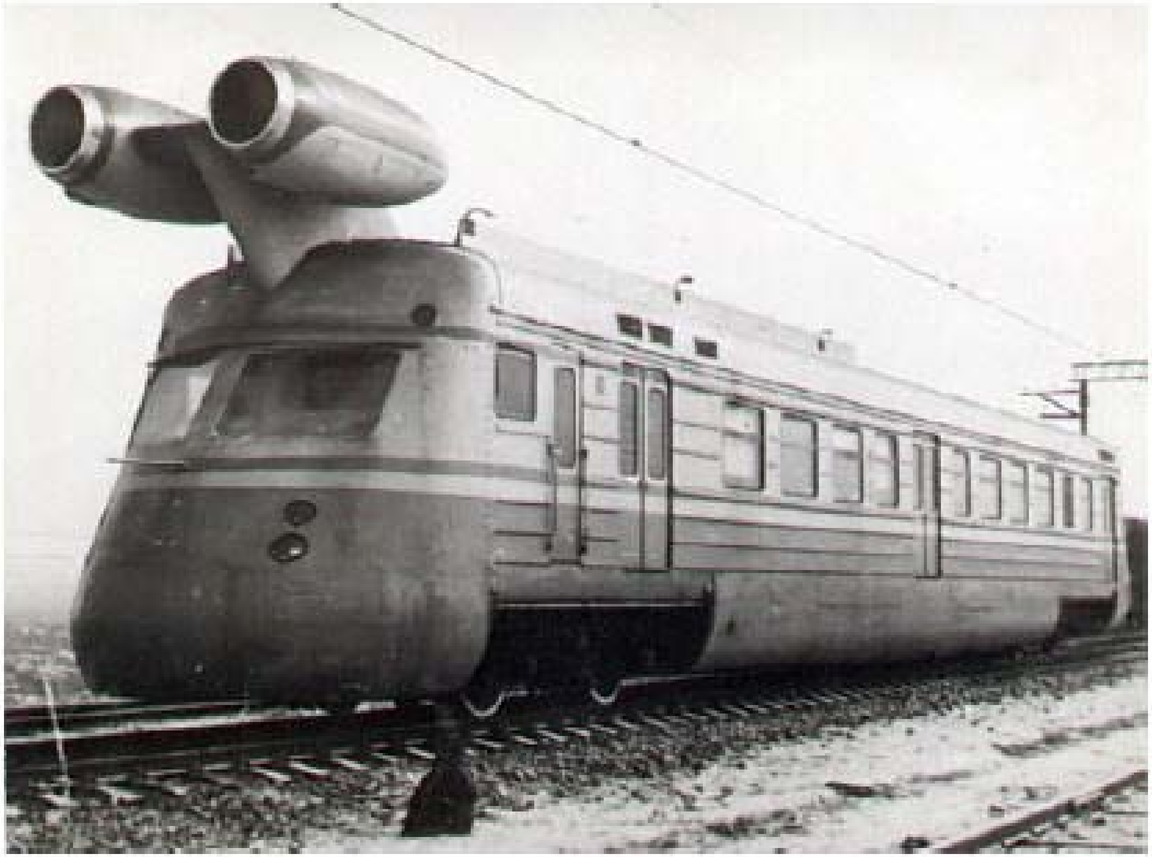

Deploying IPv6 - practical problems from the campus perspective

Tomáš Podermański, tpoder@cis.vutbr.cz
Matěj Grégr, igregr@fit.vutbr.cz

The Brno University of Technology

- http://www.vutbr.cz
- One of the largest universities in the Czech Republic
- Founded in 1899, 110th anniversary was recently celebrated
- 20,000 students and 2,000 employees
- 9 faculties
- 6 other organizational units
- Dormitory for 6,000 students

CVIS

FIT
FACULTY OF INFORMATION TECHNOLOGY

FACULTY OF ELECTRICAL ENGINEERING AND COMMUNICATION

ÚSI
INSTITUTE OF FORENSIC ENGINEERING

FACULTY OF BUSINESS AND MANAGEMENT

FAST
FACULTY OF CIVIL ENGINEERING

FSI
FACULTY OF MECHANICAL ENGINEERING

FCH
FACULTY OF CHEMISTRY

FACULTY OF ARCHITECTURE

FACULTY OF FINE ARTS

Layer 3 network

Core of the network

- Based on 10Gb/s ethernet
- Basic L3 services
- OSPF and OSPFv3
- multicast - PIM/SM

External connectivity

- Two 10Gb/s lines connecting the core to CESNET (BGP, BGP4+)
- Basic filtering (SMTP, NetBios, 445/Microsof DS)

Locality & sub-campuses

- Two 10Gb/s lines to the core
- More complex firewalls configurations are dependend on local administrators

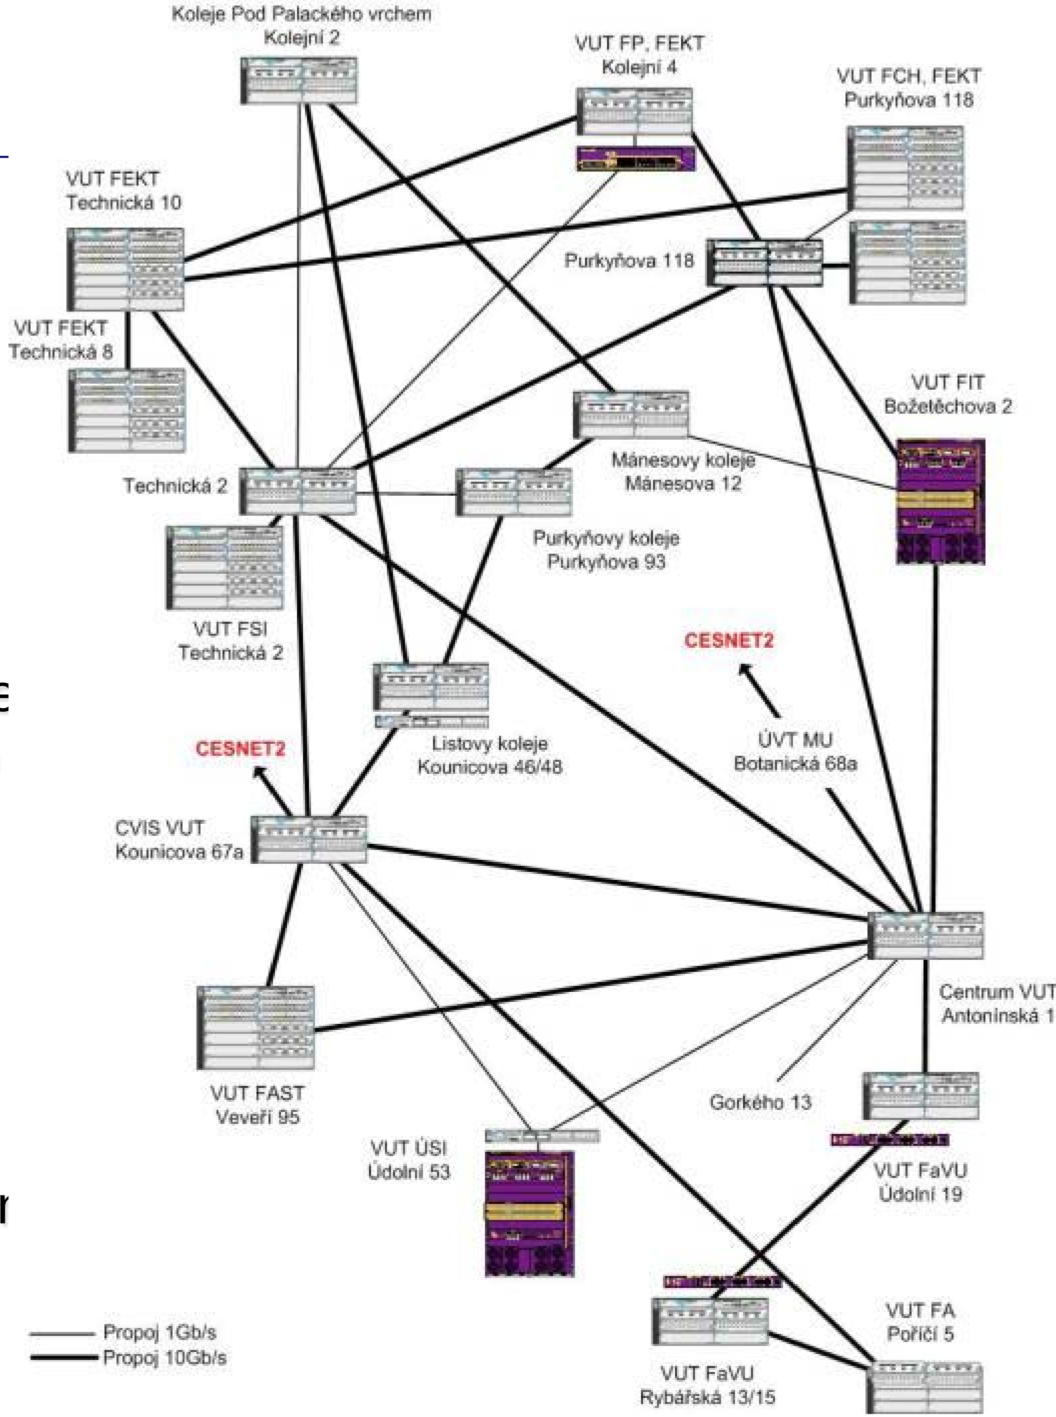

At the beginning (2008) we thought…

- IPv4 address space will be exhausted very soon
  - We have to do something
- IPv6 is here more than 13 years
  - It should be mature enough
- We have been buying devices with IPv6 support for years
  - The hardware is prepared. There is only little configuration needed.
- The Microsoft products starts supporting IPv6 by default. UNIX-like systems supports IPv6 for years
  - The major used platform is prepared

So, where could be a problem ?

