The Void
An Interesting Place For Network Security Monitoring

CIRCL
Computer Incident
Response Center
Luxembourg

Alexandre Dulaunoy, CIRCL-

TLP:WHITE

alexandre.dulaunoy@circl.lu

November 13, 2014

CIRCL, national CERT of Luxembourg

- CIRCL[1] is composed of 6 full-time incident handlers + 2 FTE back up operators.
- The team is operating as an autonomous technical team relying on its own infrastructure.
  - Operators competencies include reverse engineering, malware analysis, network and sy stem forensic, software engineering and data mining.
- CIRCL, the national CERT, is part of SMILE[2] gie (a publicly funded organization to promote information security in Luxembourg).
- In 2013, CIRCL handled more than 35000 security events and conducted more than 1000 technical investigations.

[1] `http://www.circl.lu/`

[2] `http://www.smile.public.lu/`

Motivation

- IP-darkspace is
  - Routable non-used address space of an ISP (Internet Service Provider),
  - arriving traffic is unidirectional
  - and unsolicited[3].
- Is there any traffic in those darkspaces?
- If yes, what and why does it arrive there?
  - And on purpose or by mischance?
- What's the security impact?
- What are the security recommendations?

[3] If the black-hole is not abused.

Why is there traffic?

Origins

- Attackers (and researchers) scan networks to find vulnerable systems (e.g. SSH brute-force).
- Backscatter traffic (e.g. from spoofed DoS).
- Self-replicating code using network as a vector (e.g. conficker, residual worms).
- Badly configured devices especially embedded devices (e.g. printers, server, routers).
  - $\rightarrow$ One of our IP-darkspace is especially suited for spelling errors from the RFC1918 (private networks) address space.

Why is there traffic

Typing/Spelling errors with RFC1918 networks

- While typing an IP address, different error categories might emerge:

Hit wrong key	19**2**.x.z.y →	19**3**.x.y.z
Omission of number	1**9**2.x.y.z →	12.x.y.z
Doubling of keys	10.a.b.c →	10**0**.a.b.c
	172.x.y.z	1**5**2.x.y.z

Research activities related to spelling errors

Spelling errors apply to text but also network configuration

- 34% omissions of 1 character
  - Example: Network $\rightarrow$ Netork
- 23% of all errors happen on 3rd position of a word
  - Example: Text $\rightarrow$ Test)
- 94% spellings errors are single errors in word
  - And do not reappear

References

- Pollock J. J. and Zamora A., Collection and characterization of spelling errors in scientific and scholarly text. J. Amer. Soc. Inf. Sci. 34, 1, 51 58, 1983.
- Kukich K., Techniques for automatically correcting words in text. ACM Comput. Surv. 24, 4, 377-439, 1992.

IP-Darkspace: Data Collection

Implementation

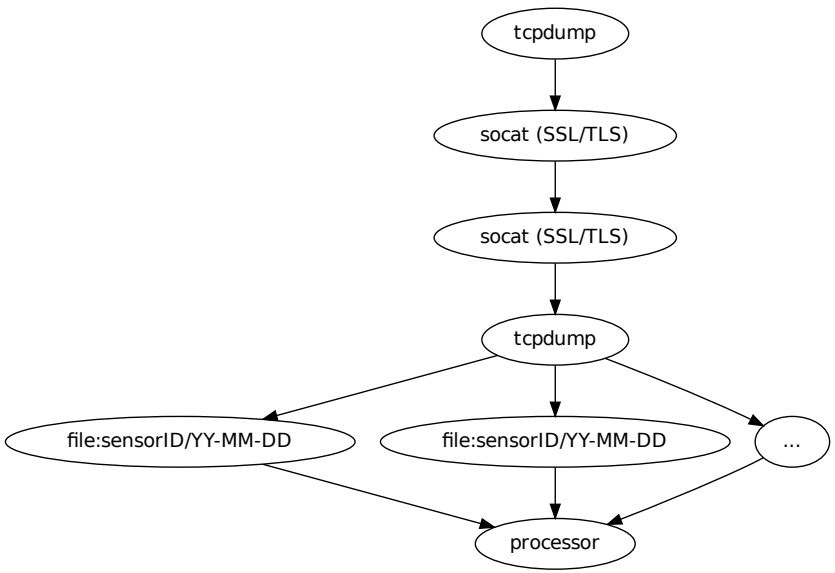

- Minimal sensor collecting IP-Darkspace networks (close to RFC1918 address space).
- Raw pcap are captured with the full payload.
- Netbeacon[a] developed to ensure consistent packet capture.

