FENIX and CESNET approach to DDoS

Ondřej Caletka

November 2015

About CESNET

- association of legal entities, est. 1996
  - public and state universities
  - Academy of Sciences
- non-profit organisation
  - development and operation of **NREN**
  - advanced network technologies and applications R&D
  - international cooperation – GNx, GN3+, GLIF, EGI, GÉANT shareholder, EGI member, Internet2 affiliate member,...
- founding member – **CZ.NIC, NIX.CZ, FENIX**

About NIX.CZ

- association of legal entities, est. 1996
- non-profit organisation
  - community driven
  - members and customers
- operator of public neutral **IXPs**
  - NIX.CZ – Prague
    - 5 PoPs
    - 136 networks
    - 1.8 Tb capacity
  - NIX.SK – Bratislava – since 2015
    - 2 PoPs
    - 29 networks
    - 150 Gb capacity

(D)DoS attacks in 2013

- between March 4th and 7th
- two waves each day: 9am – 11am, 2pm – 4pm
- targeting major Czech web sites

Monday	news portals
Tuesday	search engine `www.seznam.cz`
Wednesday	bank websites
Thursday	2 out of 3 mobile carriers

- attractive for mass media

DoS technical aspects

- sourced from transit operator RETN **via NIX.CZ**
- methods: SYN-Flood, DNS-reflection
- no harm for ISP
  - low volume (< 1 Gbps)
  - moderate packet rate (1 – 1.5 million pps)
- harmful for end sites
  - aggregation in one point
  - no SYN-cookies enabled
  - firewalls and loadbalancers up in smoke
- used solutions
  - controlled shutdown and waiting for the end of attack
  - moving service to another IP address (short DNS TTL)
  - filtration, scrubbers
  - **restricting traffic just for Czech ISPs**

Lessons learned

- NIX.CZ peering ≠ national peering
- NIX.CZ can transit spoofed traffic
- some victims misinterpreted attack transited via NIX.CZ as attack sourced from Czechia

Idea of secure peering VLAN inside NIX.CZ

- as a last resort in case of some massive attack
- for those that **trust each other**
- so **Czech users can access Czech services**

So the FENIX was born...

- club of **trustworthy operators** inside NIX.CZ which
  - avoid IP spoofing
  - take care of security incidents
- self-governed, independent of NIX.CZ
  - NIX.CZ act as an arbiter
  - new members need recommendations
  - any member can veto
- self-regulation instead of goverment regulation
- high entry threshold

Connected to trusted network

Trusted operator

FENIX organisational criteria

- Terms and Conditions allowing to disconnect customer originating malicious traffic
- 24×7 NOC, no IVR
- Trusted Introducer listed CSIRT team
- NIX.CZ member for more than 6 months
- active participation
- recommendation from 2 FENIX members, no veto

FENIX technical criteria

- **BCP38**/SSAC004 network ingress filtering
- RTBH using route servers
- fully redundant connection to NIX.CZ
- protected BGP sessions with TCP MD5
- incident reaction time less than 30 minutes
- DNS, NTP, SNMP amplification protection
- deployed IPv6 and DNSSEC
- control plane policy (RFC 6192)
- network monitoring with alerts (MTRG, NetFlow,...)

FENIX foundation

- founded by 6 operators in January 2014
  - **Active 24** (hosting)
  - **CESNET** (NREN)
  - **CZ.NIC** (TLD operator)
  - **Dial Telecom** (ISP)
  - **O2 CZ** (ISP, incumbent)
  - **Seznam.cz** (Czech Google)
- 12 operators today

Secure VLAN

- former work title for the FENIX
- separate peering VLAN of last resort
- accessible by FENIX members only
- prepared for island-mode of operation
- no data during *peace time*
- each member decides on their own when to use it

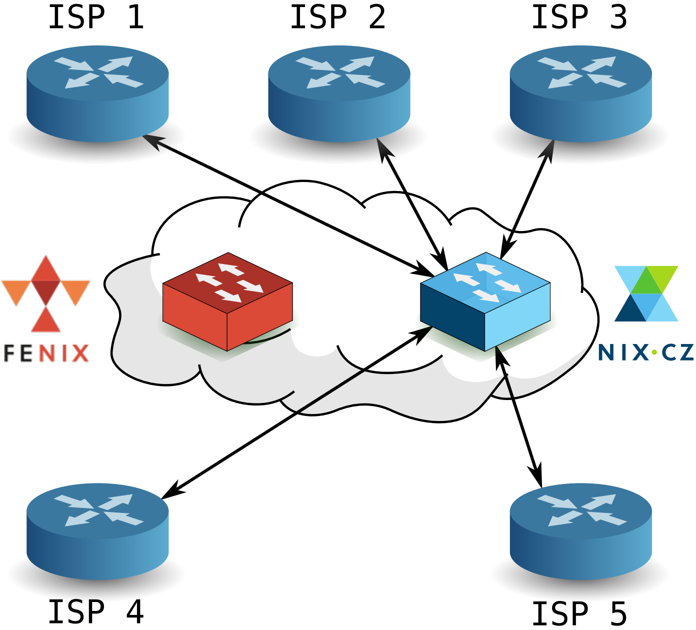
ISP 1
ISP 2
ISP 3
FENIX
NIX·CZ
ISP 4
ISP 5

