













Mapping the Dutch Critical Infrastructure

Razvan C. Oprea
Fahime Alizade

Supervised by Benno Overeinder  NLnet Labs

The initial question

Critical infrastructure sectors

- | | |
|---|--|
|  Surface water |  Food |
|  Chemical industries |  Health |
|  Public administration |  ICT |
|  Drinking water |  Transport |
|  Public order |  Finance |
|  Legal order |  Energy |



What is the network level representation of the critical infrastructure?

Previous research

Publicly accessible related research papers are scarce

PAM2012: "*Exposing a Nation-Centric View on the German Internet – A Change in Perspective on AS-Level*"

The research started from prefixes and discovered Autonomous Systems Numbers (ASNs) using RIPE database, Team Cymru and RIPE RIS

The AS interconnections were discovered using BGP dumps

Research Questions

Can we discover and map the Internet entities corresponding to the Dutch national critical infrastructure with a sufficient degree of confidence?

Our hypothesis is that the answer to the above question is affirmative

Subquestions:

What are the authoritative sources of information?

What is the resilience of Dutch critical infrastructure?

Methodology

We have no idea on organizations' physical connections to the Internet, but we are interested in the logical IP topology:

- we work at an **AS level**
 - we use two methods for discovering relevant ASNs
-

1 Bottom-up discovery approach

We discover the “Dutch” ASNs, then we identify organizations in critical sectors

2 Top-down approach

Starting from organizations in critical sectors, we identify the corresponding ASNs

3 Analysis and visualization

We combine the results of the two approaches, find interconnections and build graphs

Bottom-up Approach

We use the ASN allocation list published by the RIPE NCC

We select the ASNs allocated to organizations registered in NL or EU

Every EU ASN is queried in the RIPE WHOIS database to select NL registrations (address or description fields)

We select the organizations in the critical infrastructure sectors (domain name, KvK)

```
ripenc |*|asn|*|28184|summary
ripenc |EU|asn|1196|1|19930901|allocated
ripenc |IE|asn|1197|1|20101118|allocated
ripenc |EU|asn|1198|1|19930901|allocated
ripenc |EU|asn|1199|1|19930901|allocated
ripenc |NL|asn|1200|1|19930901|allocated
ripenc |EU|asn|1203|1|19930901|allocated
ripenc |AT|asn|1205|1|19930901|allocated
ripenc |IE|asn|1213|1|19920617|allocated
ripenc |EU|asn|1234|1|19930901|allocated
ripenc |EU|asn|1235|1|19930901|allocated
ripenc |EU|asn|1241|1|19930901|allocated
ripenc ||asn|48198|1||reserved
ripenc ||asn|48204|1||reserved
ripenc ||asn|48253|1||reserved
ripenc ||asn|12410|1||available
ripenc ||asn|15449|1||available
ripenc ||asn|15907|1||available
ripenc ||asn|16078|1||available
```

Bottom-up Approach (contd.)

Limitations

We do not know if all the ASNs of an organization relate to critical infrastructure

We have limited information on organization structure and ownership (Virtual ASNs)

The number of “Dutch” ASNs in the Internet sector is disproportionately high

We decided to keep ISPs, Data Centers, Internet Exchange Points

Observations

727 ASNs allocated to Dutch organizations

335 ASNs relate to the critical infrastructure sectors

265 ASNs relate to the Internet infrastructure sector

Top-Down Approach

We search for well-known entities in each critical sector

We find the organization name (KvK) and their domain

We search for the IP addresses corresponding to their A, AAAA and MX records

We use RIPEstat to find the prefix it is part of and the originating ASN (the “proxy” AS)

Top-Down Approach (contd.)

Limitations

We decided early on to use only public information

Complete mapping of critical sector industries requires specialized knowledge (think food chain supply)

Backup and private links are not visible

Observations

We tried to have at least few samples from every sector

In total, we hand-picked 147 organizations part of the Dutch critical infrastructure

Data analysis

We combine the result of the two approaches and obtain a “master” ASNs list.