Autoconfiguration & address assignment

DHCP, DHCPv6 & SLAAC

	DHCPv6	SLAAC	DHCP (v4)
Handle default route to a client		✔	✔
Handle address of DNS servers to a client	✔	*[1]	✔
Privacy extension or EUI64 address created by a client		✔	
Assignment IP address based on client’s MAC address			✔
Assignment IP address based on client’s DUID address	✔		

- Brand new autoconfiguration mechanisms
  - Router advertisement (doesn’t contain address of DNS servers)
    - [1] There is an extension RFC 6016 but is not widely implemented yet
  - DHCPv6 (doesn’t contain default route option)
    - There is an draft draft-ietf-mif-dhcpv6-route-option-03 but not accepted yet

- Privacy extensions
  - IPv6 addresses are created by hosts randomly
  - IPv6 addresses are periodically changed (every day, once a week)

- DUID in DHCPv6 request instead of MAC address
  - DUID is not predictable created randomly when the OS is installed
  - Impossible to tie address management with existing systems

Autoconfiguration IPv4 x IPv6

- IPv4 – DHCP, ARP

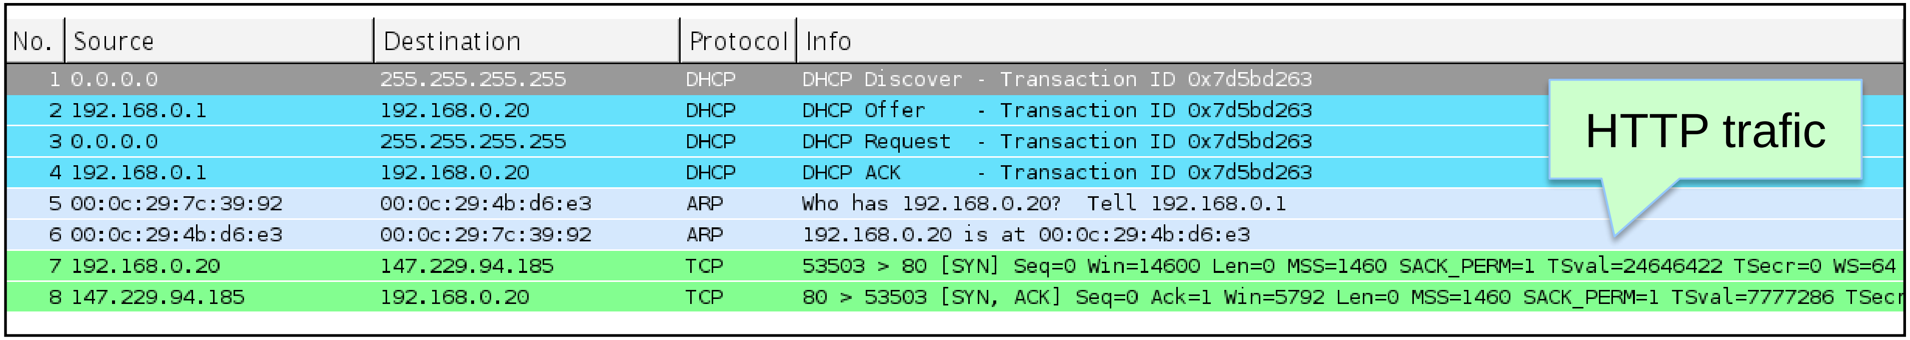

No.	Source	Destination	Protocol	Info
1	0.0.0.0	255.255.255.255	DHCP	DHCP Discover - Transaction ID 0x7d5bd263
2	192.168.0.1	192.168.0.20	DHCP	DHCP Offer - Transaction ID 0x7d5bd263
3	0.0.0.0	255.255.255.255	DHCP	DHCP Request - Transaction ID 0x7d5bd263
4	192.168.0.1	192.168.0.20	DHCP	DHCP ACK - Transaction ID 0x7d5bd263
5	00:0c:29:7c:39:92	00:0c:29:4b:d6:e3	ARP	Who has 192.168.0.20? Tell 192.168.0.1
6	00:0c:29:4b:d6:e3	00:0c:29:7c:39:92	ARP	192.168.0.20 is at 00:0c:29:4b:d6:e3
7	192.168.0.20	147.229.94.185	TCP	53503 > 80 [SYN] Seq=0 Win=14600 Len=0 MSS=1460 SACK_PERM=1 TSval=24646422 TSecr=0 WS=64
8	147.229.94.185	192.168.0.20	TCP	80 > 53503 [SYN, ACK] Seq=0 Ack=1 Win=5792 Len=0 MSS=1460 SACK_PERM=1 TSval=7777286 TSecr