[a] `www.github.com/adulau/netbeacon/`

An example of a dataset collected

- from 2012-03-12 until Today (still active).
- 260 gigabytes of raw pcap were collected.
- Constant stream of packets (150kbit/s) from two /22 network blocks.
  - no day/night profile.
- Some peaks at 2Mb/s (e.g. often TCP RST from back scatter traffic or short-term misconfiguration).

General observations

- A large part of traffic is coming from badly configured devices (e.g. RFC1918 spelling errors).
  - Printers, embedded devices, routers or even server.
  - Trying to do name resolution on non-existing DNS servers, NTP or sending syslog messages.
- Even if the black-hole is passive, payload of stateless UDP packets or even TCP (due to asymmetric routing on misspelled network) datagrams are present.
- Internal network scanning and reconnaissance tool (e.g. internal network enumeration).

Observation per AS

Traffic seen in the darknet

N	Frequency	ASN
1	4596319	4134
2	1382960	4837
3	367515	3462
4	312984	4766
5	211468	4812
6	166110	9394
7	156303	9121
8	153585	4808
9	135811	9318
10	116105	4788

- Occurrences of activities matching the proportion of hosts in a country.
- Chinese great-wall is not filtering leaked packets.

Network reconnaissance: a few machine names

ASTTF.NET	HELP.163.COM	
ASUEGYI.INFO	HP_CLIENT1	
ASUS1025C	MACBOOKAIR-CAD7	
DEFAULT	MACBOOK-B5BA66	
DELICIOUS.COM	MACBOOKPRO-5357	
DELL	MAIL.AFT20.COM	And many more ...
DELL1400	S3.QHIMG.COM	
DELL335873	SERVERWEB	
DELL7777	SERVEUR	
DELL-PC	SERVICE.QQ.COM	
DELLPOP3	SMTP.163.COM	

Network reconnaissance: NetBios machine types

23	Browser Server
4	Client?
1	Client? M <ACTIVE>
21	Domain Controller
1	Domain Controller M <ACTIVE>
11	Master Browser
1	NameType=0x00 Workstation
1	NameType=0x20 Server
105	Server
26	Unknown
1	Unknown <GROUP> B <ACTIVE>
5	Unknown <GROUP> M <ACTIVE>
1322	Workstation
1	Workstation M <ACTIVE>

Network reconnaissance (and potential misuse): DNS

```
3684 _msdcs.<companyname>.local
1232666 time.euro.apple.com
104 time.euro.apple.com.<mylocaldomain>
122 ocsp.tcs.terena.org
50000+ ocsp.<variousCA>
```

- DNS queries to an incorrect nameserver could lead to major misuse.
- A single typo in a list of 3 nameservers is usually unnoticed.
- Defeating OCSP, Moxie Marlinspinke[4].

[4] `http://safecomputing.umich.edu/events/sumit09/docs/Moxi\more2.pdf`

From passive collection to dynamic exploitation?

```
23:52:29.818155 IP 41.229.54.252.1025 > X.168.66.11.53: 21030+ A? wpad. (22)
23:53:09.073601 IP 41.229.54.252.1025 > X.168.66.10.53: 24576+ A? wpad. (22)
23:53:10.068080 IP 41.229.54.252.1025 > X.168.66.11.53: 24576+ A? wpad. (22)
23:53:11.063357 IP 41.229.54.252.1025 > X.168.66.10.53: 24576+ A? wpad. (22)
23:53:13.062686 IP 41.229.54.252.1025 > X.168.66.10.53: 24576+ A? wpad. (22)
23:53:13.068506 IP 41.229.54.252.1025 > X.168.66.11.53: 24576+ A? wpad. (22)
23:53:17.063567 IP 41.229.54.252.1025 > X.168.66.11.53: 24576+ A? wpad. (22)
23:53:17.067365 IP 41.229.54.252.1025 > X.168.66.10.53: 24576+ A? wpad. (22)
23:53:56.314674 IP 41.229.54.252.1025 > X.168.66.10.53: 57865+ A? wpad. (22)
23:53:57.317966 IP 41.229.54.252.1025 > X.168.66.11.53: 57865+ A? wpad. (22)
23:53:58.313341 IP 41.229.54.252.1025 > X.168.66.10.53: 57865+ A? wpad. (22)
23:54:00.312687 IP 41.229.54.252.1025 > X.168.66.10.53: 57865+ A? wpad. (22)
23:54:00.318675 IP 41.229.54.252.1025 > X.168.66.11.53: 57865+ A? wpad. (22)
23:54:04.312157 IP 41.229.54.252.1025 > X.168.66.10.53: 57865+ A? wpad. (22)
```

