

DDoS attacks in CESNET2

Ondřej Caletka



15th March 2016



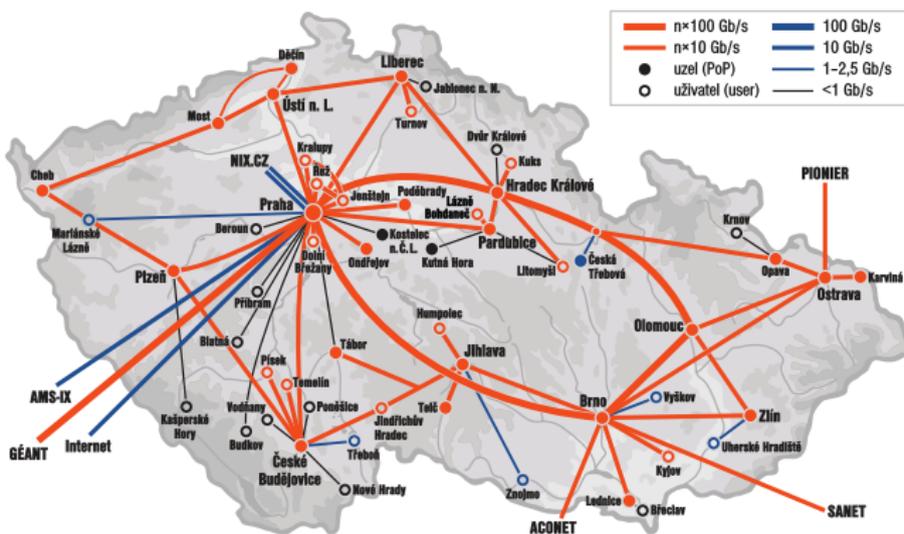
About CESNET

- association of legal entities, est. 1996
 - public and state universities
 - Academy of Sciences
- non-profit organisation
 - development and operation of **NREN** (CESNET2)
 - 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**