HTTP trafic

- IPv6 – DAD, RS/RA, DHCPv6, MLDv2, ND

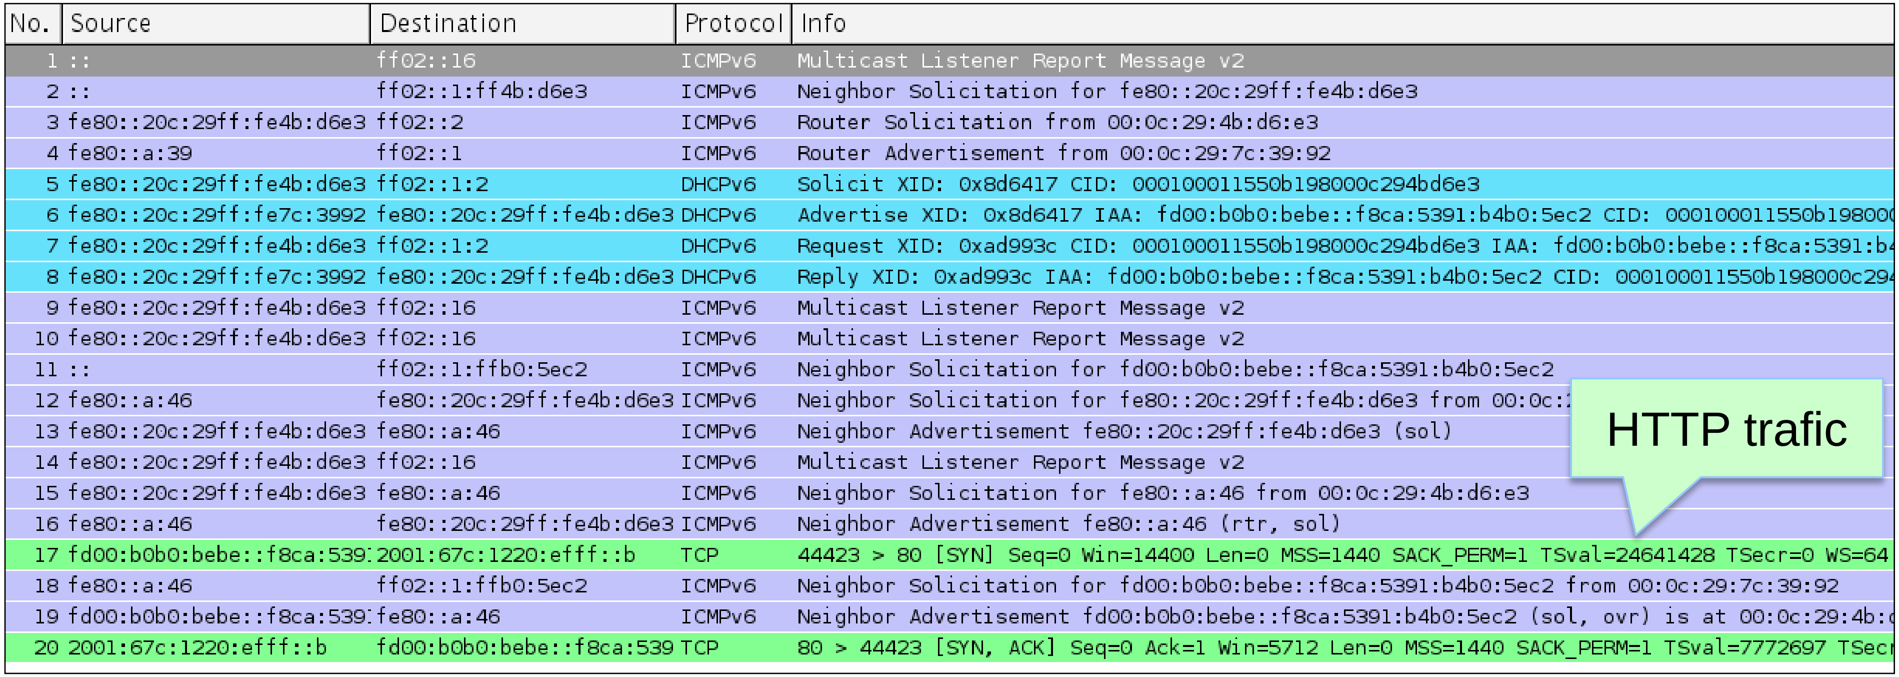

No.	Source	Destination	Protocol	Info
1	::	ff02::16	ICMPv6	Multicast Listener Report Message v2
2	::	ff02::1:ff4b:d6e3	ICMPv6	Neighbor Solicitation for fe80::20c:29ff:fe4b:d6e3
3	fe80::20c:29ff:fe4b:d6e3	ff02::2	ICMPv6	Router Solicitation from 00:0c:29:4b:d6:e3
4	fe80::a:39	ff02::1	ICMPv6	Router Advertisement from 00:0c:29:7c:39:92
5	fe80::20c:29ff:fe4b:d6e3	ff02::1:2	DHCPv6	Solicit XID: 0x8d6417 CID: 000100011550b198000c294bd6e3
6	fe80::20c:29ff:fe7c:3992	fe80::20c:29ff:fe4b:d6e3	DHCPv6	Advertise XID: 0x8d6417 IAA: fd00:b0b0:bebe::f8ca:5391:b4b0:5ec2 CID: 000100011550b19800
7	fe80::20c:29ff:fe4b:d6e3	ff02::1:2	DHCPv6	Request XID: 0xad993c CID: 000100011550b198000c294bd6e3 IAA: fd00:b0b0:bebe::f8ca:5391:b4
8	fe80::20c:29ff:fe7c:3992	fe80::20c:29ff:fe4b:d6e3	DHCPv6	Reply XID: 0xad993c IAA: fd00:b0b0:bebe::f8ca:5391:b4b0:5ec2 CID: 000100011550b198000c294
9	fe80::20c:29ff:fe4b:d6e3	ff02::16	ICMPv6	Multicast Listener Report Message v2
10	fe80::20c:29ff:fe4b:d6e3	ff02::16	ICMPv6	Multicast Listener Report Message v2
11	::	ff02::1:ffb0:5ec2	ICMPv6	Neighbor Solicitation for fd00:b0b0:bebe::f8ca:5391:b4b0:5ec2
12	fe80::a:46	fe80::20c:29ff:fe4b:d6e3	ICMPv6	Neighbor Solicitation for fe80::20c:29ff:fe4b:d6e3 from 00:0c:
13	fe80::20c:29ff:fe4b:d6e3	fe80::a:46	ICMPv6	Neighbor Advertisement fe80::20c:29ff:fe4b:d6e3 (sol)
14	fe80::20c:29ff:fe4b:d6e3	ff02::16	ICMPv6	Multicast Listener Report Message v2
15	fe80::20c:29ff:fe4b:d6e3	fe80::a:46	ICMPv6	Neighbor Solicitation for fe80::a:46 from 00:0c:29:4b:d6:e3
16	fe80::a:46	fe80::20c:29ff:fe4b:d6e3	ICMPv6	Neighbor Advertisement fe80::a:46 (rtr, sol)
17	fd00:b0b0:bebe::f8ca:539	2001:67c:1220:efff::b	TCP	44423 > 80 [SYN] Seq=0 Win=14400 Len=0 MSS=1440 SACK_PERM=1 TSval=24641428 TSecr=0 WS=64
18	fe80::a:46	ff02::1:ffb0:5ec2	ICMPv6	Neighbor Solicitation for fd00:b0b0:bebe::f8ca:5391:b4b0:5ec2 from 00:0c:29:7c:39:92
19	fd00:b0b0:bebe::f8ca:539	fe80::a:46	ICMPv6	Neighbor Advertisement fd00:b0b0:bebe::f8ca:5391:b4b0:5ec2 (sol, ovr) is at 00:0c:29:4b:
20	2001:67c:1220:efff::b	fd00:b0b0:bebe::f8ca:539	TCP	80 > 44423 [SYN, ACK] Seq=0 Ack=1 Win=5712 Len=0 MSS=1440 SACK_PERM=1 TSval=7772697 TSecr

HTTP trafic

Autoconfiguration features

	DHCPv6	SLAAC	RFC 6106	SEND	DHCP (v4)
Windows XP		✔		Only experimental version not ready to use in production enviroment	✔
Windows Vista / 7 / 8	✔	✔			✔
MAC OSX	✔	✔			✔
MAC OSX prior to Lion (2011)		✔			✔
Linux	✔	✔	✔		✔
Android		✔			✔
Windows phone					✔