- Web Proxy Autodiscovery Protocol is still used in order to find a proxy automatically.
- WPAD fetches a PAC file (JavaScript executed even if JavaScript is disabled) to give the IP address of the proxy.

Network scanning and passive collection

- Enumerating hostname in a single domain can be used for reconnaissance.
- Passive DNS collection allows to build a corpus of probable hostname.
- Then you can use the corpus in your favorite network scanner.
- Wagner, Cynthia, Jérôme François, Gérard Wagener, and Alexandre Dulaunoy. "SDBF: Smart DNS brute-forcer." In Network Operations and Management Symposium (NOMS), 2012 IEEE, pp. 1001-1007. IEEE, 2012. [5]

[5] `http://www.foo.be/papers/sdbf.pdf`

A/V Statistics from Misconfigured Resolvers

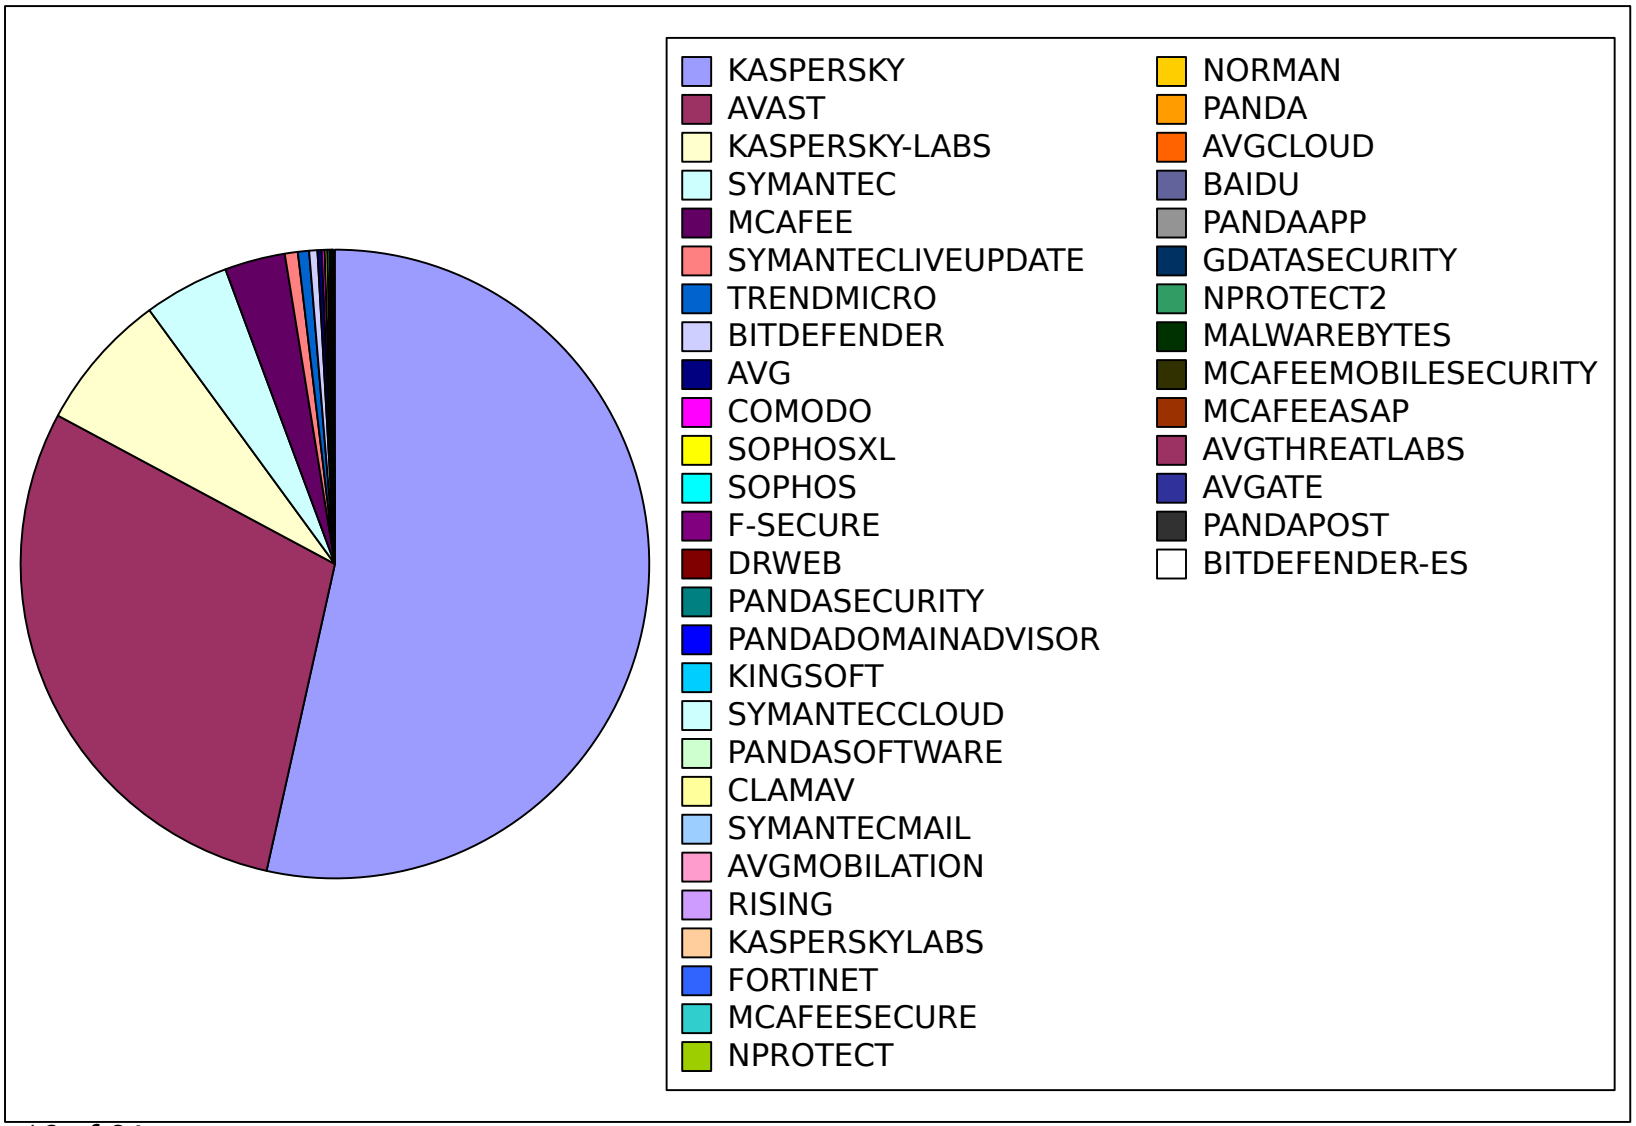

Certificate Revocation and Queries from Misconfigured Resolvers

- The increase of 5% in late 2013 might be due to certificate requirements update (e.g. key size, hashing algorithm updates)
- A lot of software assumes a certificate to be valid when OCSP or CRL are not accessibles

