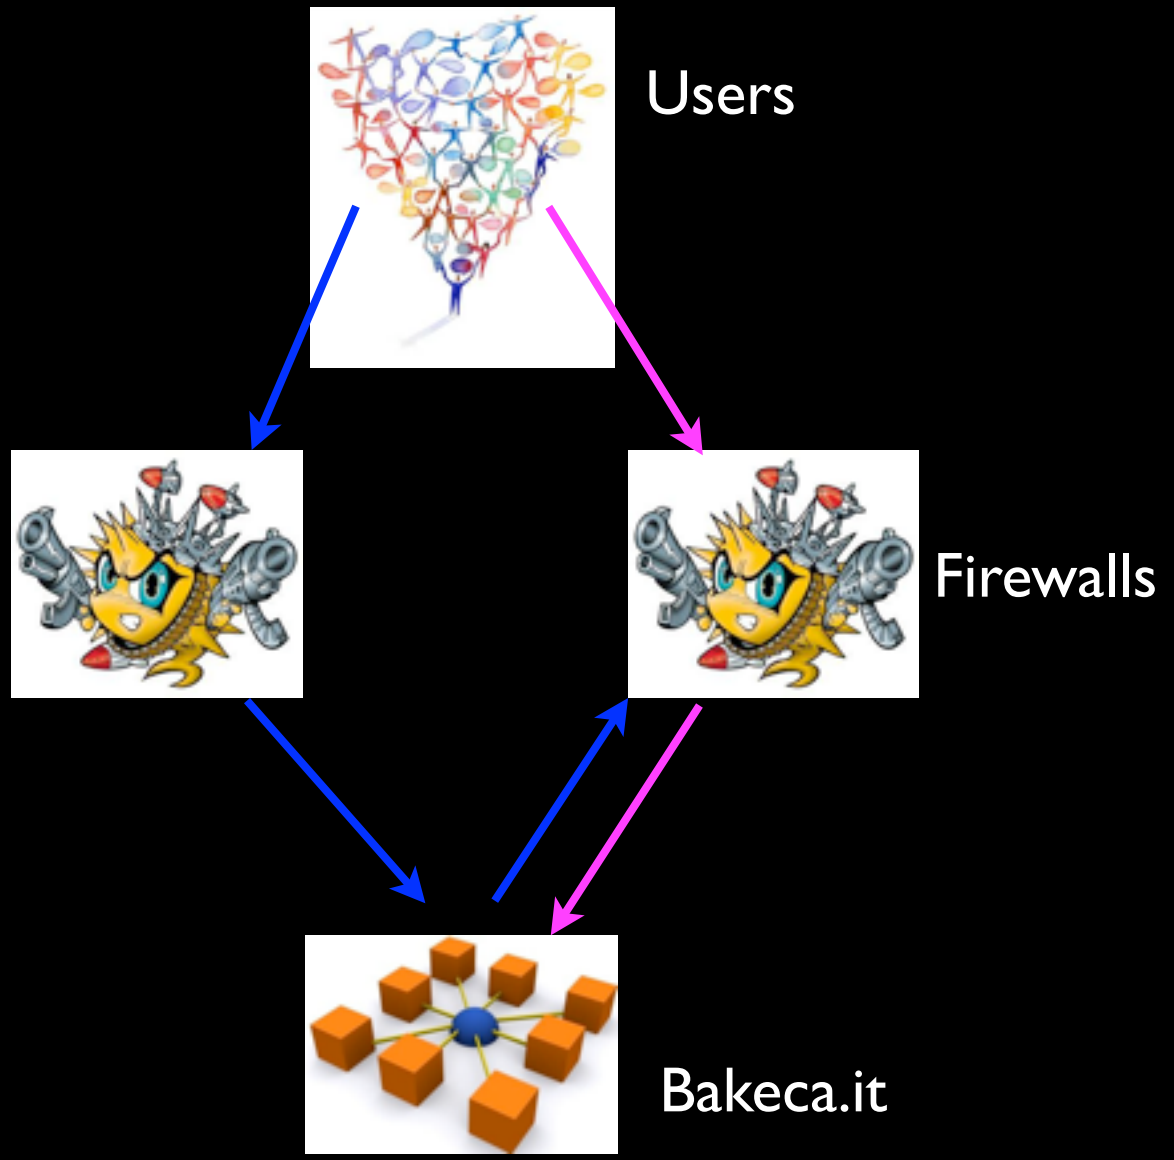

Traffic flow

Asymmetric Routing
Dropped Connection

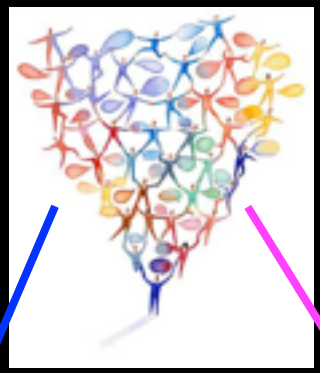

Users

Firewalls

Symmetric Routing
Allowed Connection

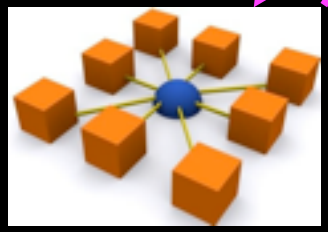

Bakeca.it

Traffic flow

Asymmetric Routing
Dropped Connection

Symmetric Routing
Allowed Connection

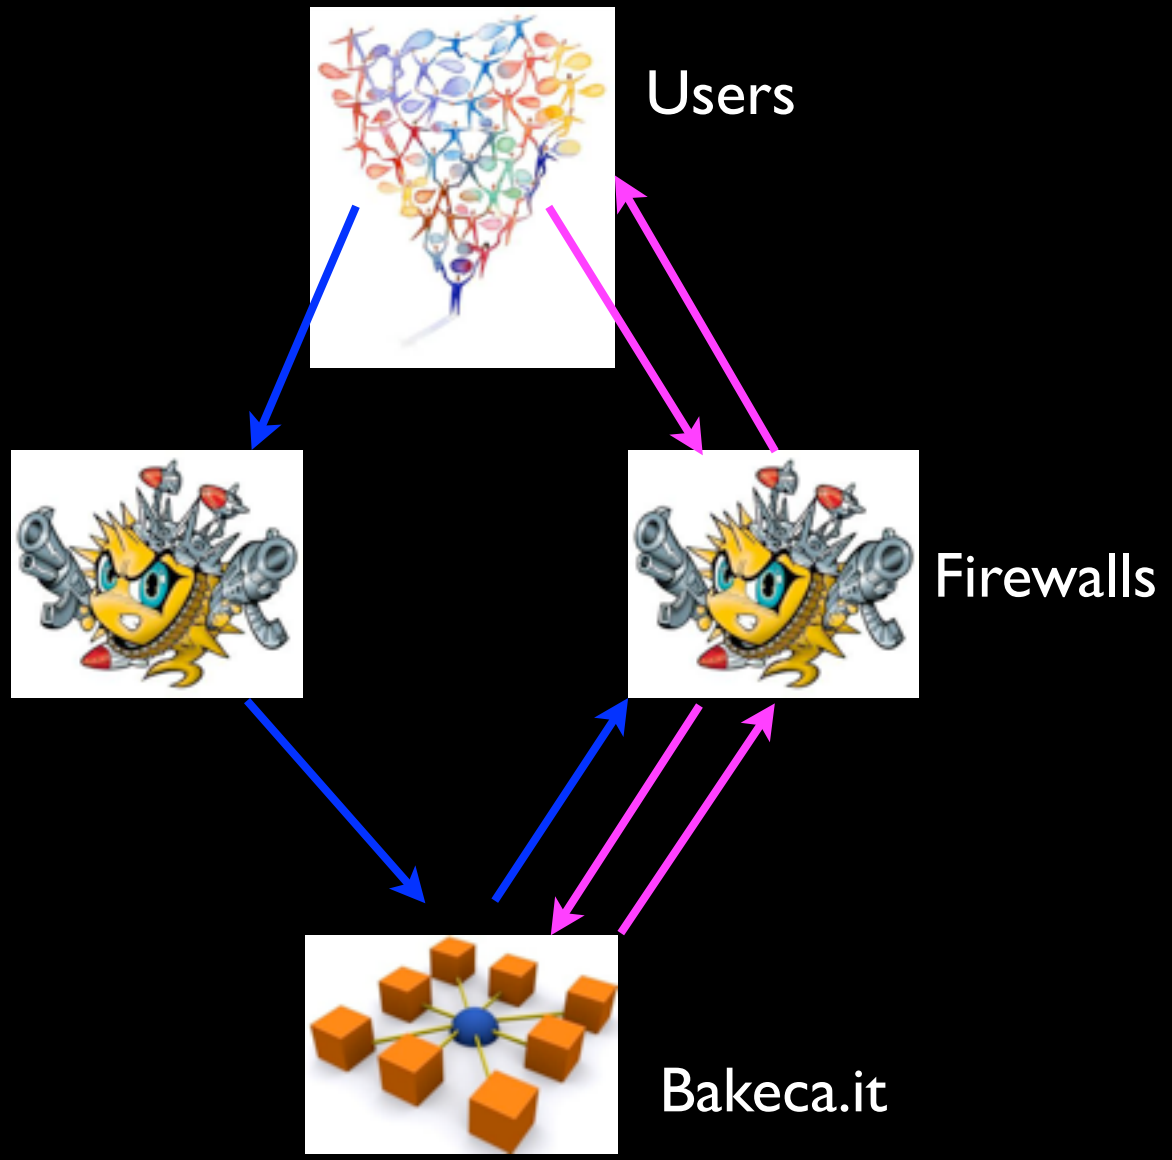

NAT as a solution

We configured **PF** to NAT the incoming traffic towards the load balancers