iOS (iPhone, iPad, iPod)		✔			✔

- Chose SLAAC ?
  - All devices supporting IPv6 will be able to use IPv6
  - Privacy extensions => troubles with tracking, accounting, security incidents and access to privileges services (based on address)
  - How to handle address of DNS servers ?
- Choose DHCPv6 ?
  - Not supported by all platforms => many clients will not have access to IPv6 network
  - Needs to solve problem with DUID and its relation to users' device/identity
  - How to handle default route ?

Catch XXII ?

Privacy extensions enabled ①

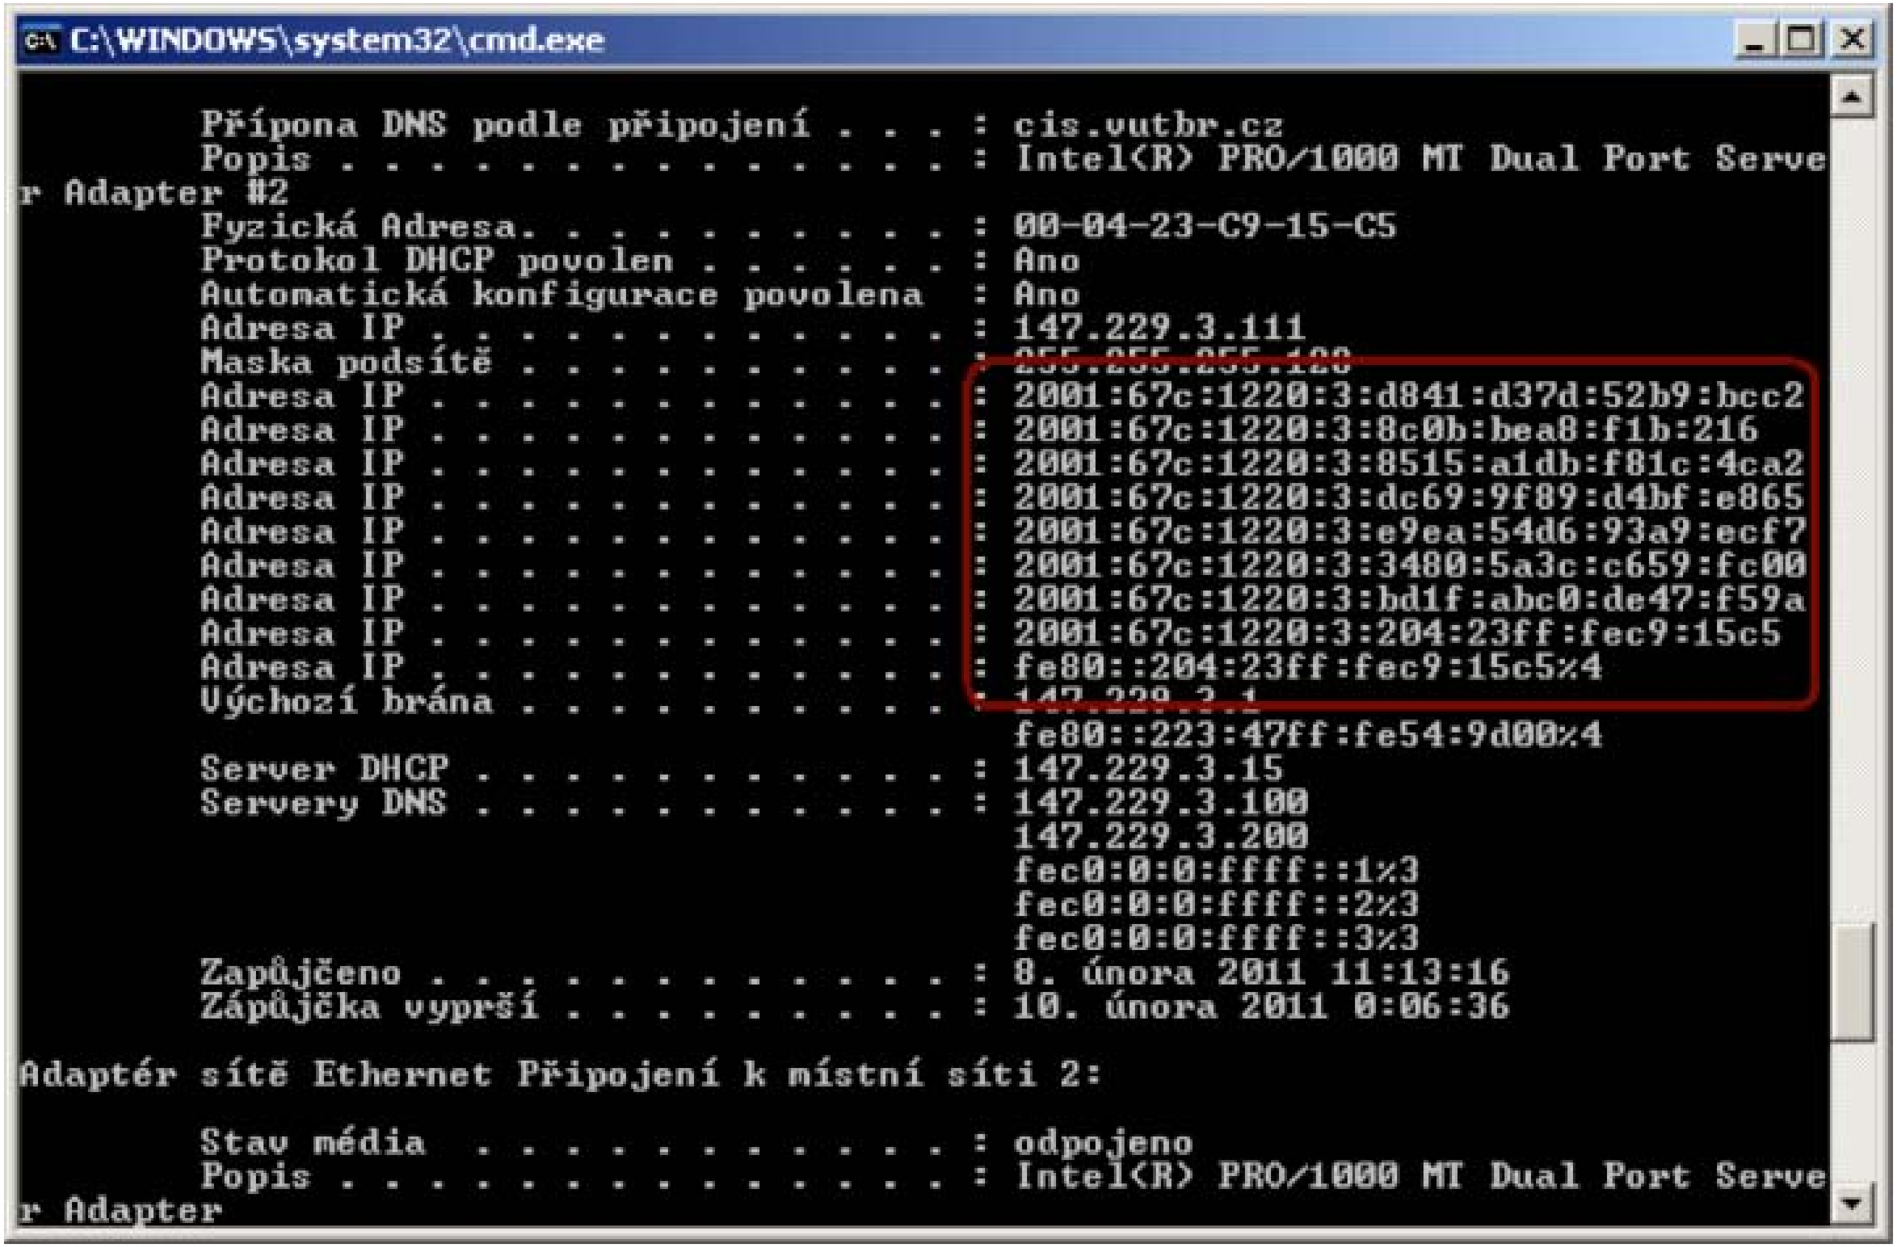

```
        Přípona DNS podle připojení . . . : cis.vutbr.cz
        Popis . . . . . . . . . . . . . . : Intel(R) PRO/1000 MT Dual Port Serve
r Adapter #2
        Fyzická Adresa. . . . . . . . . . : 00-04-23-C9-15-C5
        Protokol DHCP povolen . . . . . . : Ano
        Automatická konfigurace povolena  : Ano
        Adresa IP . . . . . . . . . . . . : 147.229.3.111
        Maska podsítě . . . . . . . . . . : 255.255.255.128
        Adresa IP . . . . . . . . . . . . : 2001:67c:1220:3:d841:d37d:52b9:bcc2
        Adresa IP . . . . . . . . . . . . : 2001:67c:1220:3:8c0b:bea8:f1b:216
        Adresa IP . . . . . . . . . . . . : 2001:67c:1220:3:8515:a1db:f81c:4ca2
        Adresa IP . . . . . . . . . . . . : 2001:67c:1220:3:dc69:9f89:d4bf:e865
        Adresa IP . . . . . . . . . . . . : 2001:67c:1220:3:e9ea:54d6:93a9:ecf7
        Adresa IP . . . . . . . . . . . . : 2001:67c:1220:3:3480:5a3c:c659:fc00
        Adresa IP . . . . . . . . . . . . : 2001:67c:1220:3:bd1f:abc0:de47:f59a
        Adresa IP . . . . . . . . . . . . : 2001:67c:1220:3:204:23ff:fec9:15c5
        Adresa IP . . . . . . . . . . . . : fe80::204:23ff:fec9:15c5%4
        Výchozí brána . . . . . . . . . . : 147.229.3.1
                                            fe80::223:47ff:fe54:9d00%4
        Server DHCP . . . . . . . . . . . : 147.229.3.15
        Servery DNS . . . . . . . . . . . : 147.229.3.100
                                            147.229.3.200
                                            fec0:0:0:ffff::1%3
                                            fec0:0:0:ffff::2%3
                                            fec0:0:0:ffff::3%3
        Zapůjčeno . . . . . . . . . . . . : 8. února 2011 11:13:16
        Zápůjčka vyprší . . . . . . . . . : 10. února 2011 0:06:36

Adaptér sítě Ethernet Připojení k místní síti 2:

        Stav média  . . . . . . . . . . . : odpojeno
        Popis . . . . . . . . . . . . . . : Intel(R) PRO/1000 MT Dual Port Serve
r Adapter
```