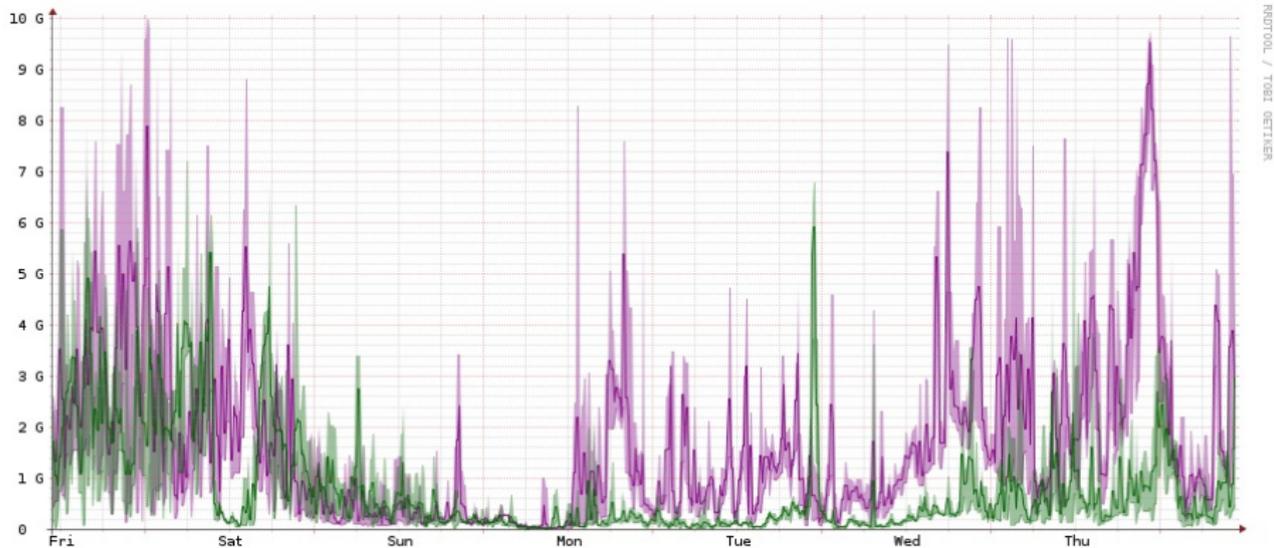
NREN specifics

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



¹unless required (BCP 38) or requested by client

Typical weekly traffic variation



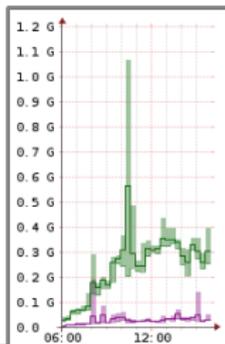
DoS as a phenomenon

- hobby of today's teenagers
 - spending allowances for DDoS-as-a-service
 - targeting classmates' internet services
 - especially gaming and TeamSpeak servers
 - big issue especially for cheap VPS providers
- hacktivism
 - DoSing unpopular services
 - possible target later this year in CZ:
on-line POS sales records collection
- shorter attack times (less than five minutes)
 - often undetected by monitoring tools
 - can break badly designed services
 - bad eyeball experience

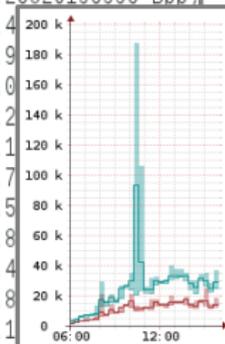
Short-lived UDP flood example

```
Notification      : UDP from external networks and port 0,53,123,161 to internal IPS,
bytes>=1024, targets - DETECTED traffic anomaly
Detected          : 195.113. [REDACTED] (dest. IP) - found 878 (limit 250) flows within
period of 5 seconds

Flows time range [GMT]      : 15/12/17 09:33:02-15/12/17 09:34:03
Flows time range [local]   : 15/12/17 10:33:02-15/12/17 10:34:03
```



50.241.253.129	udp(17)/0	--> 195.113.[REDACTED]	udp(17)/0	: 6113015 B,	5436 p,	1124.54286239882 Bpp,
186.215.207.138	udp(17)/0	--> 195.113.[REDACTED]	udp(17)/0	: 4732764 B,	4694 p,	1008.25820195995 Bpp,
46.188.59.141	udp(17)/0	--> 195.113.[REDACTED]	udp(17)/0	: 1576588 B,	1330 p,	1185.4 Bpp,
61.85.1.79	udp(17)/0	--> 195.113.[REDACTED]	udp(17)/0	: 3466937 B,	3132 p,	1106.9 Bpp,
202.114.238.116	udp(17)/0	--> 195.113.[REDACTED]	udp(17)/0	: 13477647 B,	11460 p,	1176.0 Bpp,
117.102.65.174	udp(17)/0	--> 195.113.[REDACTED]	udp(17)/0	: 9227596 B,	7946 p,	1161.2 Bpp,
217.128.111.66	udp(17)/0	--> 195.113.[REDACTED]	udp(17)/0	: 367793 B,	331 p,	1111.1 Bpp,
219.222.224.6	udp(17)/0	--> 195.113.[REDACTED]	udp(17)/0	: 13262153 B,	11194 p,	1184.7 Bpp,
163.29.216.61	udp(17)/0	--> 195.113.[REDACTED]	udp(17)/0	: 4589202 B,	4038 p,	1136.5 Bpp,
96.90.226.11	udp(17)/0	--> 195.113.[REDACTED]	udp(17)/0	: 2143761 B,	2038 p,	1051.8 Bpp,
112.160.7.249	udp(17)/0	--> 195.113.[REDACTED]	udp(17)/0	: 952083 B,	819 p,	1162.4 Bpp,
219.223.18.20	udp(17)/0	--> 195.113.[REDACTED]	udp(17)/0	: 13569874 B,	11424 p,	1187.8 Bpp,
123.57.177.192	udp(17)/0	--> 195.113.[REDACTED]	udp(17)/0	: 2836849 B,	2499 p,	1135.1 Bpp,
218.244.142.146	udp(17)/0	--> 195.113.[REDACTED]	udp(17)/0	: 749959 B,	688 p,	1090.05668604651 Bpp,
24.111.41.107	udp(17)/0	--> 195.113.[REDACTED]	udp(17)/0	: 5840435 B,	5325 p,	1096.79530516432 Bpp,
187.141.38.238	udp(17)/0	--> 195.113.[REDACTED]	udp(17)/0	: 620221 B,	479 p,	1294.82463465553 Bpp,
60.167.222.146	udp(17)/0	--> 195.113.[REDACTED]	udp(17)/0	: 9041951 B,	7314 p,	1156.12125504195 Bpp,