All traffic appeared to be generated by the private IPs of the firewalls

Traffic flow after NAT

IP traffic with
user IP as SRC IP

Users

Firewalls
with NAT

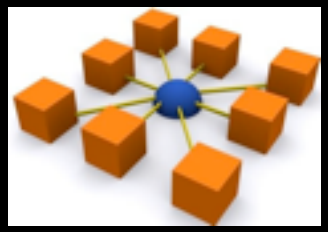

Bakeca.it

Traffic flow after NAT

IP traffic with
user IP as SRC IP

Users

Firewalls
with NAT

IP traffic with
firewall's IP as SRC IP

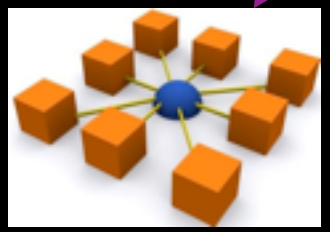

Bakeca.it

Traffic flow after NAT

IP traffic with
user IP as SRC IP

Users

Firewalls
with NAT

IP traffic with
firewall's IP as SRC IP

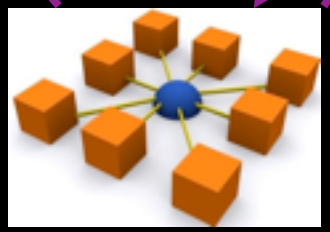

Bakeca.it

Traffic flow after NAT

IP traffic with
user IP as SRC IP

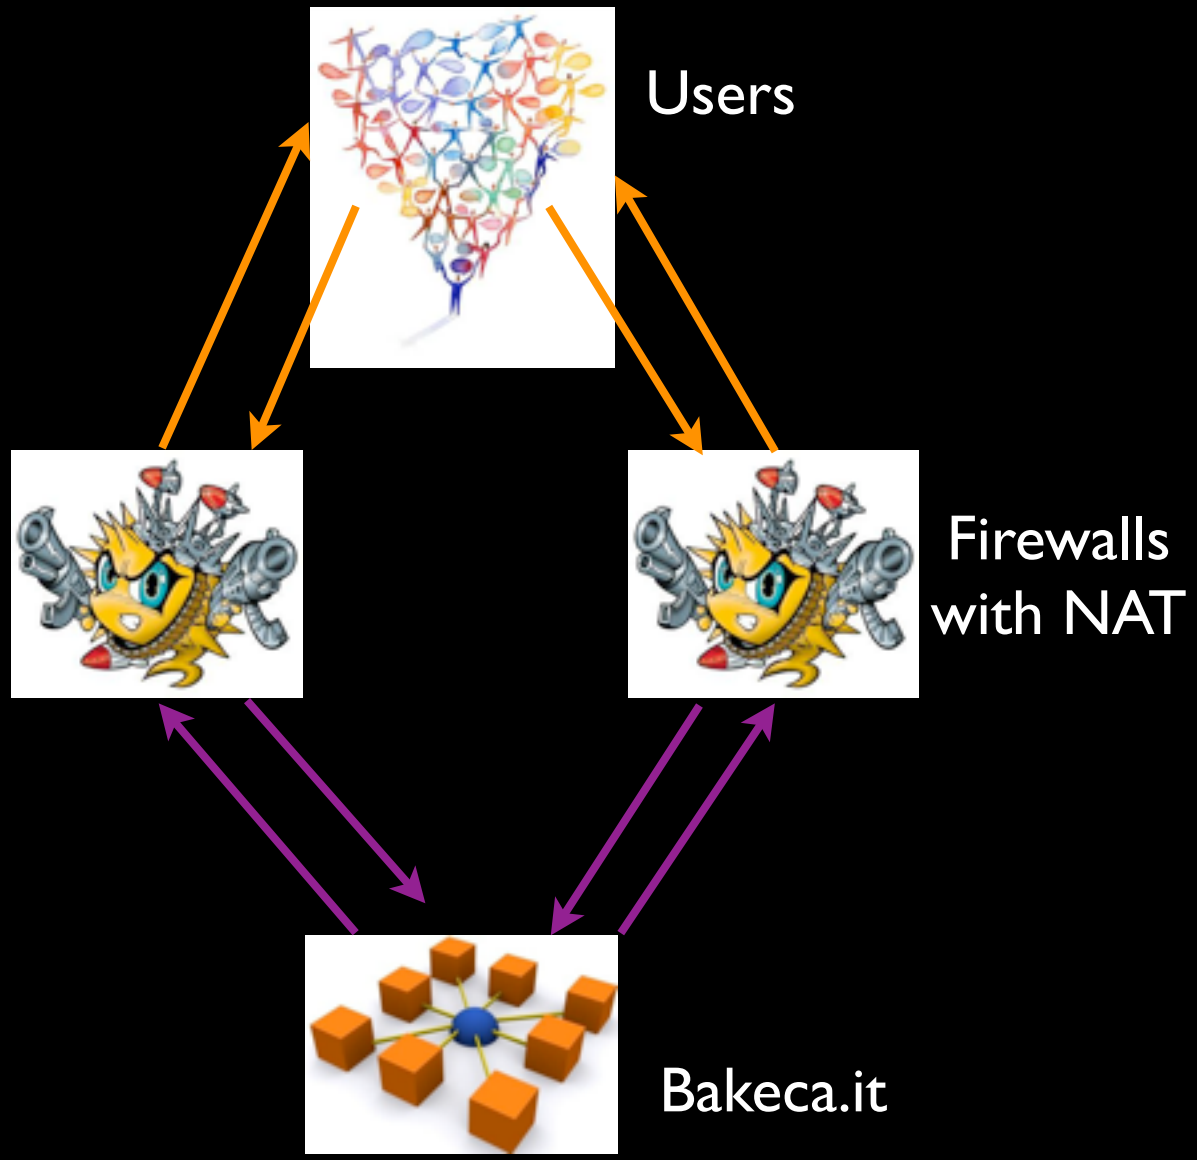

IP traffic with
firewall's IP as SRC IP

Optimizing traffic management

Purpose: Increase total throughput the firewall can handle

kern.maxclusters=128000
net.inet.icmp.errppslimit=1000
net.inet.icmp.errppslimit=1000
net.inet.tcp.rfc1323=1
net.inet.tcp.sack=1
net.inet.ip.ifq.len=0
net.inet.ip.ifq.maxlen=2500
net.inet.tcp.recvspace=262144
net.inet.tcp.sendspace=262144
net.inet.udp.recvspace=262144
net.inet.udp.sendspace=262144

On-line again

The international traffic was enabled again…

Bingo!

Everything were working fine…

The ISP upgraded the available bandwidth to 500 Mb/s

We were managing more than 200 Mb/s of SYN Flood!!

Bang again...

The traffic started raising again, and again…

At about 300 Mb/s of incoming traffic both firewalls were unresponsive...

Replicate, replicate now!