The inter-AS relationships is visible in BGP dumps, but it's better to have multiple viewpoints for accuracy

RIPE RIS, RouteViews, Route Servers, Looking Glasses all offer multi-views on the BGP links

traceroute is not a viable option since the IP address space used by organizations is privileged information

We considered the aggregated data offered by UCLA IRL, CAIDA and University of Washington and we ultimately chose UCLA

Data analysis (contd.)

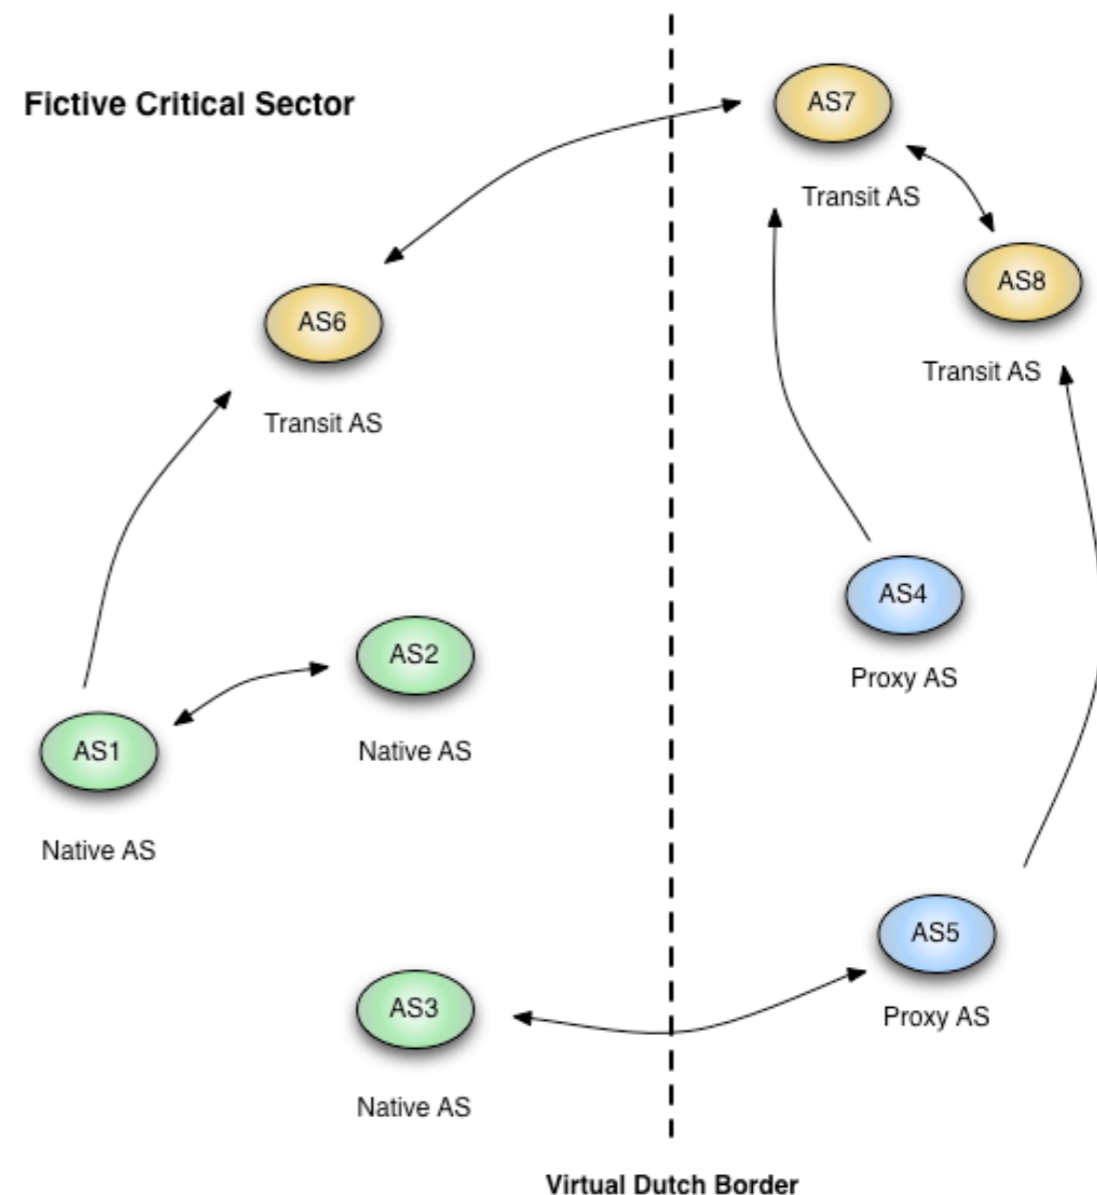
Many nodes (ASNs) are abroad

The initial graphs show many disconnected nodes

Which ASNs to include to show relevant links?

We choose to include the providers of the native and proxy ASNs

We then built the full mash of the AS and provider list based on UCLA data



Visualization Methods

To display and present high number of AS numbers and their relations, **HTML canvas**, **Javascript** and **jQuery** are chosen.


We need an interactive presentation of graph to zoom-in and to see labels.

Different Javascript libraries are taken into account:
D3.js, **Sigma.js**

D3.js

Data Driven Documents

We formatted our dataset in two Json files:
Nodes and **Links**



Node positioning: **Force Layout**

By modifying links constraints the layout finds the best-fitted position for each node.

```
[  
{  
  "as": "286",  
  "company": "Brabant Water",  
  "sector": "C1",  
  "input": ["proxy",  
            {"record": "A",  
             "company": "KPN",  
             "country": "NL"}]  
}  
]
```

Sigma.js

We chose **Sigma.js**, which is an open source Javascript library.


We could parse Json files using **jQuery**

In contrast to D3.js, positioning layouts are not provided.

Nodes with the higher degree are put in inner levels.