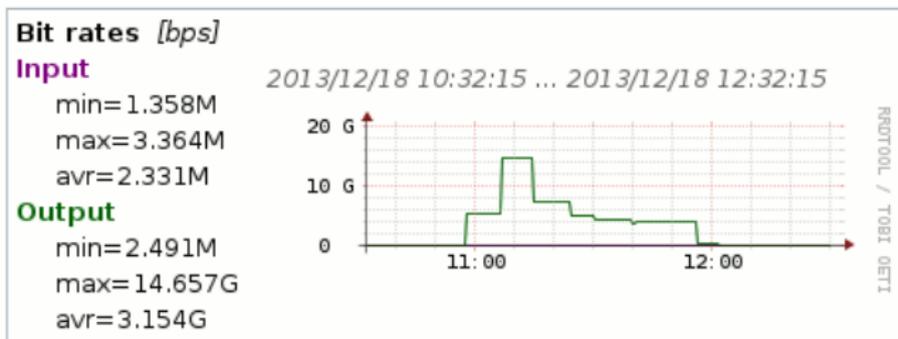
- using connection-less protocol and IP source spoofing for reflection/amplification attack
 - TCP SYN flood
 - DNS
 - NTP
 - SNMP
 - SSDP
- DNS random subdomain queries
 - using spoofed source IP or botnet and open resolvers
 - targets authoritative DNS servers
 - eating resources on recursive DNS servers

Why only few operators deploy BCP 38?

- the closer to the edge the simpler to deploy
- simple automatic urpf - checks don't work well with multihoming
- network equipment vendors still don't offer an easy to deploy solution for multihomed clients (Feasible Reverse Path Forwarding - BCP 84)
- loose RPF has no use against most spoofing attacks
- our solution: ACLs on customer ports
 - managed manually
 - prone to operator errors
 - probably too resource demanding for many ISPs

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



Mitigation strategies in CESNET

- RTBH for clients
 - attacks targetted to small number of IP addresses
 - without RTBH, the last mile link is easily saturated with malicious traffic
 - BGP 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-edge of the core network
 - crucial service for *eyeball* experience
 - hard to recognize attack on the perimeter
 - filtering UDP packets without either port 53



The FENIX project of NIX.CZ

- response to DoS attacks to major websites in 2013
- attack sourced from transit carrier RETN via NIX.CZ
- no technical nor organisational countermeasures available inside the IXP at that time