We started a massive deployment of **OpenBSD** Firewall boxes

8 hosts, all configured in the same way

The DNS A records were reconfigured to point at every host in the stack

The ISP upgraded our bandwidth to 1 Gb/s

Standing on our feet!

And the traffic continuously grew up…

and grew up…

We reached 850 Mb/s and the cluster infrastructure was working, the attacker seemed to had finished the bandwidth!

Attack traffic escalation

Attack traffic escalation

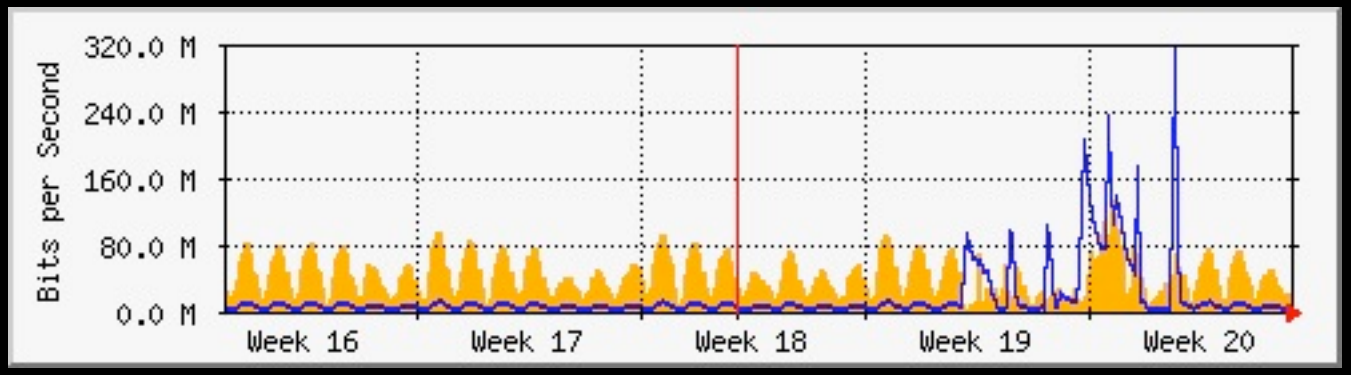

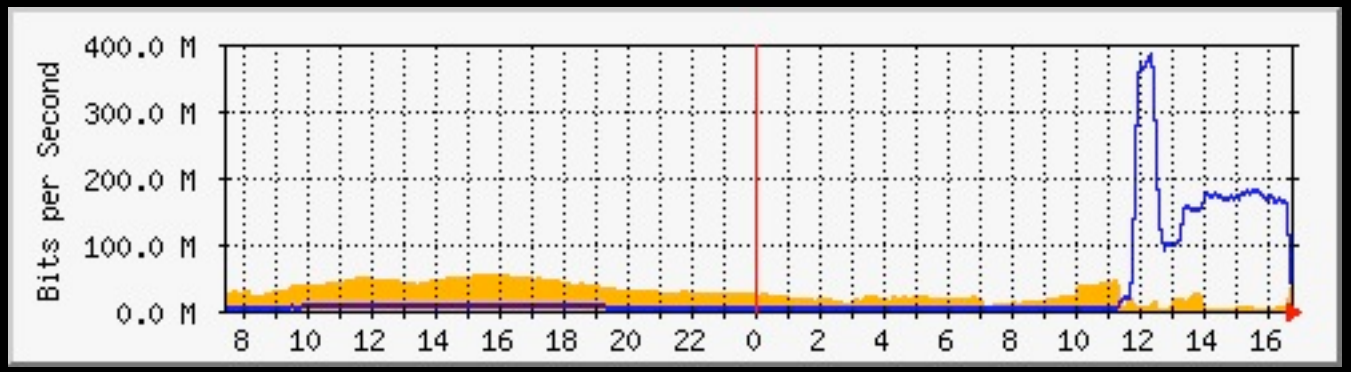

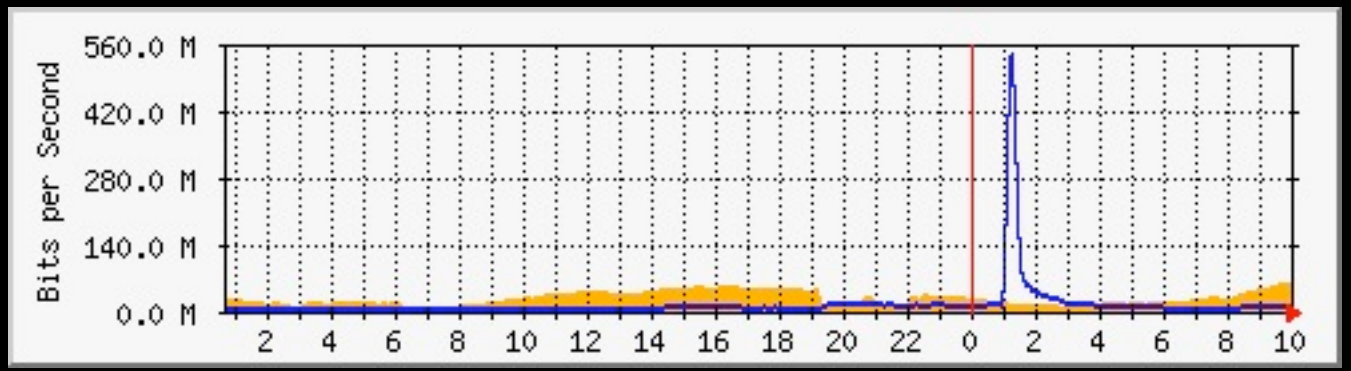

GET Flood

And then the SYN Flood started disappearing

Then a strange activity on the database started…

Everything was slow… and then stopped working again :(

Mitigate it

Rate limiting connections helped us to avoid too many HTTP GET query to reach the load balancers, and everything started working again

PF rate limiting connections

```
http_rate ="(source-track rule, max-src-states 100, \
        max-src-conn-rate 100/60, \
            overload <BLACKLIST> flush global)"

table <BLACKLIST> persist file "/etc/blacklist"

block in quick on $outside from <BLACKLIST>
```

Specific GET Flood

The rate limit allowed only some GET (connections) per second from the same host

Then the GET start being less time-intensive, but most of the requests were directed to the two slower and more CPU/IO-intensive pages of the public sites (*Rent on Milan*)

Keep in mind:

We were managing traffic from about 20.000 hosts, plus the normal hosts we were used to manage before the attack

We need time!

Our engineers at **EasyBit** asked for some more time while engineering an algoritm to mitigate the attack…

It was during the week-end

It was two weeks that we were working 24/7!

Traffic laundry

The customer decided to invest some money

They stipulated a contract with some external companies: they asked us to point our DNS on their filters

We would have back only the clean traffic

Worst than before

We tried two companies

Both promised, none maintained

No traffic, or too much, was arriving

Worst than before

We tried two companies

Both promised, none maintained

No traffic, or too much, was arriving

So they started talking about A.I. and neural network, more money needing, and some complex setup to do on their side ...

