

# Passive DNS Replication

Florian Weimer

17<sup>th</sup> Annual FIRST Conference, Singapore, 2005

# Outline

A very brief introduction to DNS

Case Study: Botnet mitigation

Architecture and implementation

Results

# Outline

A very brief introduction to DNS

Case Study: Botnet mitigation

Architecture and implementation

Results

# DNS as a huge table

www.enyo.de	IN	A	212.9.189.164
static.enyo.de	IN	A	212.9.189.164
mail.enyo.de	IN	A	212.9.189.167
enyo.de	IN	MX	10 mail.enyo.de
...			
www.first.org	IN	A	163.1.2.77
www.first.org	IN	A	192.25.206.20
www.first.org	IN	A	210.148.223.8
...			
164.189.9.212.in-addr.arpa	IN	PTR	www.enyo.de

# DNS as a huge table

www.enyo.de	IN	A	212.9.189.164
static.enyo.de	IN	A	212.9.189.164
mail.enyo.de	IN	A	212.9.189.167
enyo.de	IN	MX	10 mail.enyo.de
...			
www.first.org	IN	A	163.1.2.77
www.first.org	IN	A	192.25.206.20
www.first.org	IN	A	210.148.223.8
...			
164.189.9.212.in-addr.arpa	IN	PTR	www.enyo.de

the domain name

# DNS as a huge table

www.enyo.de	IN	A	212.9.189.164
static.enyo.de	IN	A	212.9.189.164
mail.enyo.de	IN	A	212.9.189.167
enyo.de	IN	MX	10 mail.enyo.de
...			
www.first.org	IN	A	163.1.2.77
www.first.org	IN	A	192.25.206.20
www.first.org	IN	A	210.148.223.8
...			
164.189.9.212.in-addr.arpa	IN	PTR	www.enyo.de

the record class (mostly unused)

# DNS as a huge table

www.enyo.de	IN	A	212.9.189.164
static.enyo.de	IN	A	212.9.189.164
mail.enyo.de	IN	A	212.9.189.167
enyo.de	IN	MX	10 mail.enyo.de
...			
www.first.org	IN	A	163.1.2.77
www.first.org	IN	A	192.25.206.20
www.first.org	IN	A	210.148.223.8
...			
164.189.9.212.in-addr.arpa	IN	PTR	www.enyo.de

the record type

## DNS as a huge table

www.enyo.de	IN	A	212.9.189.164
static.enyo.de	IN	A	212.9.189.164
mail.enyo.de	IN	A	212.9.189.167
enyo.de	IN	MX	10 mail.enyo.de
...			
www.first.org	IN	A	163.1.2.77
www.first.org	IN	A	192.25.206.20
www.first.org	IN	A	210.148.223.8
...			
164.189.9.212.in-addr.arpa	IN	PTR	www.enyo.de

the record data (as determined by the record type)



## DNS as a huge table

www.enyo.de	IN	A	212.9.189.164
static.enyo.de	IN	A	212.9.189.164
mail.enyo.de	IN	A	212.9.189.167
enyo.de	IN	MX	10 mail.enyo.de
...			
www.first.org	IN	A	163.1.2.77
www.first.org	IN	A	192.25.206.20
www.first.org	IN	A	210.148.223.8
...			
164.189.9.212.in-addr.arpa	IN	PTR	www.enyo.de

the primary key (the only selection criterion)

# DNS as a huge table

www.enyo.de	IN	A	212.9.189.164
static.enyo.de	IN	A	212.9.189.164
mail.enyo.de	IN	A	212.9.189.167
enyo.de	IN	MX	10 mail.enyo.de
...			
www.first.org	IN	A	163.1.2.77
www.first.org	IN	A	192.25.206.20
www.first.org	IN	A	210.148.223.8
...			
164.189.9.212.in-addr.arpa	IN	PTR	www.enyo.de