Solution: gathering information from the network

- IP address ➔ MAC address: neighbor cache (NC), arp table
- IP address ➔DHCPv6 DUID: Lease db on DHCPv6 server
  - eg. http://metanav.uninett.no/
- MAC address ➔ Switch port: forwarding database (FDB)
- MAC address ➔ Login : radius server
- A system that can collect proper information has to be deployed

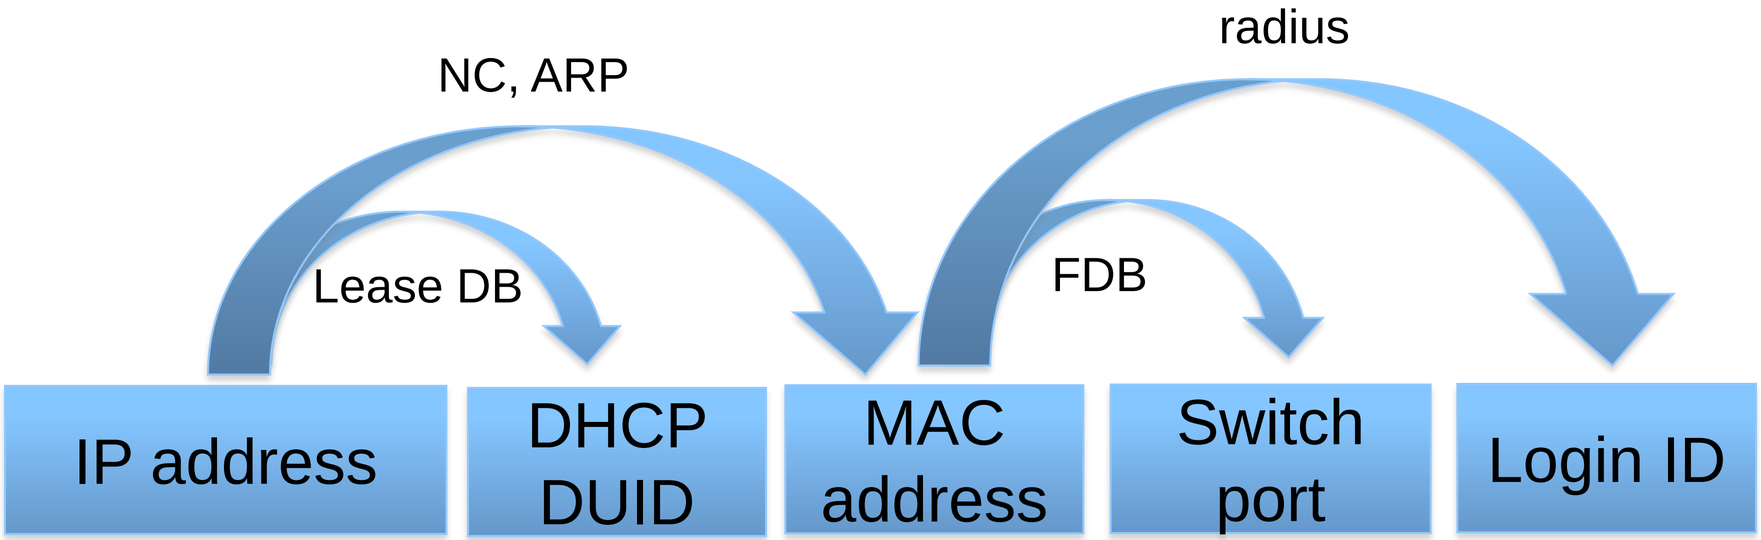

First hop security

First hop security

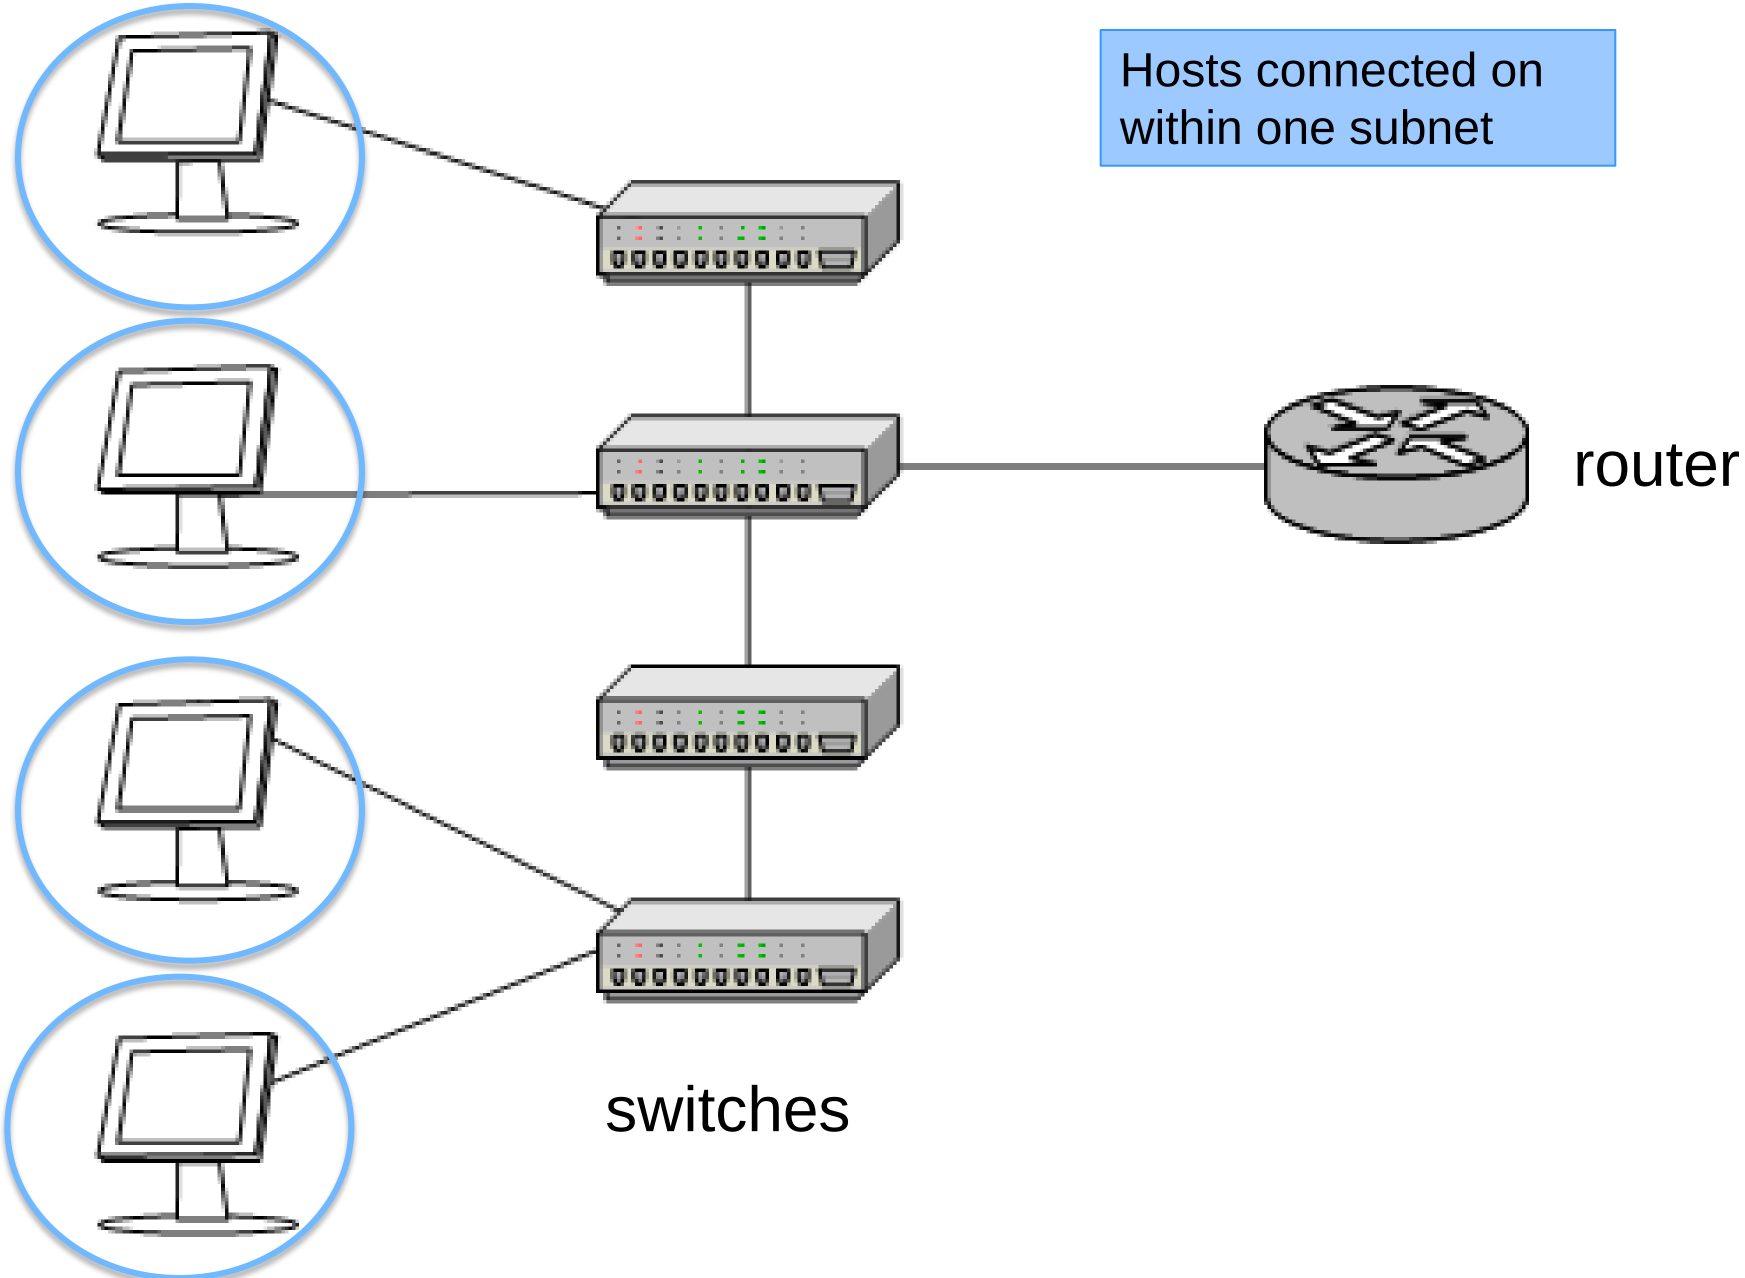

Where and why is first hop security important

- “Stealing” IP address assigned to another device
- Source IP address spoofing
- Source MAC address spoofing
- ARP poisoning
- Rogue DHCP server
  - **Trojan.Flush.M,**
  - **Trojan:W32/DNSChanger**
- The trend is to build and operate “flat” networks
- 2 or 4 C-blocks joined into one subnet