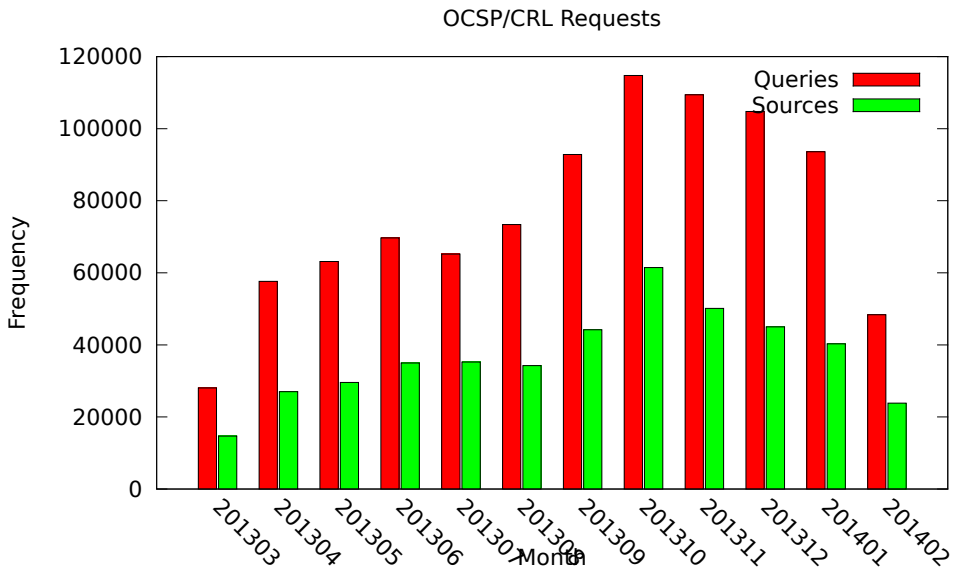

Software Updates/Queries from Misconfigured Resolvers

- Discovering software usage (and vulnerabilities) can be easily done with passive reconnaissance
- Are the software update process ensuring the integrity of the updates?

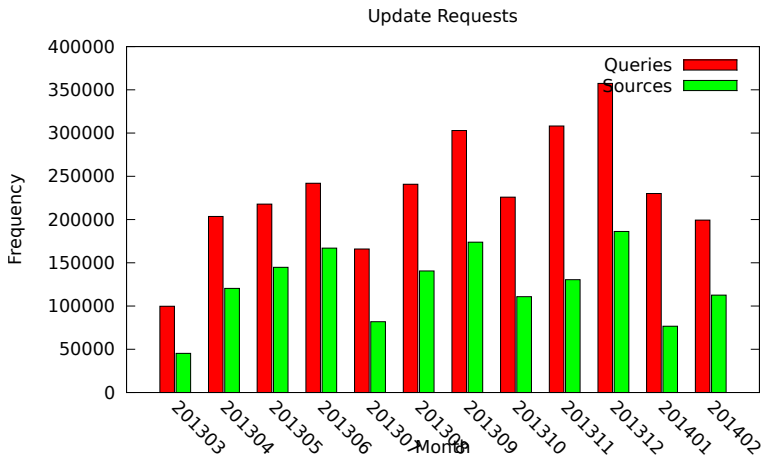

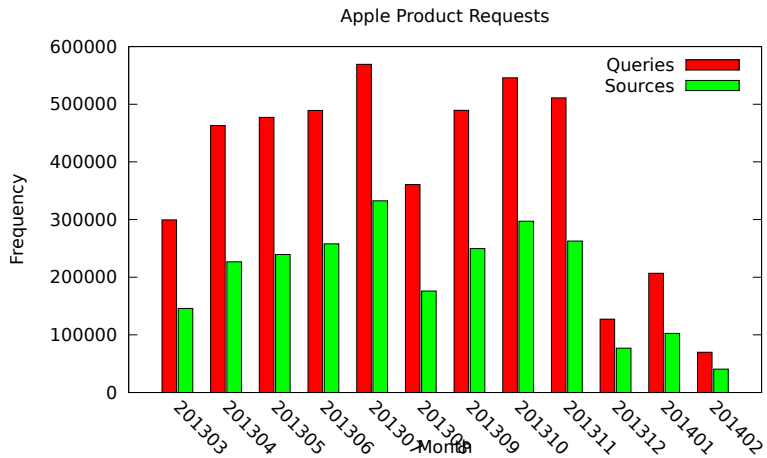

Printer syslog to the world

or how to tell to the world your printer status

```
2012-03-12 18:00:42
   SYSLOG lpr.error printer: offline
     or intervention needed
2012-03-23 21:51:24.985290
   SYSLOG lpr.error printer: paper out
   ...
2012-08-06 19:14:57.248337
   SYSLOG lpr.error printer: paper jam
```

- Printers are just an example out of many syslog messages from various devices.
- Information leaked could be used by attackers to gain more information or improve targeted attacks.