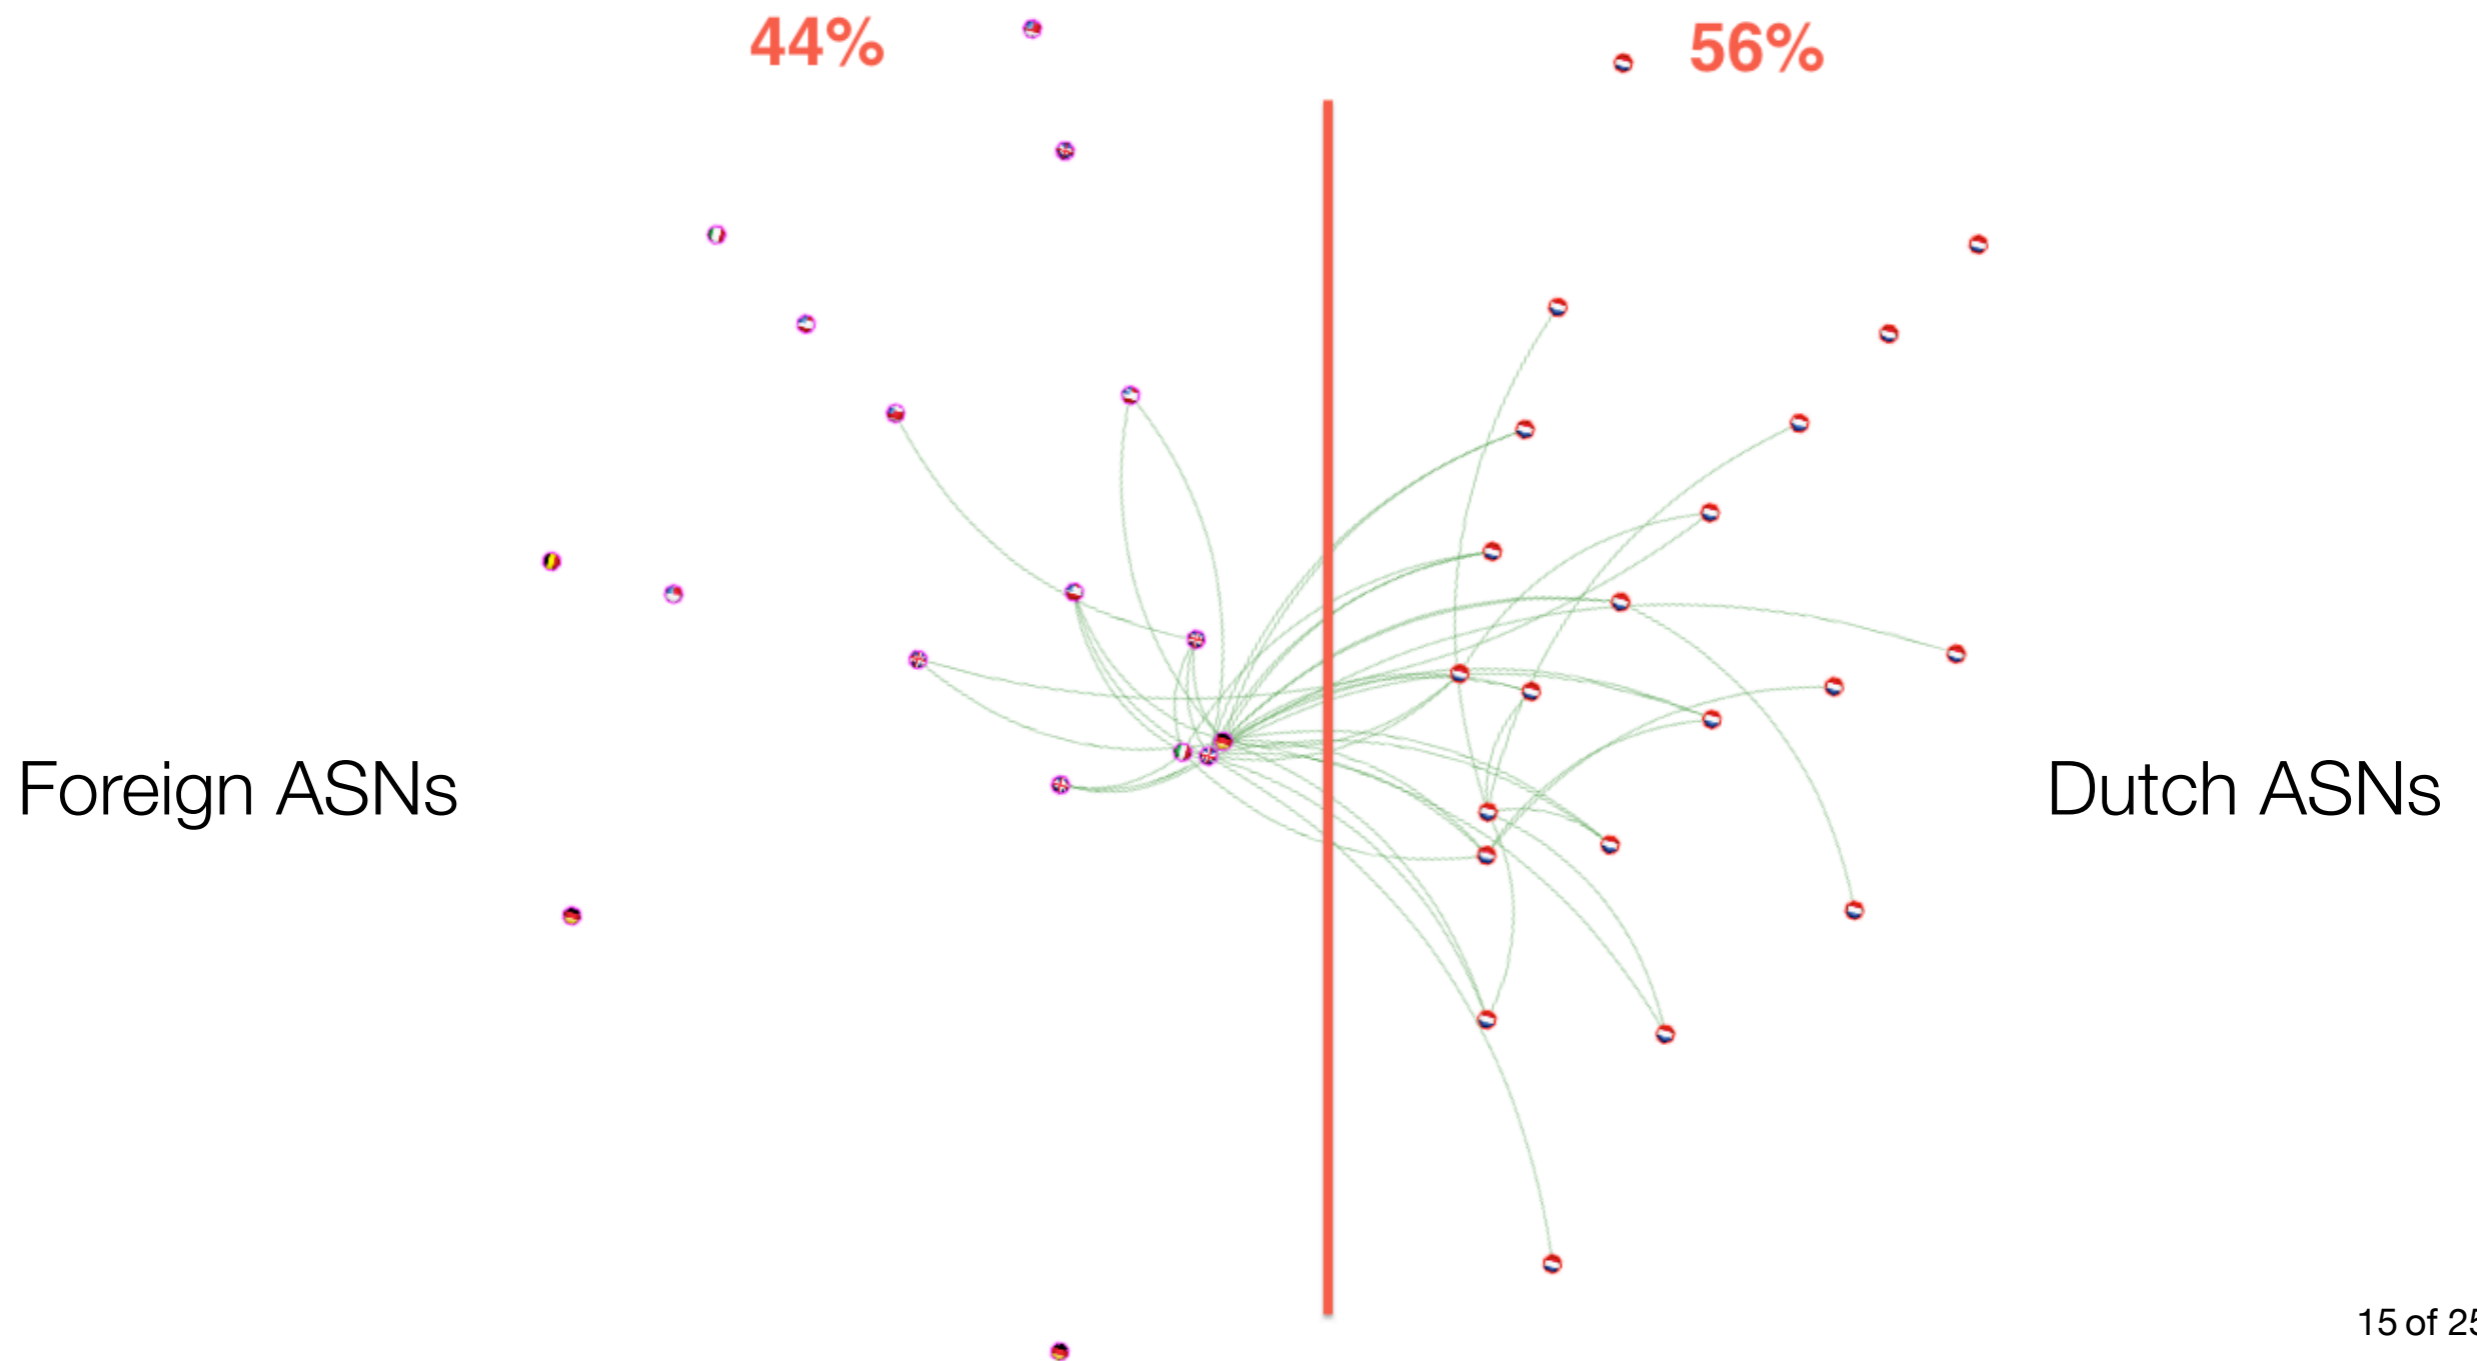
$$y = \text{Cos}(a) * r$$

$$x = \text{Sin}(a) * r$$




Visualization and conclusions

Energy Sector - no providers




Visualization and conclusions (contd.)

Energy Sector - with providers



Visualization and conclusions (contd.)

Food Sector - with providers



Observations

- 1 Related companies/industries choose sometimes the same providers: **NS and ProRail (BT), Royal Dutch Shell, Gasunie and Argos Energies (Microsoft Corp.)**
- 2 Some organizations have their own ASN, but they still outsource their email and website hosting (**Alliander**).
- 3 The biggest providers (mail) are **MessageLabs (UK & US), KPN, Microsoft, Tele2 Nederland and Ziggo.**

Observations (contd.)