Idea of secure peering VLAN inside NIX.CZ

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



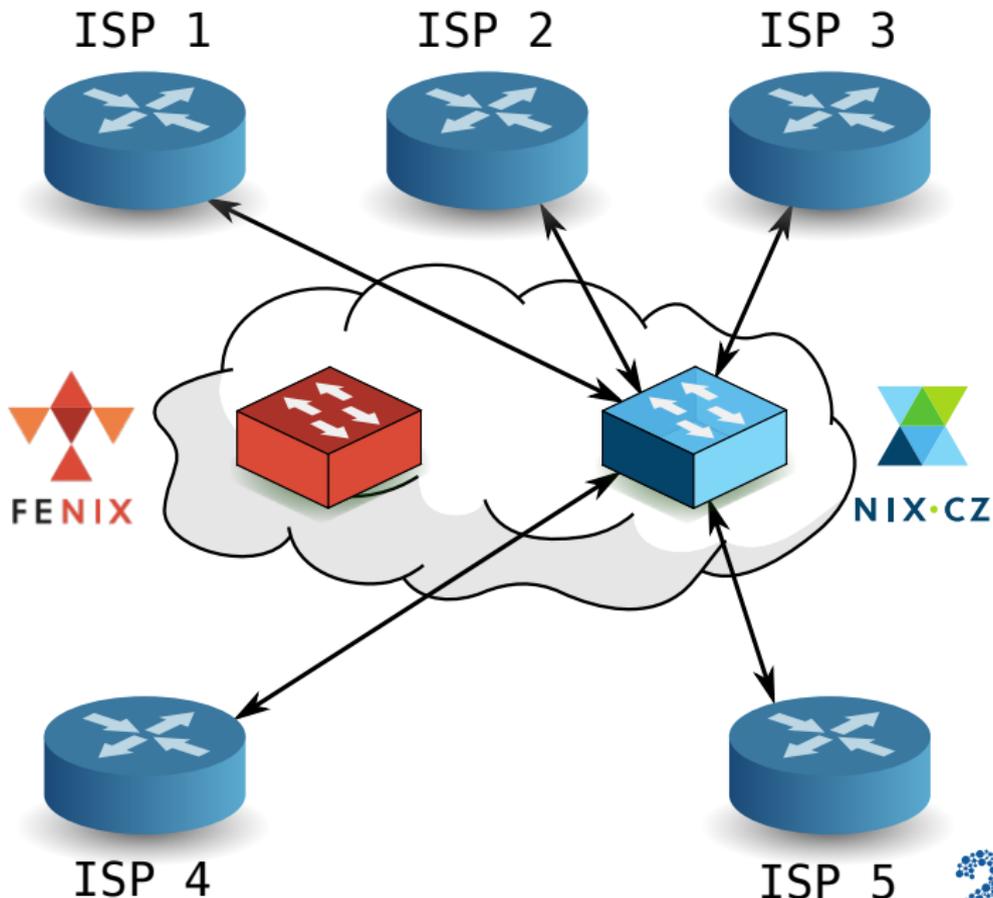
- founded by 6 operators in January 2014
 - **Active 24** (hosting)
 - **CESNET** (NREN)
 - **CZ.NIC** (TLD operator)
 - **Dial Telecom** (ISP)
 - **O2 CZ** (ISP, incumbent)
 - **Seznam.cz** (Czech Google)
- 13 members today