How to configure your router (without security)

Enable command logging and send the logs to a random syslog server

```
Aug 13 10:11:51 M6000-G5 command-log:[10:11:51 08-13-2012
  VtyNo: vty1  UserName: XXX IP: XXX ReturnCode: 1
  CMDLine: show subscriber interface gei-0/2/1/12.60
Aug 13 10:46:05 M6000-G5 command-log:[10:46:05 08-13-2012
  VtyNo: vty2  UserName: XXX IP: XXX  ReturnCode: 1
  CMDLine: conf t ]
Aug 13 10:46:10 M6000-G5 command-log:[10:46:10 08-13-2012
  VtyNo: vty2  UserName: XXX IP: XXX  ReturnCode: 1  CMD
Line: aaa-authentication-template 1100 ]
...
```

We will let you guess the sensitive part afterwards...

Misconfigured network interception in Iran for 2 hours?

- On April 08, 2013, a peak of ICMP time exceeded in-transit were received during 2 hours
- IP sources allocated in Iran with a nice distribution among Iranian Internet providers

```
12:29:49.255942 IP 93.126.56.1 > a.b.100.1: ICMP time exceeded in-transit, length 36
12:29:49.255957 IP 80.191.114.59 > a.b.100.1: ICMP time exceeded in-transit, length 36
12:29:49.255963 IP 80.191.114.59 > a.b.100.1: ICMP time exceeded in-transit, length 36
12:29:49.256144 IP 93.126.56.1 > a.b.100.1: ICMP time exceeded in-transit, length 36
12:29:49.256172 IP 93.126.56.1 > a.b.100.1: ICMP time exceeded in-transit, length 36
12:29:49.256481 IP 93.126.56.1 > a.b.100.1: ICMP time exceeded in-transit, length 36
12:29:49.256568 IP 93.126.56.1 > a.b.100.1: ICMP time exceeded in-transit, length 36
12:29:49.257086 IP 80.191.114.59 > a.b.100.1: ICMP time exceeded in-transit, length 36
12:29:49.257098 IP 80.191.114.59 > a.b.100.1: ICMP time exceeded in-transit, length 36
12:29:49.257470 IP 93.126.56.1 > a.b.100.1: ICMP time exceeded in-transit, length 36
12:29:49.257565 IP 80.191.114.59 > a.b.100.1: ICMP time exceeded in-transit, length 36
12:29:49.257603 IP 80.191.114.59 > a.b.100.1: ICMP time exceeded in-transit, length 36
12:29:49.258575 IP 178.173.128.245 > a.b.100.1: ICMP time exceeded in-transit, length 36
12:29:49.258657 IP 178.173.128.245 > a.b.100.1: ICMP time exceeded in-transit, length 36
12:29:49.258669 IP 178.173.128.245 > a.b.100.1: ICMP time exceeded in-transit, length 36
12:29:49.258677 IP 178.173.128.245 > a.b.100.1: ICMP time exceeded in-transit, length 36
```

Research Opportunities

- Analysis of noise traffic in order to discover patterns or similarities among collectors.
- Network packet data storage, indexing and fast lookup (e.g. bitindex, bloomfilter, privacy-preserving dataset).
- Detecting abuse of black-hole sensors.
- Analysis of country-wide Interception from noise traffic.
- Automatic exploitation using passive reconnaissance.

Conclusions

- Security recommendations
  - **Default routing/NAT to Internet in operational network is evil**.
  - Use fully qualified domain names.
  - Double check syslog exports via UDP (e.g. information leakage is easy).
  - Verify any default configuration with SNMP (e.g. enable by default on some embedded devices).
- Offensive usage? What does it happen if a malicious Internet operator is responding to misspelled RFC1918 addresses? (e.g. DNS/NTP requests, software update or proxy request).

- Interested in a research project on similar dataset? or an internship on some technically interesting project?
- $\rightarrow$ alexandre.dulaunoy@circl.lu
- PGP: 3B12 DCC2 82FA 2931 2F5B 709A 09E2 CD49 44E6 CBCD