4 What do **ABN AMRO**, **Triodos Bank**, **AkzoNobel**, **GGD** have in common: all their mails come through the same provider: MessageLabs Ltd., UK

Nine other companies in the critical sectors use the services of MessageLabs Inc., US

5 In fact, **MessageLabs** (a division of Symantec Corp.) is the single biggest messaging provider in our list

Observations (contd.)

Sector	Dutch Provider	Foreign Provider	Top 1 Foreign Provider
Energy	56%	44%	Microsoft Corp. ,US
ICT	96%	4%	Websense hosted, UK
Drinking water	61%	39%	MessageLabs Inc., US
Food	63%	37%	There is no biggest one!
Health	75%	25%	MessageLabs Ltd. ,UK
Finance	81%	9%	MessageLabs Ltd. ,UK
Surface water	57% (no Native)	43%	Microsoft Corp. ,US
Public order	92%	8%	ClaraNET Ltd. ,UK
Legal order	67%	33%	BT PLC, UK
Public	74%	26%	MessageLabs Ltd., UK
Transport	61%	39%	BT PLC, UK
Chemical	36%	64%	MessageLabs Inc., US

Table 1. Distribution of Mail providers in each sector

Observations (Dutch government)

Dutch ministries accessible through two umbrella domains:

- 1
 - government.nl - A (Prolocation, NL), MX (MessageLabs, UK and MessageLabs, US)
 - rijksoverheid.nl - A (Prolocation B.V., NL), MX (KPN, NL)

2 Courts of Justice accessible through one umbrella domain:

- rechtspraak.nl - A (ASP4ALL Hosting, NL), MX (Tele 2 Nederland, NL)

3 Ministry of Defense website is accessible via the rijksoverheid.nl domain

However, military branches (like infantry, marine, aviation) use their own infrastructure (domain and AS)

Conclusions

We do not see physical, private and back-up links.

We could discover the representative Dutch critical infrastructure organizations using the two discovery methods (bottom-up and top-down).

The discovered organizations were verified manually one-by-one so we have a high degree of confidence.

A more comprehensive list of organizations can only be obtained with specialized and preferably privileged information.

Conclusions (contd.)

Many critical infrastructure organizations have reliable connections to the Internet, but rely a lot on foreign providers for their communication needs

It is worth discussing the security and privacy implications of having email and websites hosted with entities from outside the EU

We do not see that critical infrastructure organizations regard their network infrastructure as being of national critical importance

Any questions?

Thank you for your attention!