The traffic in the laundry

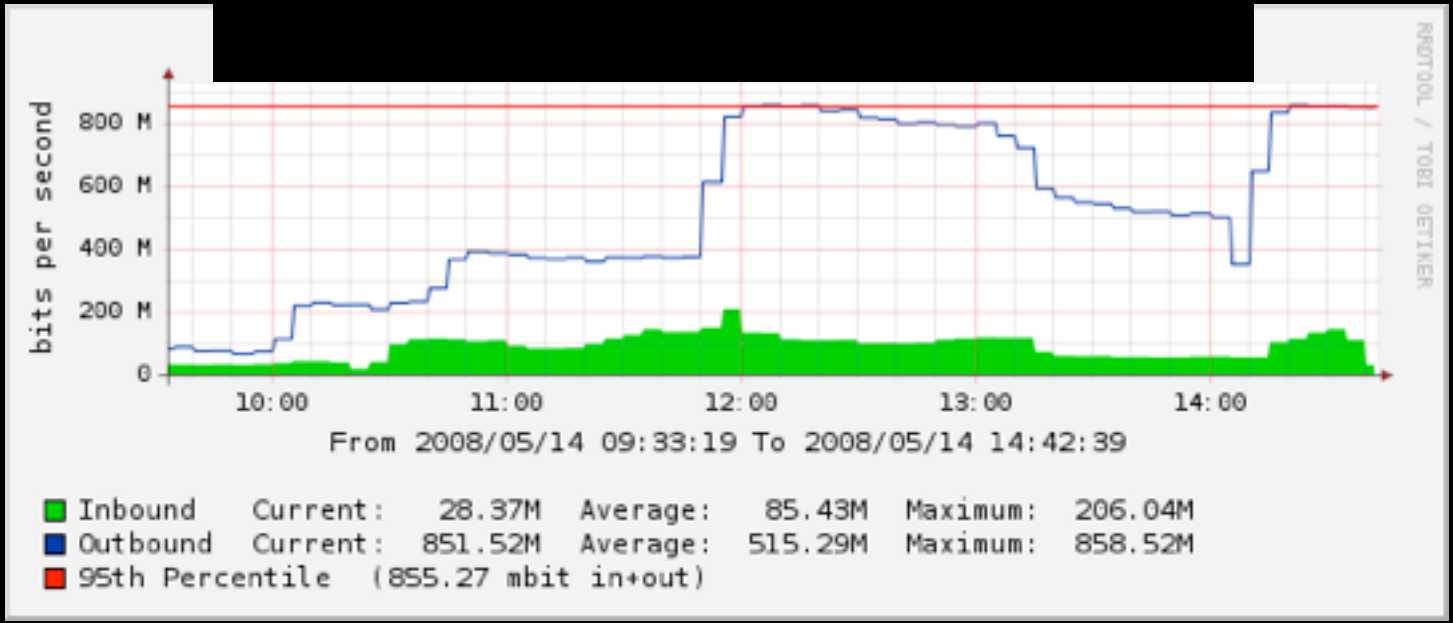

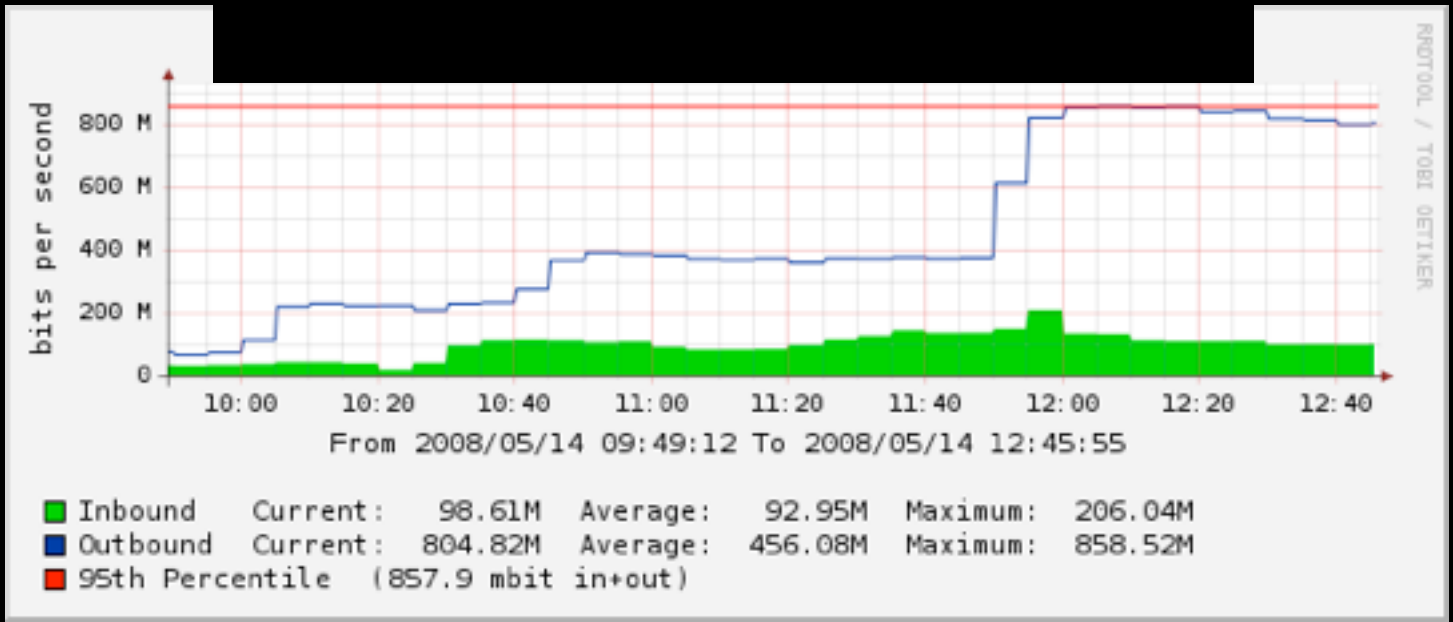

We were faster!