query for www.first.org IN A

## DNS as a huge table

www.enyo.de	IN	A	212.9.189.164
static.enyo.de	IN	A	212.9.189.164
mail.enyo.de	IN	A	212.9.189.167
enyo.de	IN	MX	10 mail.enyo.de
...			
www.first.org	IN	A	163.1.2.77
www.first.org	IN	A	192.25.206.20
www.first.org	IN	A	210.148.223.8
...			
164.189.9.212.in-addr.arpa	IN	PTR	www.enyo.de

response to www.first.org IN A

# DNS as a huge table

www.enyo.de	IN	A	212.9.189.164
static.enyo.de	IN	A	212.9.189.164
mail.enyo.de	IN	A	212.9.189.167
enyo.de	IN	MX	10 mail.enyo.de
...			
www.first.org	IN	A	163.1.2.77
www.first.org	IN	A	192.25.206.20
www.first.org	IN	A	210.148.223.8
...			
164.189.9.212.in-addr.arpa	IN	PTR	www.enyo.de

reverse lookup for 212.9.189.164, rewritten to  
164.189.9.212.in-addr.arpa IN PTR

## DNS as a huge table

www.enyo.de	IN	A	212.9.189.164
static.enyo.de	IN	A	212.9.189.164
mail.enyo.de	IN	A	212.9.189.167
enyo.de	IN	MX	10 mail.enyo.de
...			
www.first.org	IN	A	163.1.2.77
www.first.org	IN	A	192.25.206.20
www.first.org	IN	A	210.148.223.8
...			
164.189.9.212.in-addr.arpa	IN	PTR	www.enyo.de

reverse lookup result for 212.9.189.164

## DNS as a huge table

www.enyo.de	IN	A	212.9.189.164
static.enyo.de	IN	A	212.9.189.164
mail.enyo.de	IN	A	212.9.189.167
enyo.de	IN	MX	10 mail.enyo.de
...			
www.first.org	IN	A	163.1.2.77
www.first.org	IN	A	192.25.206.20
www.first.org	IN	A	210.148.223.8
...			
164.189.9.212.in-addr.arpa	IN	PTR	www.enyo.de

reverse lookup result for 212.9.189.164  
does not include static.enyo.de

## DNS summary

- ▶ You can query only by the primary key, the domain/class/type triple.
- ▶ Queries on secondary keys can be emulated if the key is encoded in a domain name (as in `164.189.9.212.in-addr.arpa`).
- ▶ There are no consistency guarantees.
- ▶ Reverse lookups (based on PTR records) are optional and not reliable: Both `www` and `static` point to `212.9.189.164`, but there is only one PTR record.

# Outline

A very brief introduction to DNS

Case Study: Botnet mitigation

Architecture and implementation

Results



# Slowing down botnet propagation

Controller 1

Victim 1

## Slowing down botnet propagation



Controller sends command `".advscan dcom135 100 5 0 -r"`

## Slowing down botnet propagation

Controller 1


Victim 1



Perimeter filter causes victim to drop out of botnet


# Slowing down botnet propagation

Controller 1



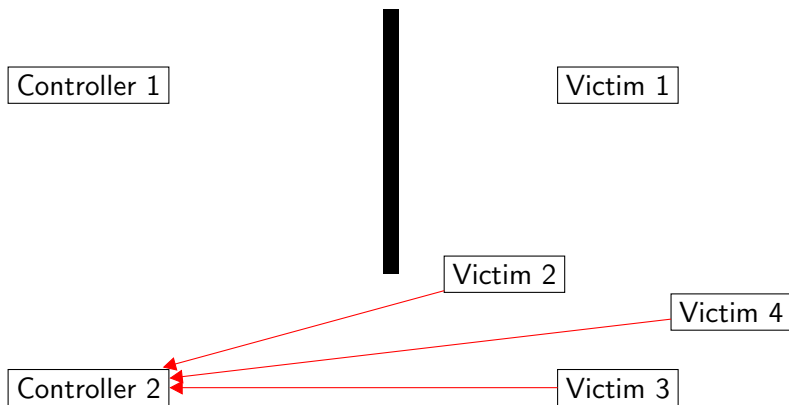
First victim infects more hosts on the internal network

## Slowing down botnet propagation



New victims cannot connect to first controller,  
thanks to the perimeter filter

# Slowing down botnet propagation



But victims may join the botnet on a second controller

# Contacting the botnet controller

- ▶ The bot may contain one or more domain names instead of hard-coded IP addresses.
- ▶ Each domain can resolve to multiple IP addresses.
- ▶ Blocking a single IP address often does not prevent hosts from joining the botnet.
- ▶ If you know the domain name, better filters are possible.
  - ▶ You can adjust the filters when the IP addresses change.
  - ▶ You can filter the domain name on your resolvers (in theory).