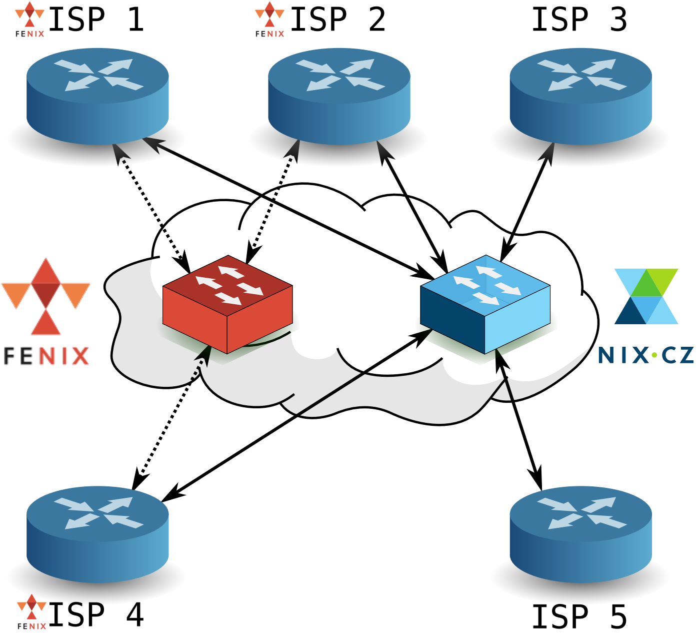
FENIX
ISP 1
FENIX
ISP 2
ISP 3
FENIX
NIX·CZ
FENIX
ISP 4
ISP 5

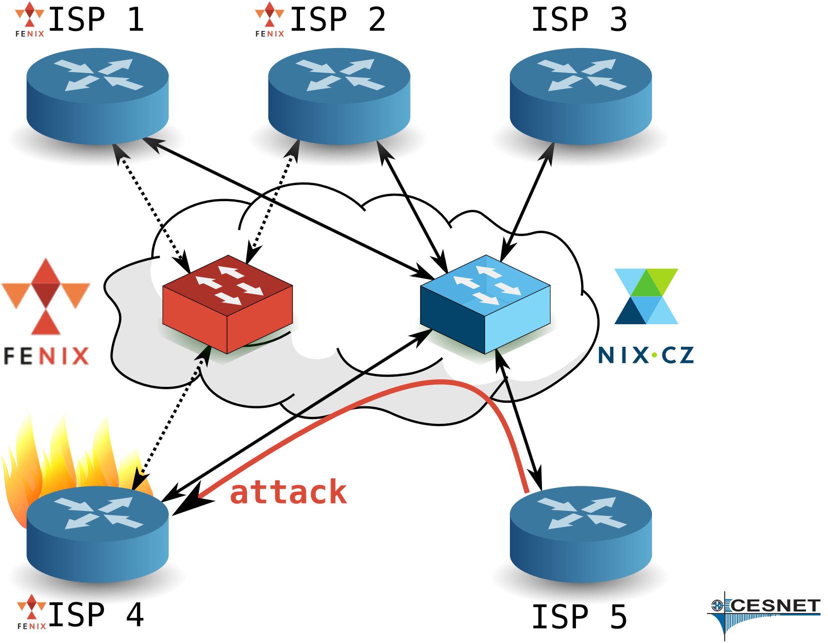
ISP 1
ISP 2
ISP 3
FENIX
NIX·CZ
attack
ISP 4
ISP 5
CESNET