During those dramatic tests EasyBit never stopped working to analyze and implement the algoritm to mitigate the GET flood

It was ready, we took back the traffic, and everything started working again!

The applicative filter

On the Linux load balancers were implemented:

selective HTTP deflector, based on URL and User-Agent

some URL rewriting rules

some GET rate-limiting filters

The Backend

The host managing the database was clustered in two nodes, both replicating and balancing all the queries

This allowed not only to avoid a SPoF, but also helped in mitigating the attack

Sleep needing

Everyone needed some sleep hours

But during night of May 26th...

DNS Flood

The DNS servers were not in the same server farm. They were, temporary, on a secondary network, with slow bandwidth and no **OpenBSD** cluster to protect them…

And the attacker started flooding with random traffic (UDP/ICMP) that network!

Protect the DNS

We moved to the same WEB farm also the DNS server, that started working fine, protected by the **OpenBSD** PF stack!

How to post on Bakeca

You post trough a web form

An e-mail confirms the post

Then you confirm the mail and the post is approved

SMTP Flood

The attacker inserted thousands of new posts

All the e-mails were in the queue of the mail server (many thousands)

Its default gateway was not able to handle all incoming and outgoing traffic

SMTP Relay

Every **OpenBSD** host started using sendmail (8) to relay internal mails to the world

The mail server was using the stack hosts as relay servers in Round-Robin

The queue was empty in a couple of hours

The mediatic campaign

http://web-pulito.seolab.it/

200 support messages in less than 1 month!

We were lucky...

The attack was DNS based

Bakeca is a solid and clever company that invested a lot of money to improve the service

All partners were smart

Scripts

Managing a stack of **OpenBSD** hosts was not a problem anyway

We created some hand-made scripts to modify the same file on every host automagically

(think about *pf.conf*...)

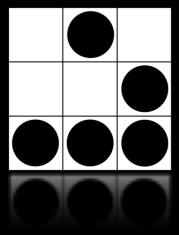

Conclusions

The results

May 30th, 8 **OpenBSD** with PF with capabilities of act as a SYN proxy, connections rate limiting, incoming connections' NAT, relaying mails with sendmail(8)

About 850 Mb/s of traffic, over 20.000 hosts

Anyway...

Anyway...

DDoS are always a nightmare

This was an incredible adventure, very long and hard, but we can now say:

Anyway...

DDoS are always a nightmare

This was an incredible adventure, very long and hard, but we can now say:

the evil forces have been defeated!

Thanks to...

Paolo Geymonat … for trusting us :)

Roberto Emanuele for working so hard

Everyone at Backeca, **SEOLab** and **EasyBit** for supporting us, no matter which hour of day or night was :)

Obviously all the friends of **OpenBSD**

Also thanks to...

All the hackers that listened to all our rants in those days and gave us some precious advices:

Guido "Zen" Bolognesi

Daniele "Cyrax" Martini

People at E-Privacy and LinuxPerSec3

Remeber: don't be evil :)

Remeber: don't be evil :)

Stitch as Emperor Palpatine

Alessio L.R. Pennasilico

mayhem@alba.st

JUST 4MEETING

Information Security Meeting

Workshops and Training Sessions

25th to 27th.June.2010

TORINO

Bakeca.it

Albast

Questions?

Alessio L.R. Pennasilico

mayhem@alba.st

Information Security Meeting

Workshops and Training Sessions

25th to 27th.June.2010

JUST 4MEETING

TORINO

Bakeca.it

Albast

Questions?
Obrigado!

Alessio L.R. Pennasilico

mayhem@alba.st

Information Security Meeting

Workshops and Training Sessions

25th to 27th.June.2010

JUST 4MEETING