# How to recover domain names from IP addresses

- ▶ Reverse engineer the bot.
  - ▶ Disassembling needs time and expertise.
  - ▶ And a copy of the malware.
- ▶ The security team often cannot access the caching resolvers which store a copy of the DNS record.
- ▶ Zone file transfers do not work.
  - ▶ Usually considered a security vulnerability and therefore disabled.
  - ▶ Which zone?



## From domain names to IP addresses

- ▶ Capture DNS packets and look for the IP address you are interested in.
- ▶ DNS caches may delay the reappearance of resource records for hours.
- ▶ Idea: Capture DNS records in advance and store them in a database for later reference.
- ▶ This leads to “passive DNS replication”.

# Outline


A very brief introduction to DNS

Case Study: Botnet mitigation

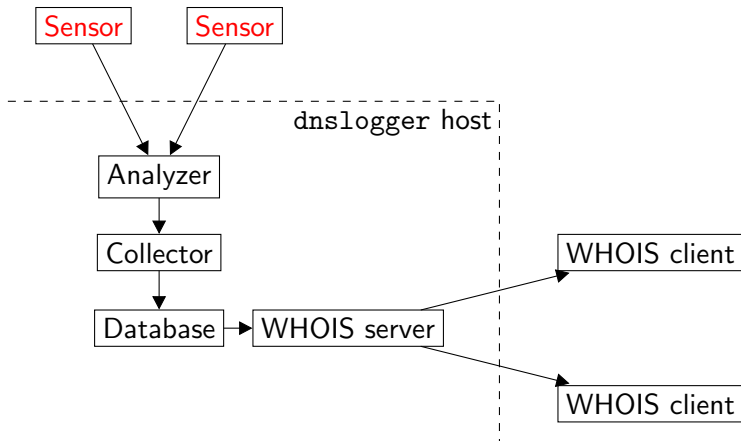
Architecture and implementation

Results

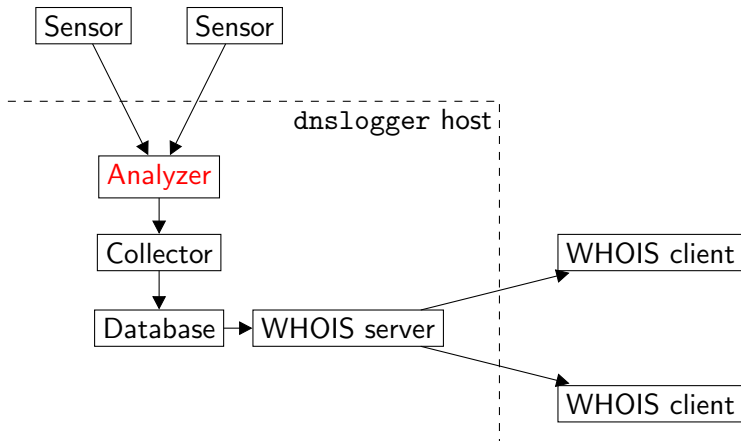
## dnslogger architecture




# dnslogger architecture




## dnslogger architecture



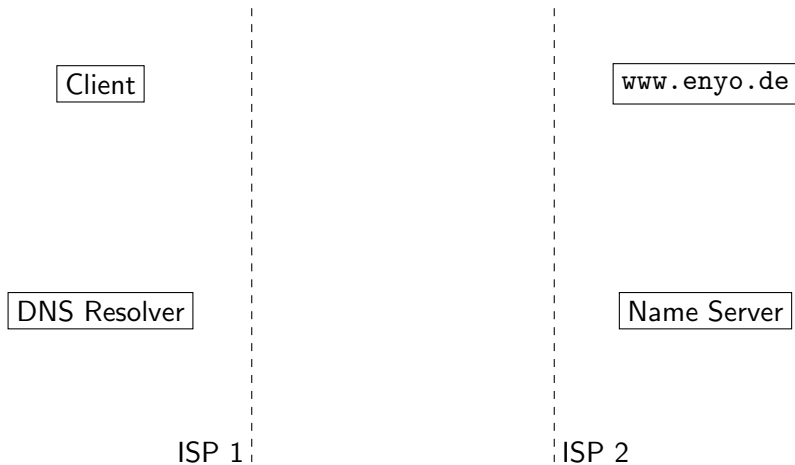
## dnslogger architecture



## dnslogger architecture




# Sensor placement

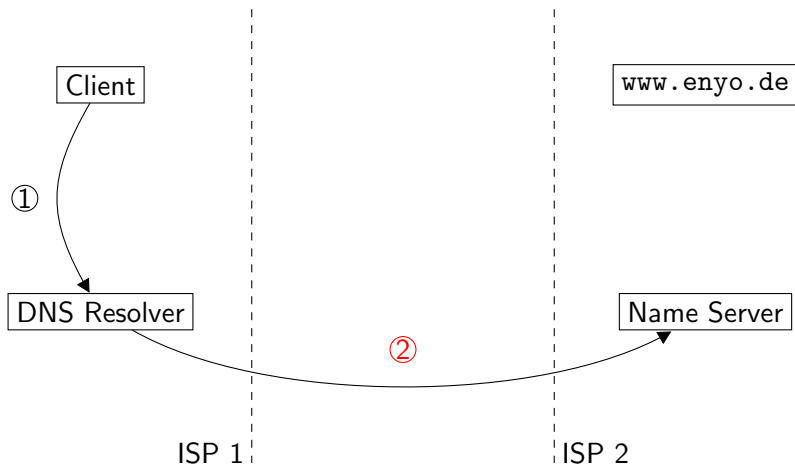




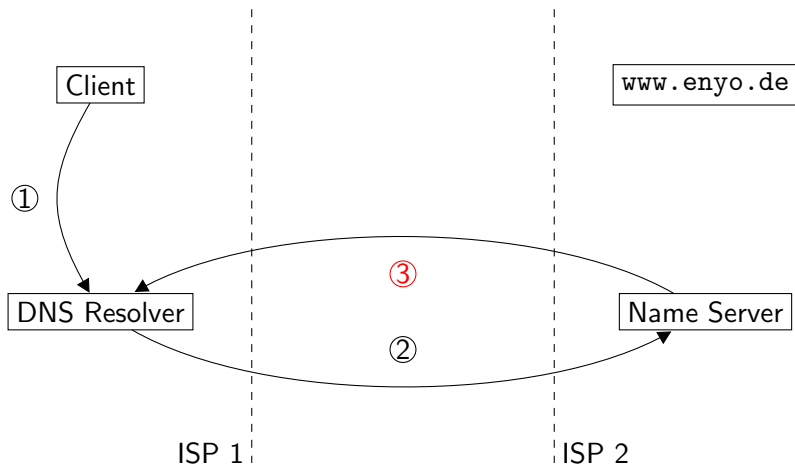
# Sensor placement




# Sensor placement




# Sensor placement



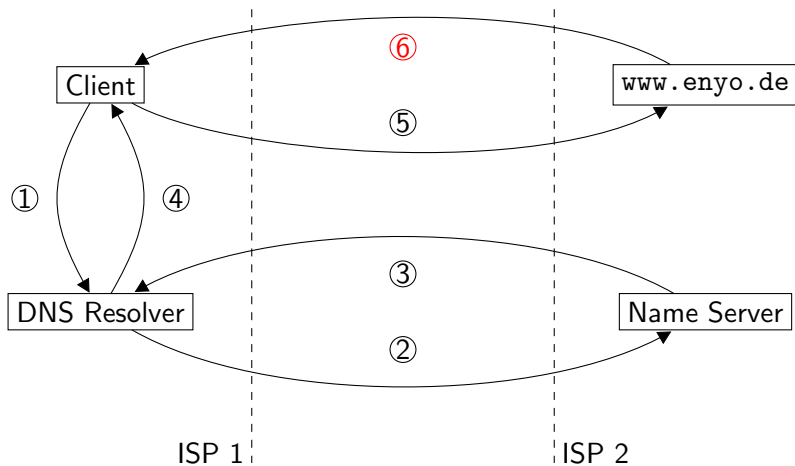
# Sensor placement




# Sensor placement



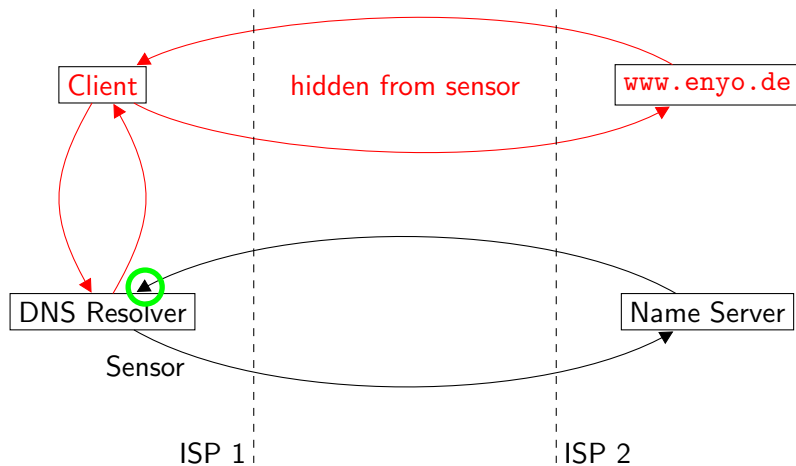
## Sensor placement



## Sensor placement

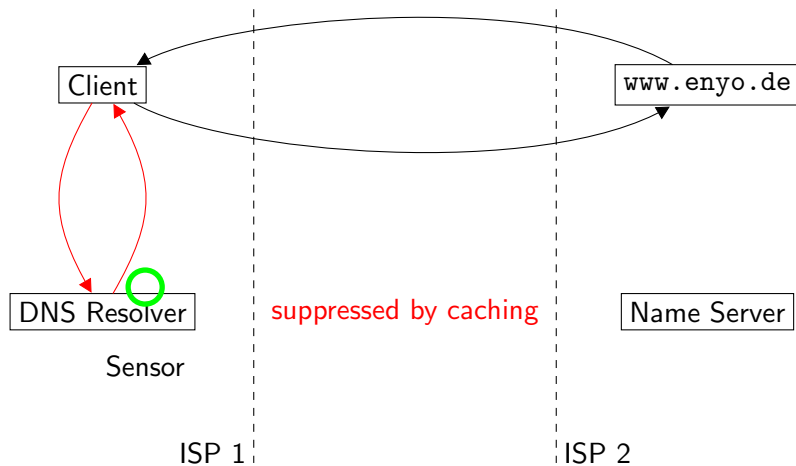


# Sensor placement





# Sensor placement



# Challenges

- ▶ Privacy concerns
- ▶ Security concerns
- ▶ DNS-specific problems
  - ▶ Truncated and EDNS0 responses
  - ▶ What about bogus DNS data captured by the sensors?
- ▶ Data volume
  - ▶ The data rate itself is fairly low on medium-sized campus networks.
  - ▶ But keeping lots of historic data is problematic.

# dnslogger implementation

- ▶ Two sensor implementations:
  - ▶ Perl: simple and obviously correct
  - ▶ C: higher performance, fewer dependencies
- ▶ The remaining parts of the `dnslogger` software are written in Ada.
- ▶ Berkeley DB from Sleepycat is used as the underlying database technology.
  - ▶ The schema design is highly denormalized and clustered on reversed domain names.
  - ▶ All database updates are idempotent and commute, which makes replication particularly easy.