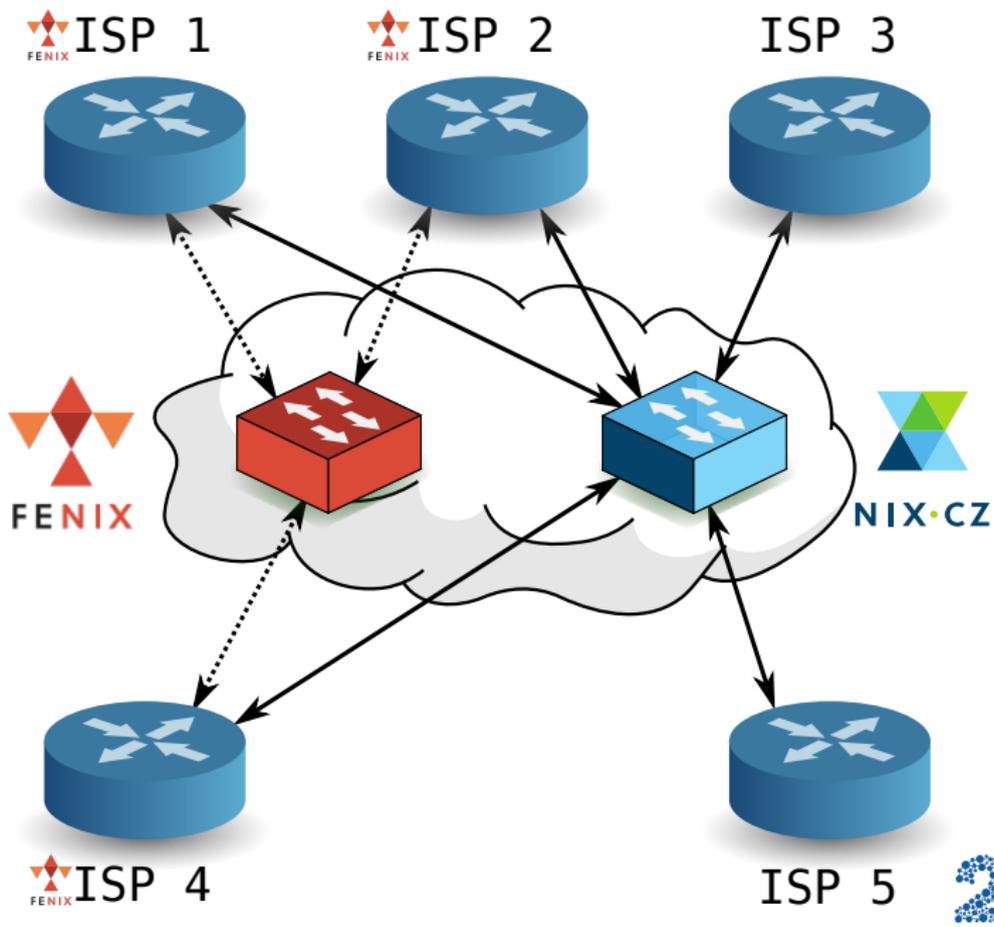
FENIX criteria

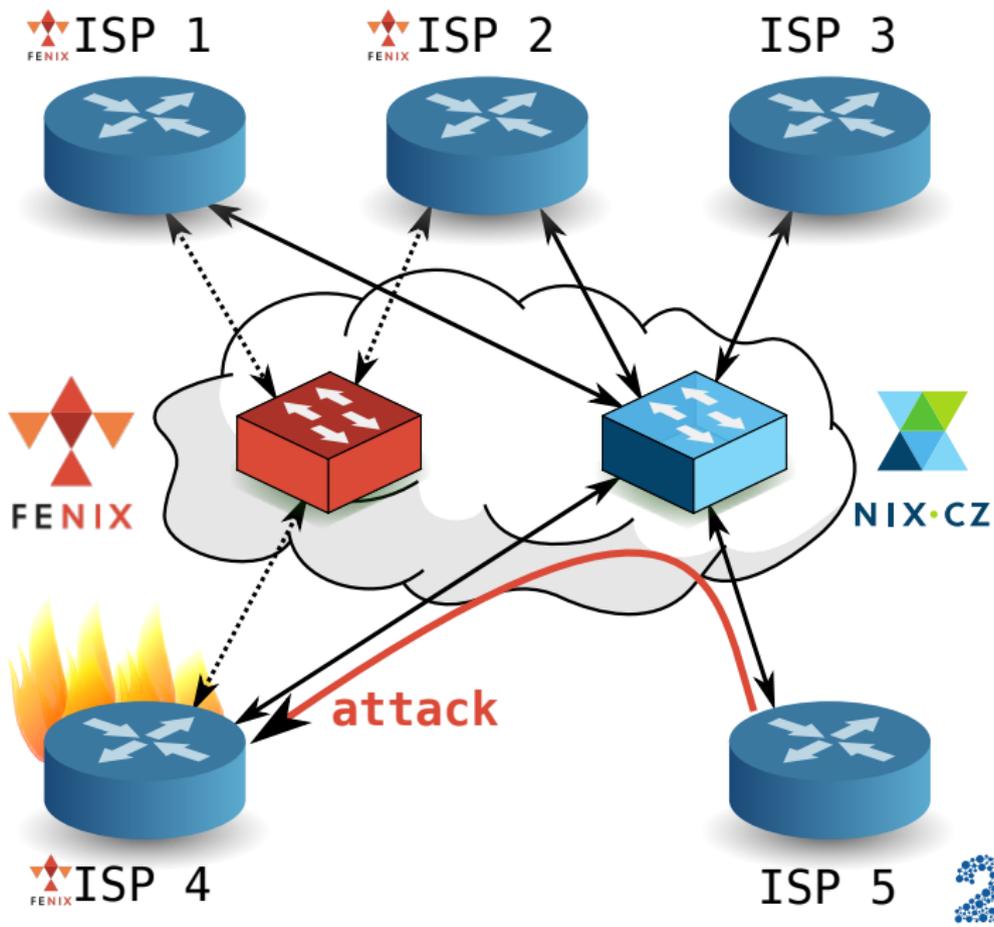
- Terms and Conditions allowing to disconnect customer originating malicious traffic
- 24×7 NOC, no Interactive Voice Response machine
- Trusted Introducer listed CSIRT team
- recommendation from 2 FENIX members, no veto
- **BCP 38**/SSAC 004 network ingress filtering
- RTBH using route servers
- fully redundant connection to NIX.CZ
- protected BGP sessions with TCP MD5
- DNS, NTP, SNMP amplification protection
- deployed IPv6 and DNSSEC