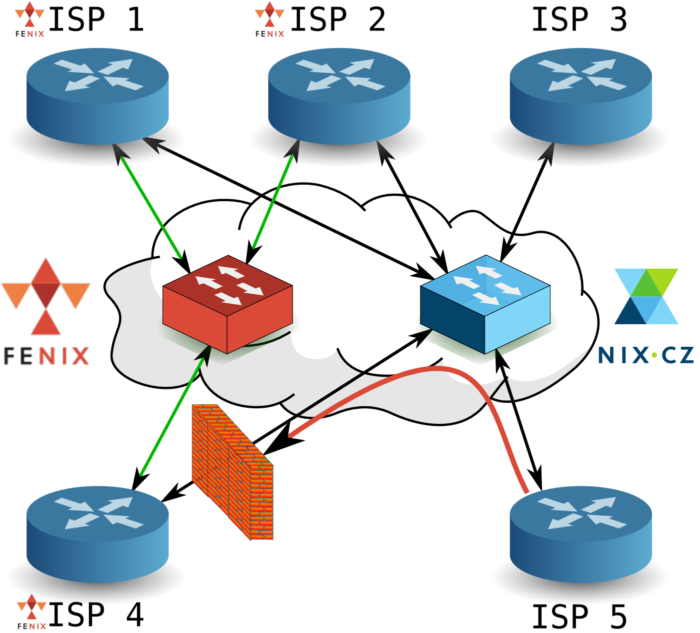
FENIX
ISP 1
FENIX
ISP 2
ISP 3
FENIX
NIX·CZ
FENIX
ISP 4
ISP 5

Key concepts of FENIX VLAN

- only prefixes guaranteed to be clean of spoofing can be announced into FENIX VLAN
- public peering VLAN used for everything by default
- once a FENIX member decides to switch to island mode, they start attracting trafic from other FENIX members via FENIX VLAN
- public peering VLAN **should not be disconected** otherwise the attack would spill over to transit connectivity
- malicious traffic could be blackholed or sent to a scrubber/filter device

CESNET mission in FENIX

- we believe in FENIX principles
  - which brings benefits to **every single network**
- we are pushing our clients to adopt similar rules
  - IP spoofing protection – do not rely on upstream to do the filtering
  - amplification attack protection
  - incident handling
- we do our best **not to source** or support any attack
  - as we could be dangerous to other networks
- we offer tools for monitoring clients’ networks – **Security Tools as a Service**

NREN specifics

- very well provisioned backbone
- big variation of legitimate traffic
- **no filtering by default**[1]

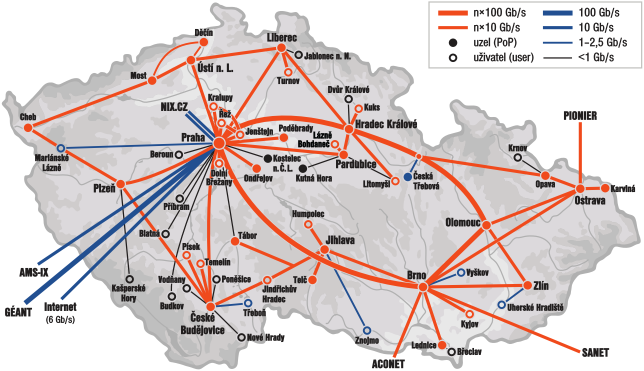

CESNET

[1] unless required (BCP38) or requested by client

DoS experience in CESNET

- client router announces /16 but only /17 is routed
  - packets to remaining /17 ping-pongs between routers
  - last mile link saturated
- received UDP floods from transit can saturate target’s 10Gbps link

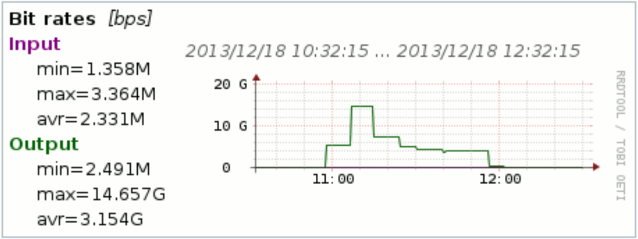

Bit rates *[bps]*

Input
min=1.358M
max=3.364M
avr=2.331M

Output
min=2.491M
max=14.657G
avr=3.154G

Mitigation strategies in CESNET

- RTBH for clients
  - attacks targetted to small number of IP addresses
  - flowspec-based RTBH in development
- per-protocol QoS on the network perimeter
  - for connection-less protocols like NTP, SNMP,...
  - sum of NTP flows typical ∼2 Mbps
  - different packet sizes of legitimate and attack flows
- DNS QoS on the inner-egde of the core network
  - crucial service for *eyeball* experience
  - hard to recognize attack on the perimeter
  - filtering UDP packet without either port 53

Conclusion

- fallback to FENIX VLAN is the very last resort
  - a lot of things will break down
  - but at least *something will work*
- FENIX membership itself *very* useful
  - tighten the community
  - consensual view
  - mutual help and assistence
  - **personal trust**
- higher standards make networks **more reliable**
  - avoids possible goverment regulation
  - making the whole industry a better place

Thank You!

Ondřej Caletka
Ondrej.Caletka@cesnet.cz
https://Ondřej.Caletka.cz

https://www.ces.net
http://fe.nix.cz/en