First hop security

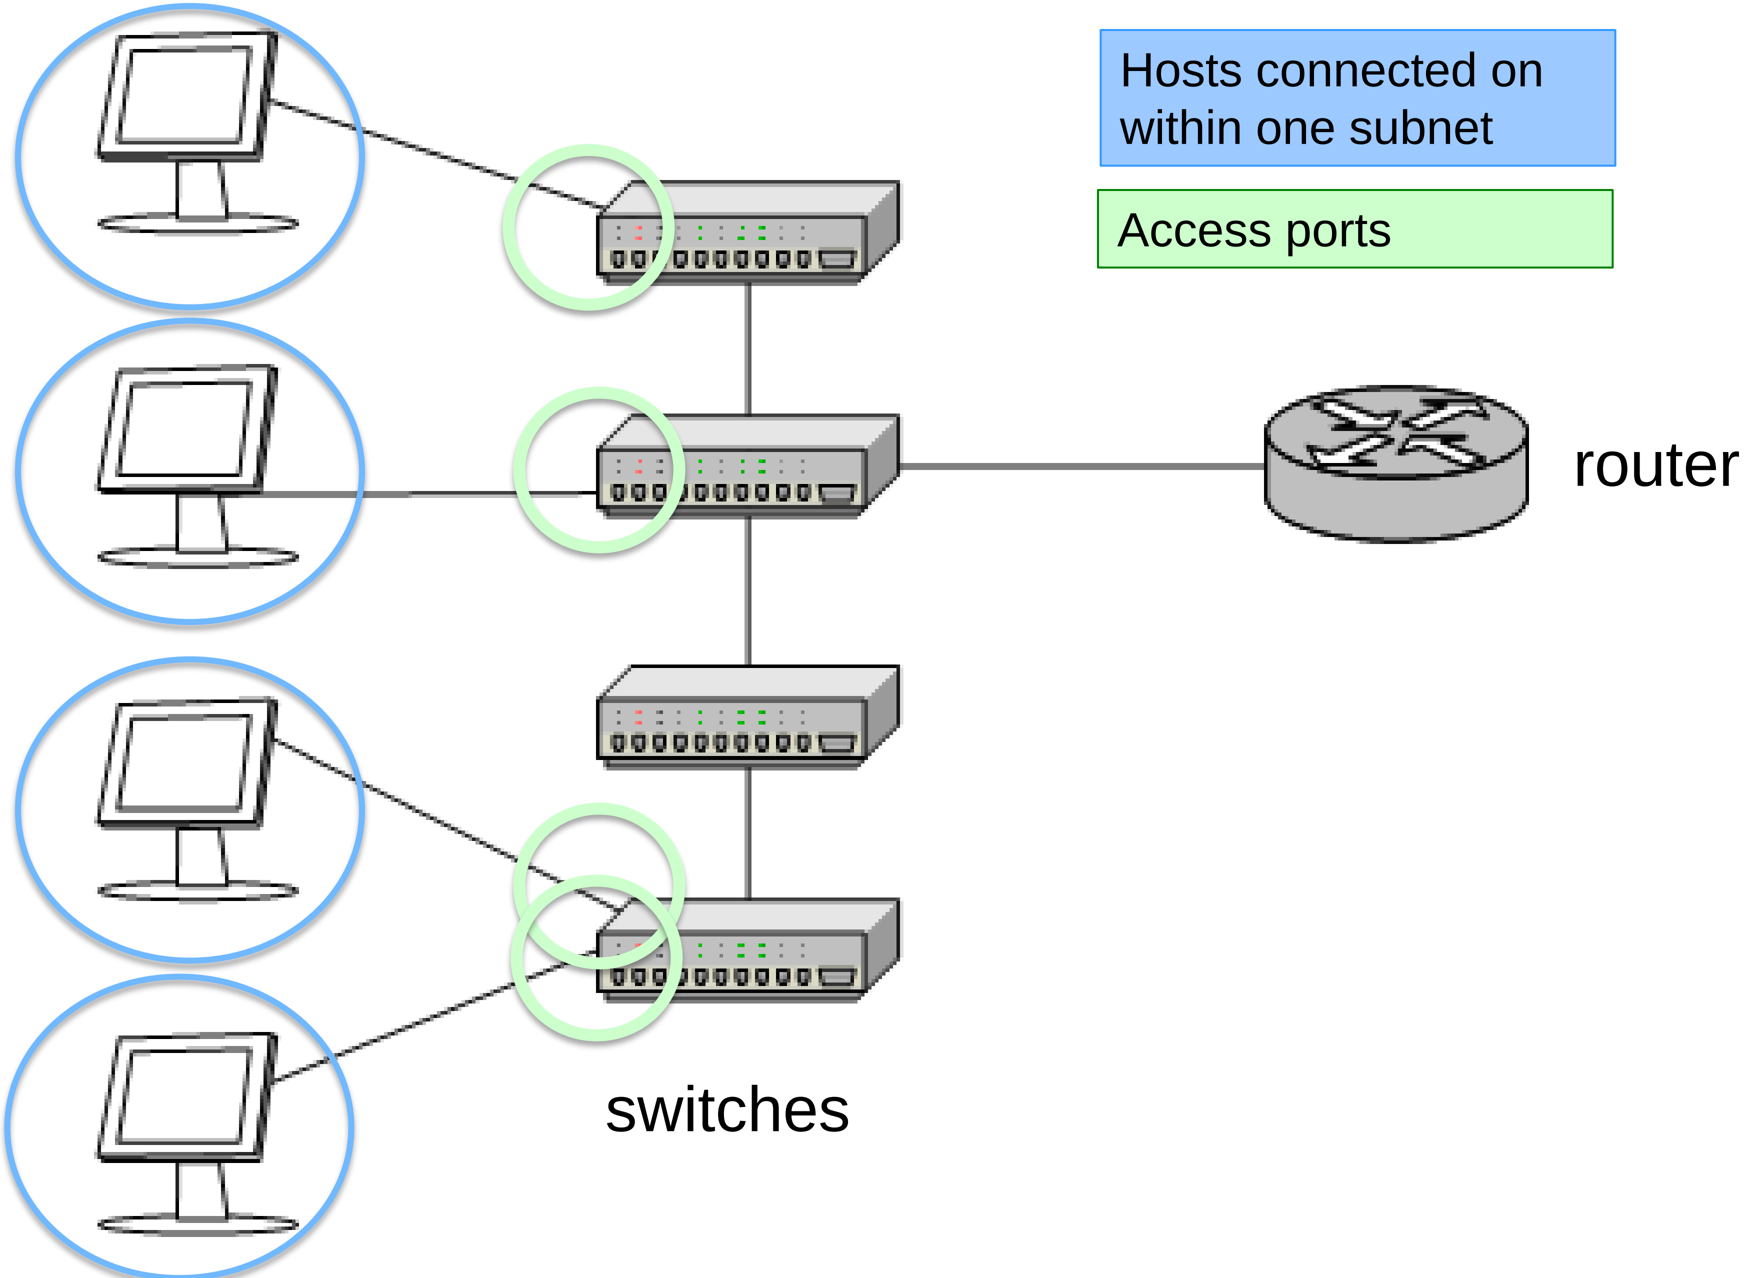

Hosts connected on within one subnet

Access ports

First hop security in IPv4

	DHCP snooping
Rogue DHCP server	✔
ARP poisoning	
Forces users to use DHCP	
Source IPv4 address spoofing	
Source MAC address spoofing	
Require support on access port	✔

- IPv4 autoconfiguration = DHCP + protection on switches
  - **DHCP snooping**
    - **Blocking DHCP answers on access port**
    - **Building binding database (MAC-IP) related to access port**
  - **Dynamic ARP protection/ARP inspection**
    - MAC-IP address database based on DHCP leases
    - Checking content of ARP packets on client access port
  - **Dynamic lock down, IP source guard**
    - The MAC-IP database is used for inspection of client source MAC and IP address.

First hop security threats in IPv6

- Neighbor cache spoofing
  - Very similar to ARP spoofing
  - The spoofed address can be kept in the NC longer
- DoS - Duplicate Address Detection (DAD)
  - Nodes usually create own address (EUI 64, Privacy Extensions)
  - Optimistic DAD – “sorry, the address is mine, choose another one”
- Neighbor Cache table overload
  - Big address space (64 bits – 1.8e+19 address)
  - Many records in the NC for non existing clients
- Rogue Router Advertisement
  - I am a router for this network – use me as a default router
  - The real router is not a valid anymore – zero lifetime
- Rogue DHCPv6 Server
  - I am a DHCPv6 sever for this network. Use my options (DNS)

Rogues IPv6 routers

- Flooding IPv6 hosts with Router Advertises
- thc-toolkit (http://thc.org/thc-ipv6/):

```
# ./flood_router6 eth0
```

- All Windows Vista/7 boxes will freeze
- Other platforms will have problems with IPv6 connectivity
- The problem is known for more than 3 years – no patch/update yet.
- More info:

  http://samsclass.info/ipv6/proj/flood-router6a.htm

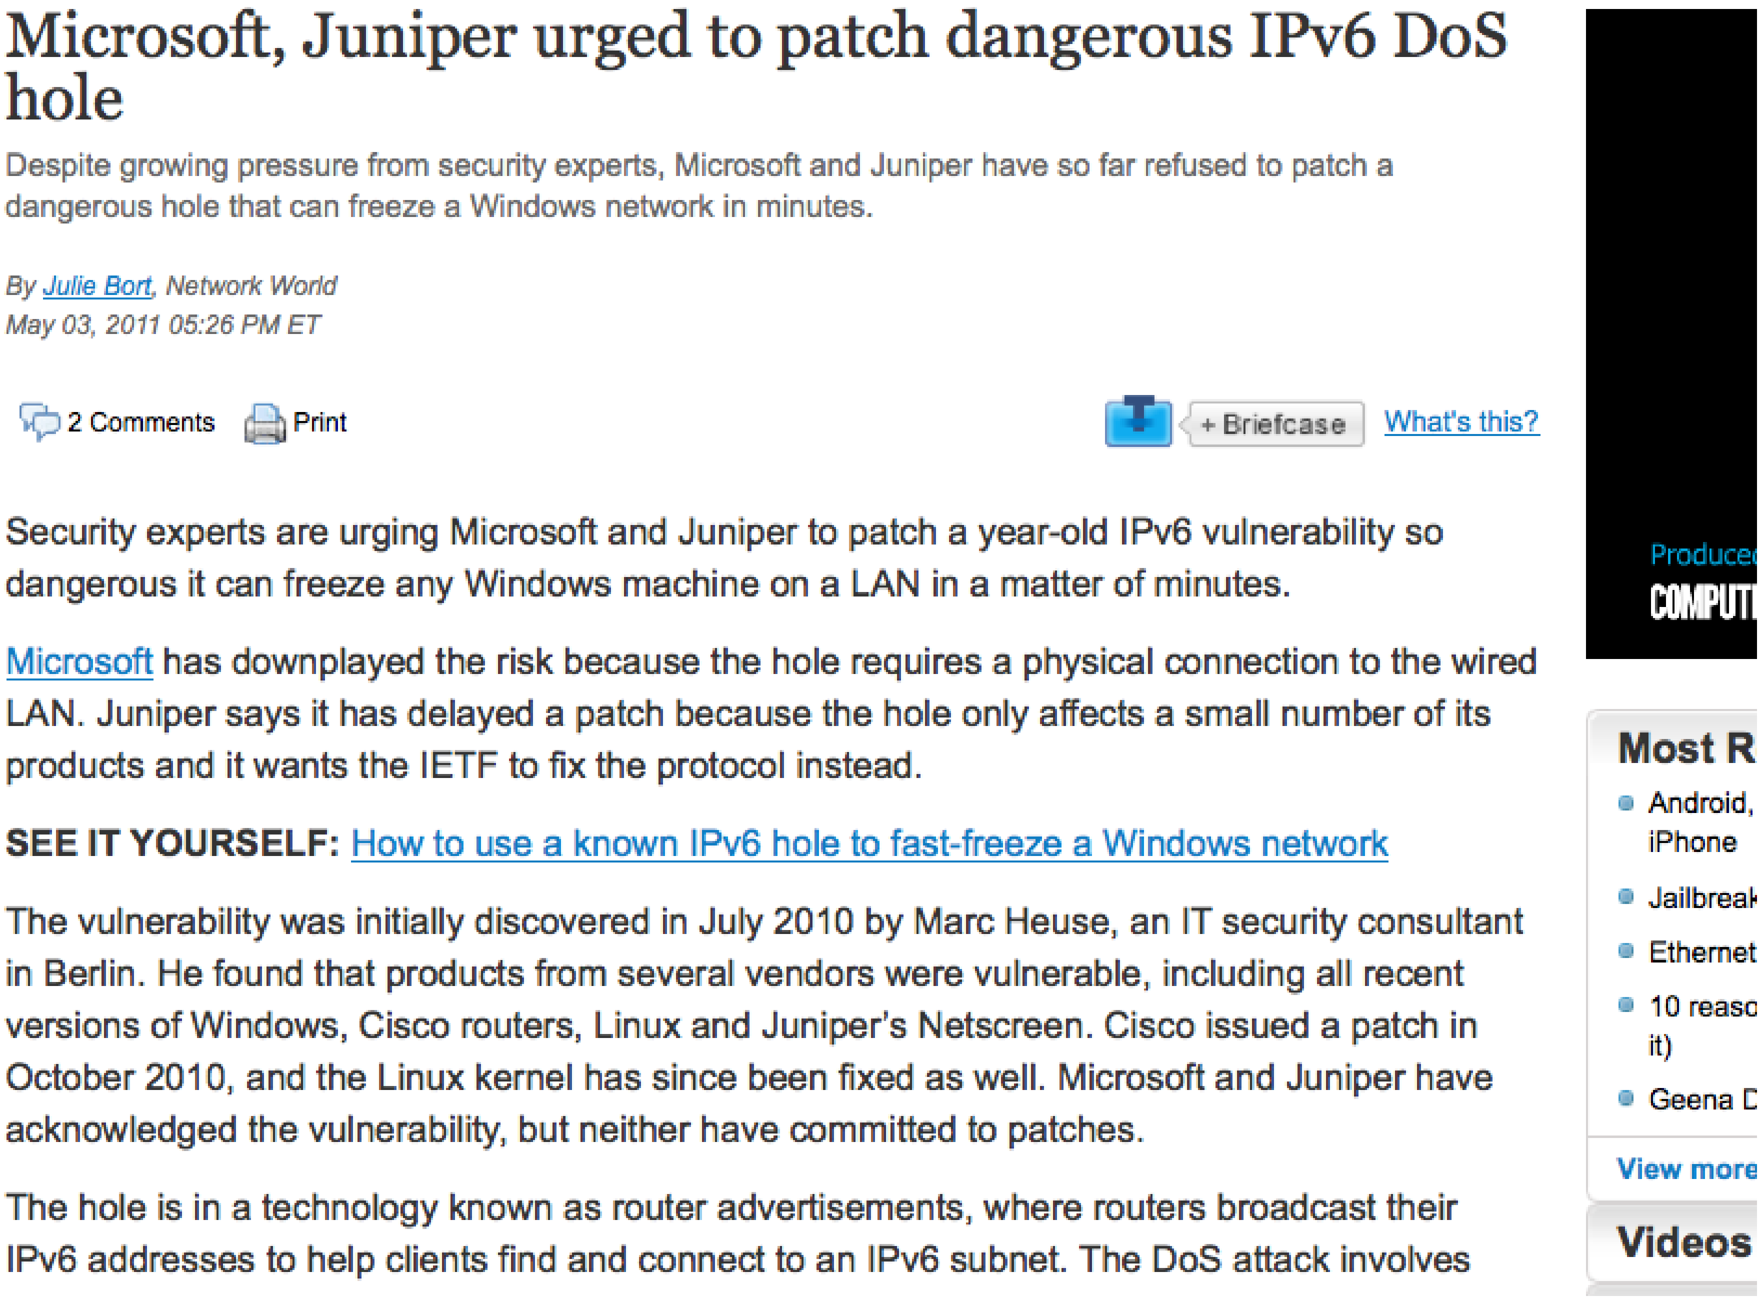

Microsoft, Juniper urged to patch dangerous IPv6 DoS hole

Despite growing pressure from security experts, Microsoft and Juniper have so far refused to patch a dangerous hole that can freeze a Windows network in minutes.