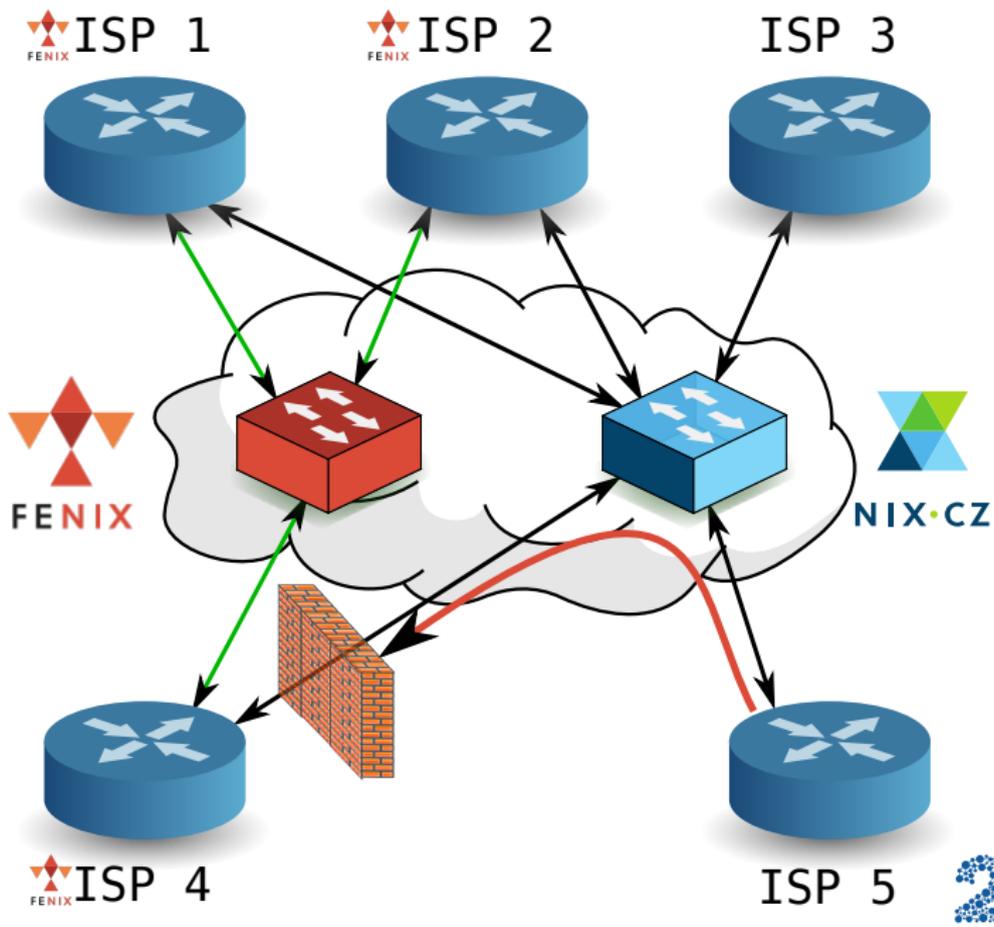
Secure peering VLAN

- 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









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 disconnect all other peerings – traffic flows only from/to other FENIX members via the FENIX VLAN
- instead of disconnecting, malicious traffic could be blackholed or sent to a scrubber/filter device

- 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

Conclusion

- DDoSes are more and more common
 - shorter attack times make them harder to mitigate
 - the future is probably in *automatic DDoS mitigation*
- FENIX-like communities very useful
 - consensual view
 - mutual help and assistance
 - sharing best practices
 - **personal trust**
- higher standards make networks **more reliable**
 - avoids possible government regulation
 - making the whole industry a better place

Thank You!

Ondřej Caletka

Ondrej.Caletka@cesnet.cz

<https://Ondrej.Caletka.cz>



<https://www.ces.net>

<http://fe.nix.cz/en>