# Outline

A very brief introduction to DNS

Case Study: Botnet mitigation

Architecture and implementation

Results

## Example: The kimble.org fiasco

First seen	Domain	Type	Data
2004-06-23 13:58:51	kimble.org	A	127.0.0.1
2004-08-07 16:14:00	kimble.org	A	207.234.155.17
2004-10-20 07:15:58	kimble.org	A	212.100.234.54
2004-10-20 16:12:56	kimble.org	A	64.203.97.121
2004-10-21 17:15:01	kimble.org	A	212.113.74.58
2004-10-21 17:45:01	kimble.org	A	195.130.152.100
2004-10-31 14:45:01	kimble.org	A	195.225.218.59
2004-11-02 23:15:01	kimble.org	A	206.132.83.2
2004-11-04 18:15:01	kimble.org	A	213.139.139.206
2004-11-21 03:15:02	kimble.org	A	216.7.173.212
2004-11-25 22:45:02	kimble.org	A	38.112.165.60

## Example: Hijacking of ebay.de

First seen	Domain	Type	Data
2004-06-23 08:21:57	ebay.de	NS	crocodile.ebay.com
2004-06-23 08:21:57	ebay.de	NS	sjc-dns1.ebaydns.com
2004-06-23 08:21:57	ebay.de	NS	sjc-dns2.ebaydns.com
2004-08-28 05:34:01	ebay.de	NS	ns1.goracer.de
2004-08-28 05:34:01	ebay.de	NS	ns2.goracer.de

## Example: Network Solution's "Site Finder Light"

First seen	Domain	Type	Data
2004-09-19 05:01:53	misslink.net	CNAME	†
2004-09-19 05:57:49	ns.bighornent.com	CNAME	†
2004-09-19 06:13:44	ns13.magnum-inap4.net	CNAME	†
2004-09-19 06:24:28	host2.7thgate.com	CNAME	†
2004-09-19 07:25:26	www.zydigo.com	CNAME	†
2004-09-19 08:08:33	muslimsonline.com	CNAME	†
2004-09-19 08:28:26	www.animatiehuis.com	CNAME	†
2004-09-19 08:57:19	www.urbanvoicesonline.com	CNAME	†
...			

† = resalehost.networksolutions.com

## Example: Correlating domains

- ▶ An email messages references `dkchaotichigh.com` (“MegaPowerPills.com”).
- ▶ An ordinary DNS lookup reveals that `ns1.m-dns.us` is used as a name server.
- ▶ Additional domains are hosted on this name server:

<b>Domain</b>	<b>Type</b>	<b>Data</b>
<code>outfacegood.com</code>	NS	<code>ns1.m-dns.us</code>
<code>outregood.com</code>	NS	<code>ns1.m-dns.us</code>
<code>megalithgood.com</code>	NS	<code>ns1.m-dns.us</code>
<code>medverdigrisgood.com</code>	NS	<code>ns1.m-dns.us</code>
...		



## Example: Unauthorized name servers for .com

First seen	Domain	Type	Data
2004-06-24 00:52:37	com	NS	ns1.hi2000.net
2004-06-24 23:04:11	com	NS	ns1.vertical-inc.net
2004-06-30 23:26:21	com	NS	tempsvr.wam.wamusa.com
2004-07-01 04:32:18	com	NS	ns7.domainredirect.com
2004-07-05 04:18:36	com	NS	ns1.infoglobe.net
2004-07-05 07:35:14	com	NS	ns1.cntrading.com
2004-07-05 16:40:27	com	NS	ns1.spacesurfer.com
2004-07-08 00:34:29	com	NS	ns.tradenames.com
...			

# Observations

- ▶ DNS usage is very localized and specific to the network in which the sensor is placed.
- ▶ For many applications, you have to run your own sensor, instead of using data collected on other networks.
- ▶ But sharing the data with others does not hurt.

# Summary

- ▶ Passive DNS replication provides new ways to access and process DNS data.
- ▶ This data can support various security-related processes.
- ▶ It also provides new insights into the operation of the domain name system.

Thank you for your attention!

Questions?