By Julie Bort, Network World
May 03, 2011 05:26 PM ET

Security experts are urging Microsoft and Juniper to patch a year-old IPv6 vulnerability so dangerous it can freeze any Windows machine on a LAN in a matter of minutes.

Microsoft has downplayed the risk because the hole requires a physical connection to the wired LAN. Juniper says it has delayed a patch because the hole only affects a small number of its products and it wants the IETF to fix the protocol instead.

SEE IT YOURSELF: How to use a known IPv6 hole to fast-freeze a Windows network

The vulnerability was initially discovered in July 2010 by Marc Heuse, an IT security consultant in Berlin. He found that products from several vendors were vulnerable, including all recent versions of Windows, Cisco routers, Linux and Juniper's Netscreen. Cisco issued a patch in October 2010, and the Linux kernel has since been fixed as well. Microsoft and Juniper have acknowledged the vulnerability, but neither have committed to patches.

The hole is in a technology known as router advertisements, where routers broadcast their IPv6 addresses to help clients find and connect to an IPv6 subnet. The DoS attack involves

Number of rogues IPv6 routers

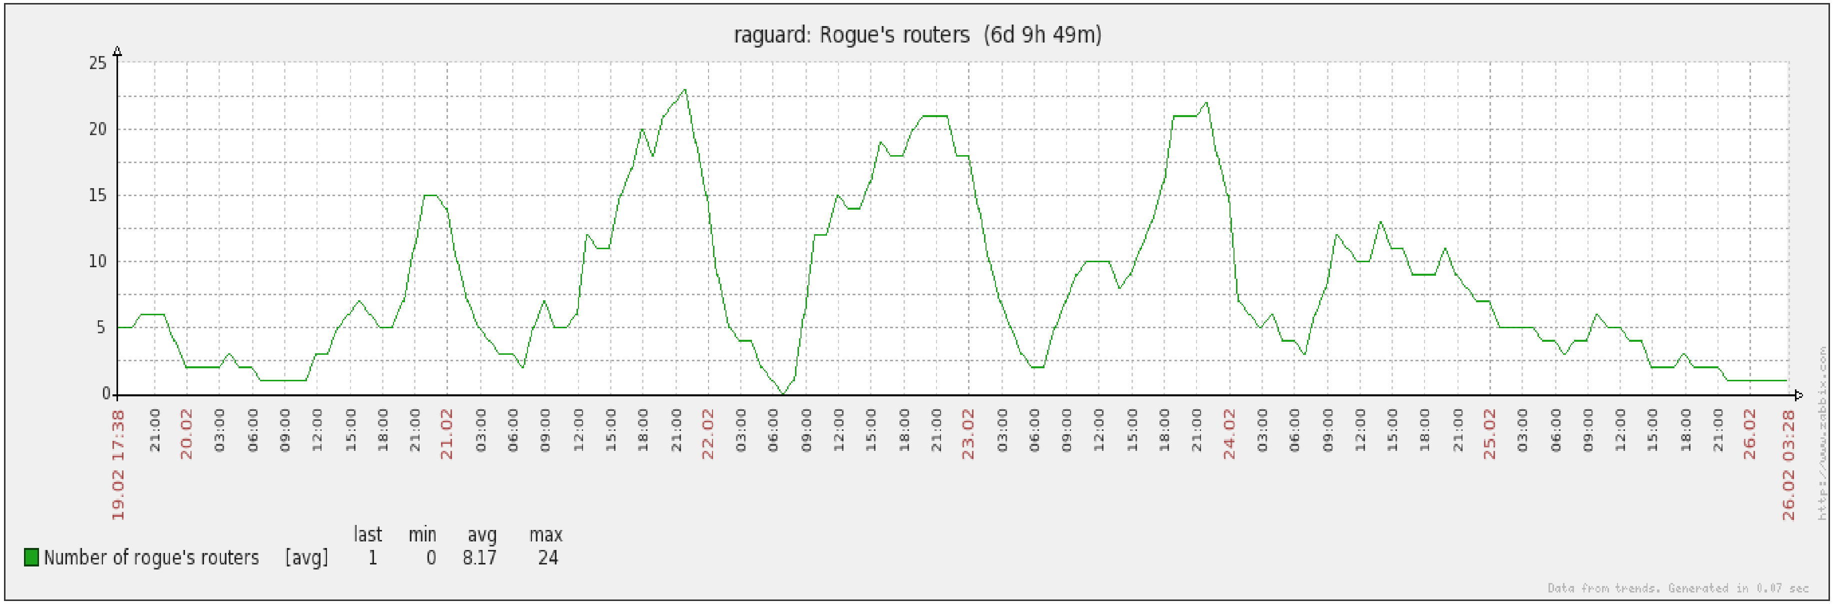

- Usually caused by **Internet Connection Sharing** in Window Vista (Windows 7 treats RA much better)
- Up to 20 rogue routers within network with 4000 hosts
  - Similar (even worse) problem as 20 rogue DHCPv4 servers
- Single host with enabled ICS can destroy IPv6 connectivity for all hosts on the subnet

IPv4 traffic redirection using IPv6

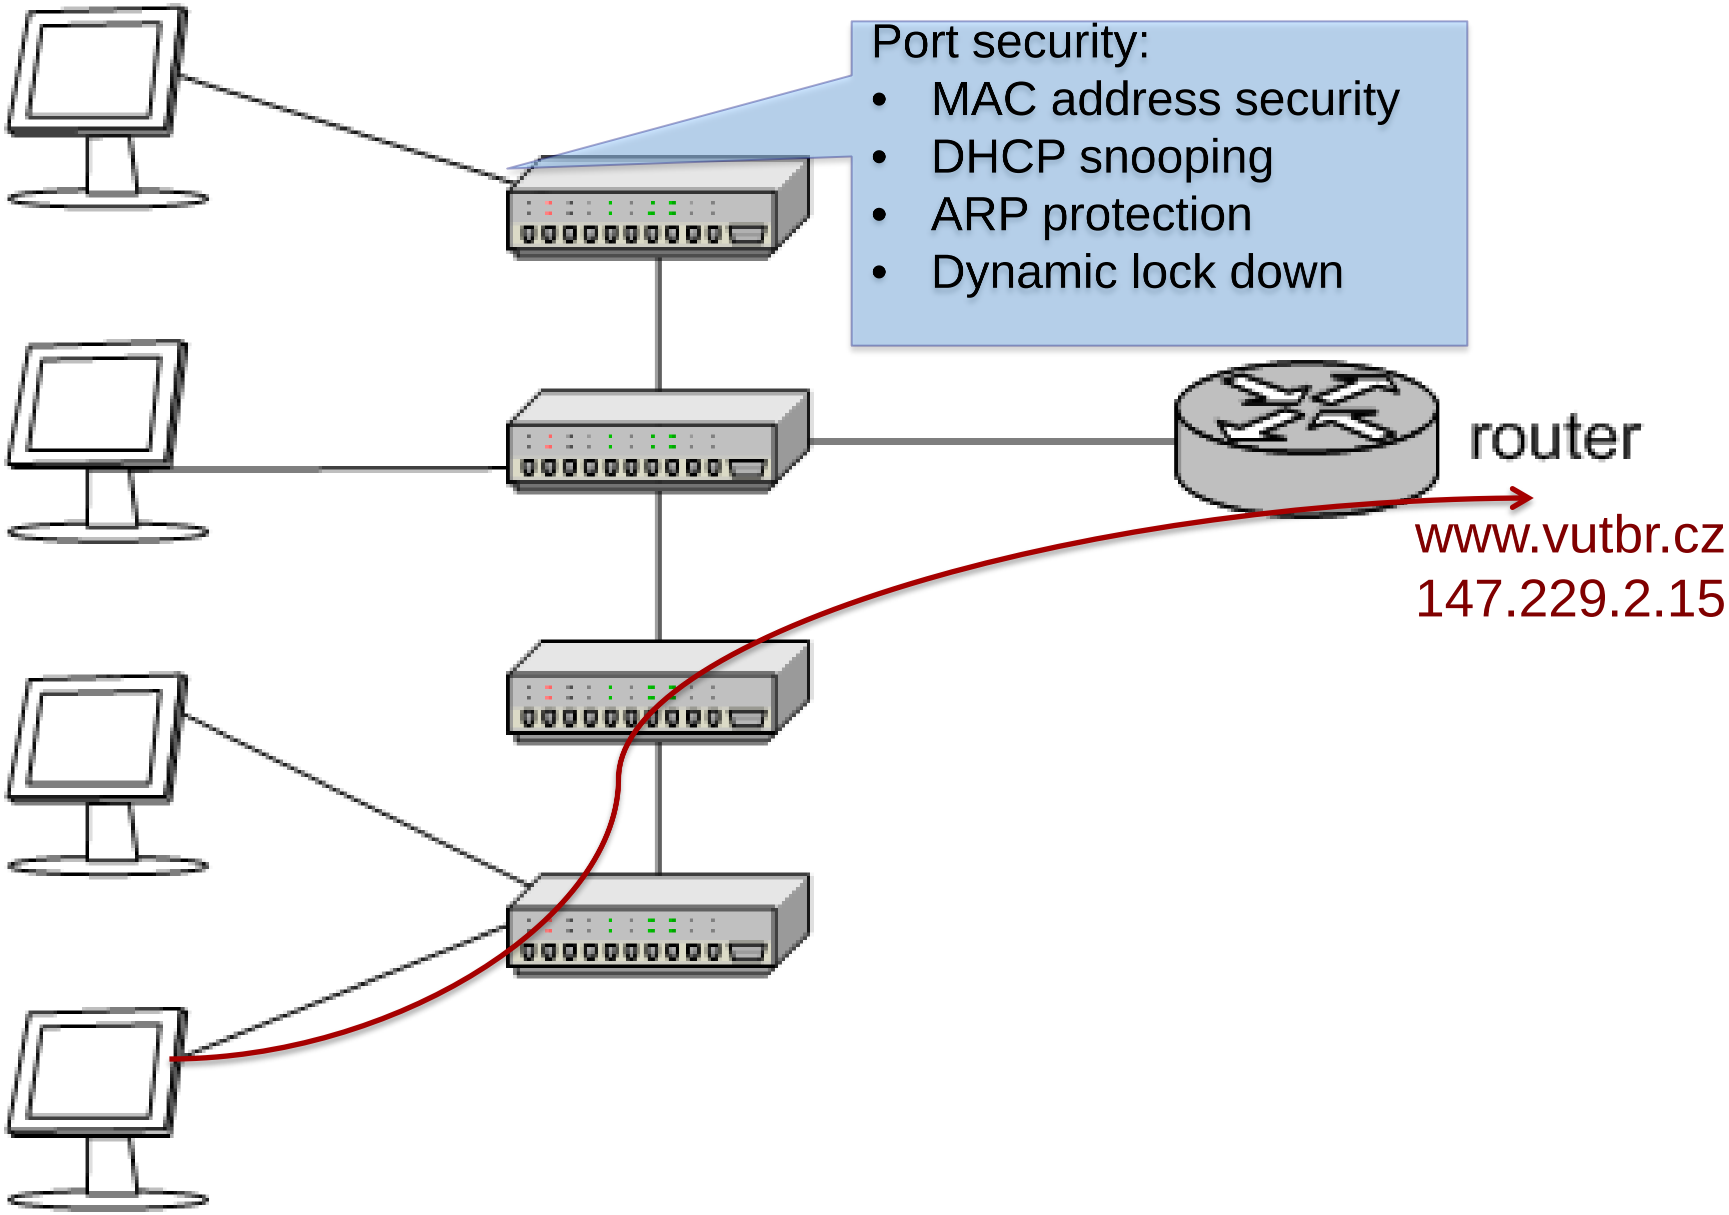

Port security:

- MAC address security
- DHCP snooping
- ARP protection
- Dynamic lock down

IPv4 traffic redirection using IPv6

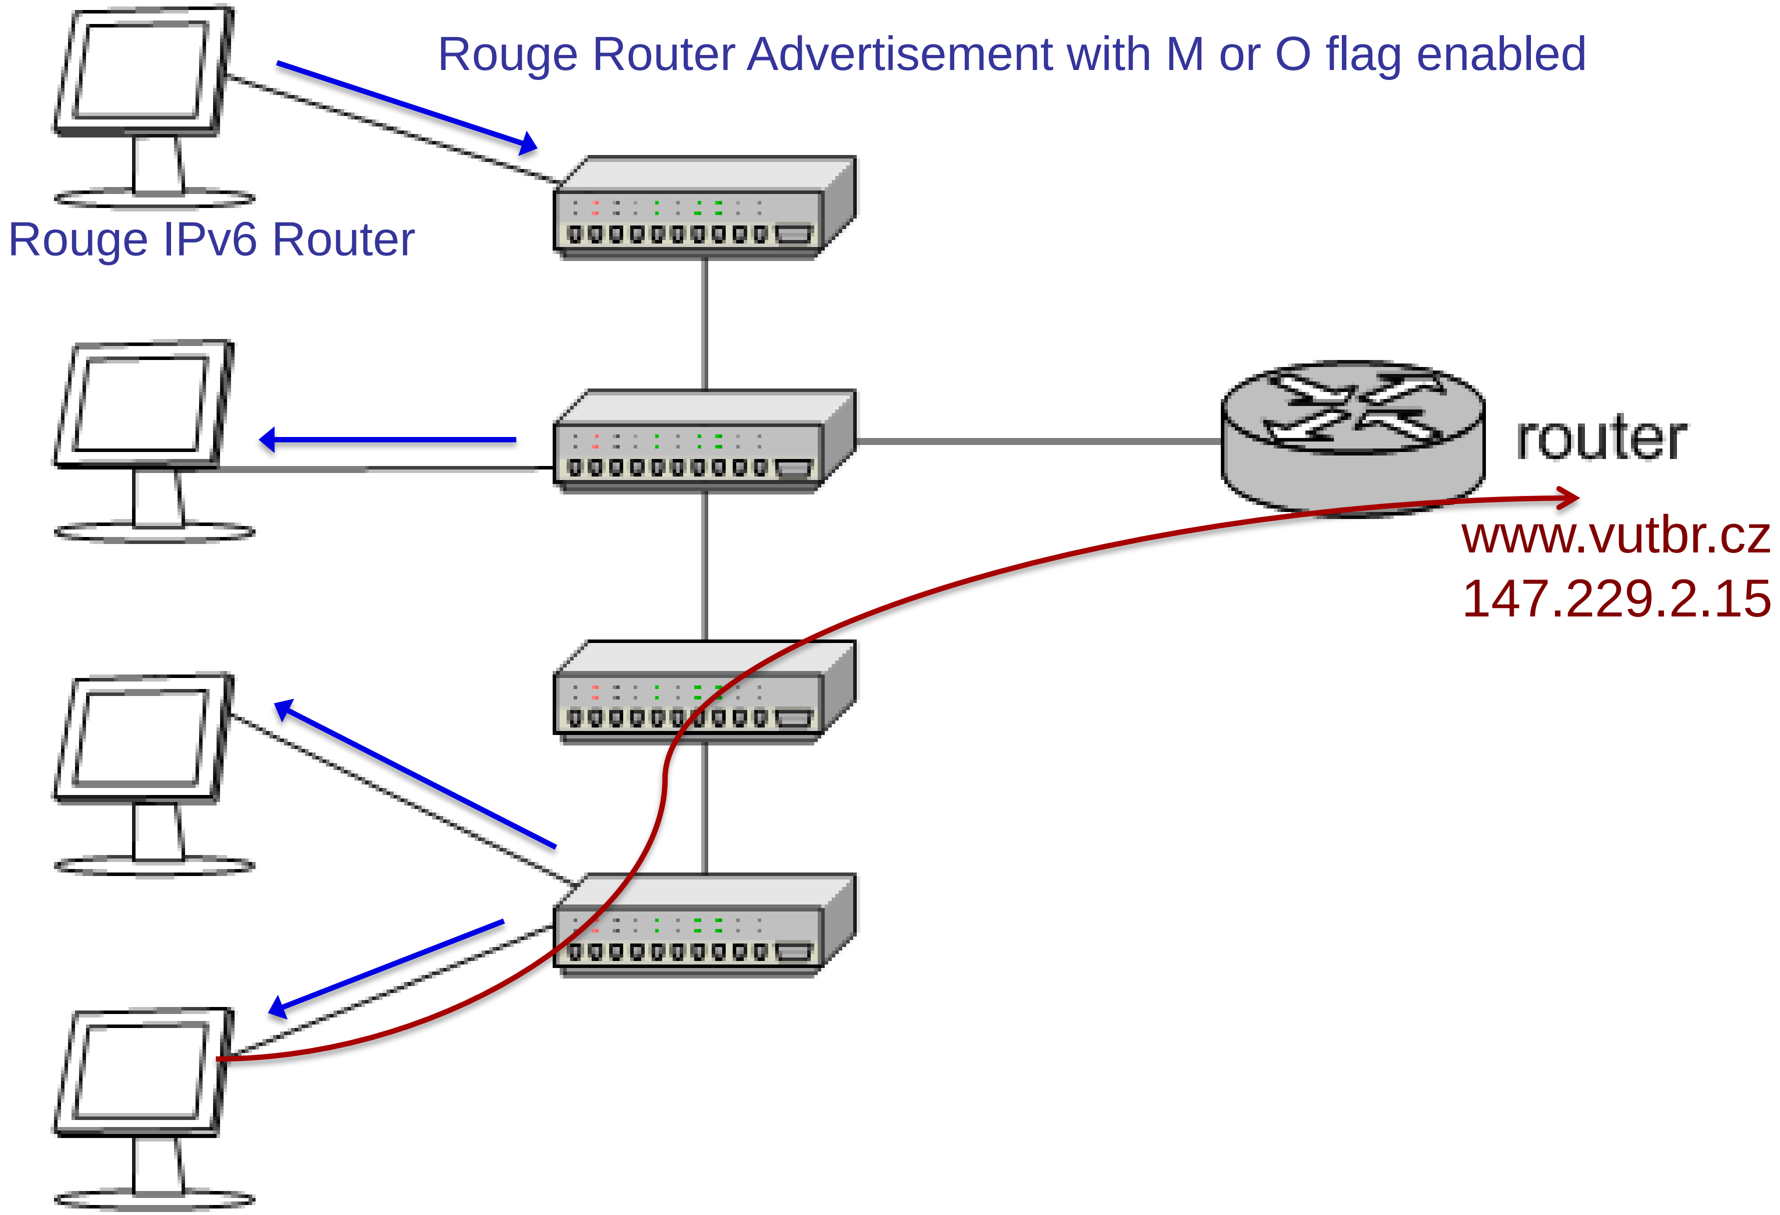

IPv4 traffic redirection using IPv6

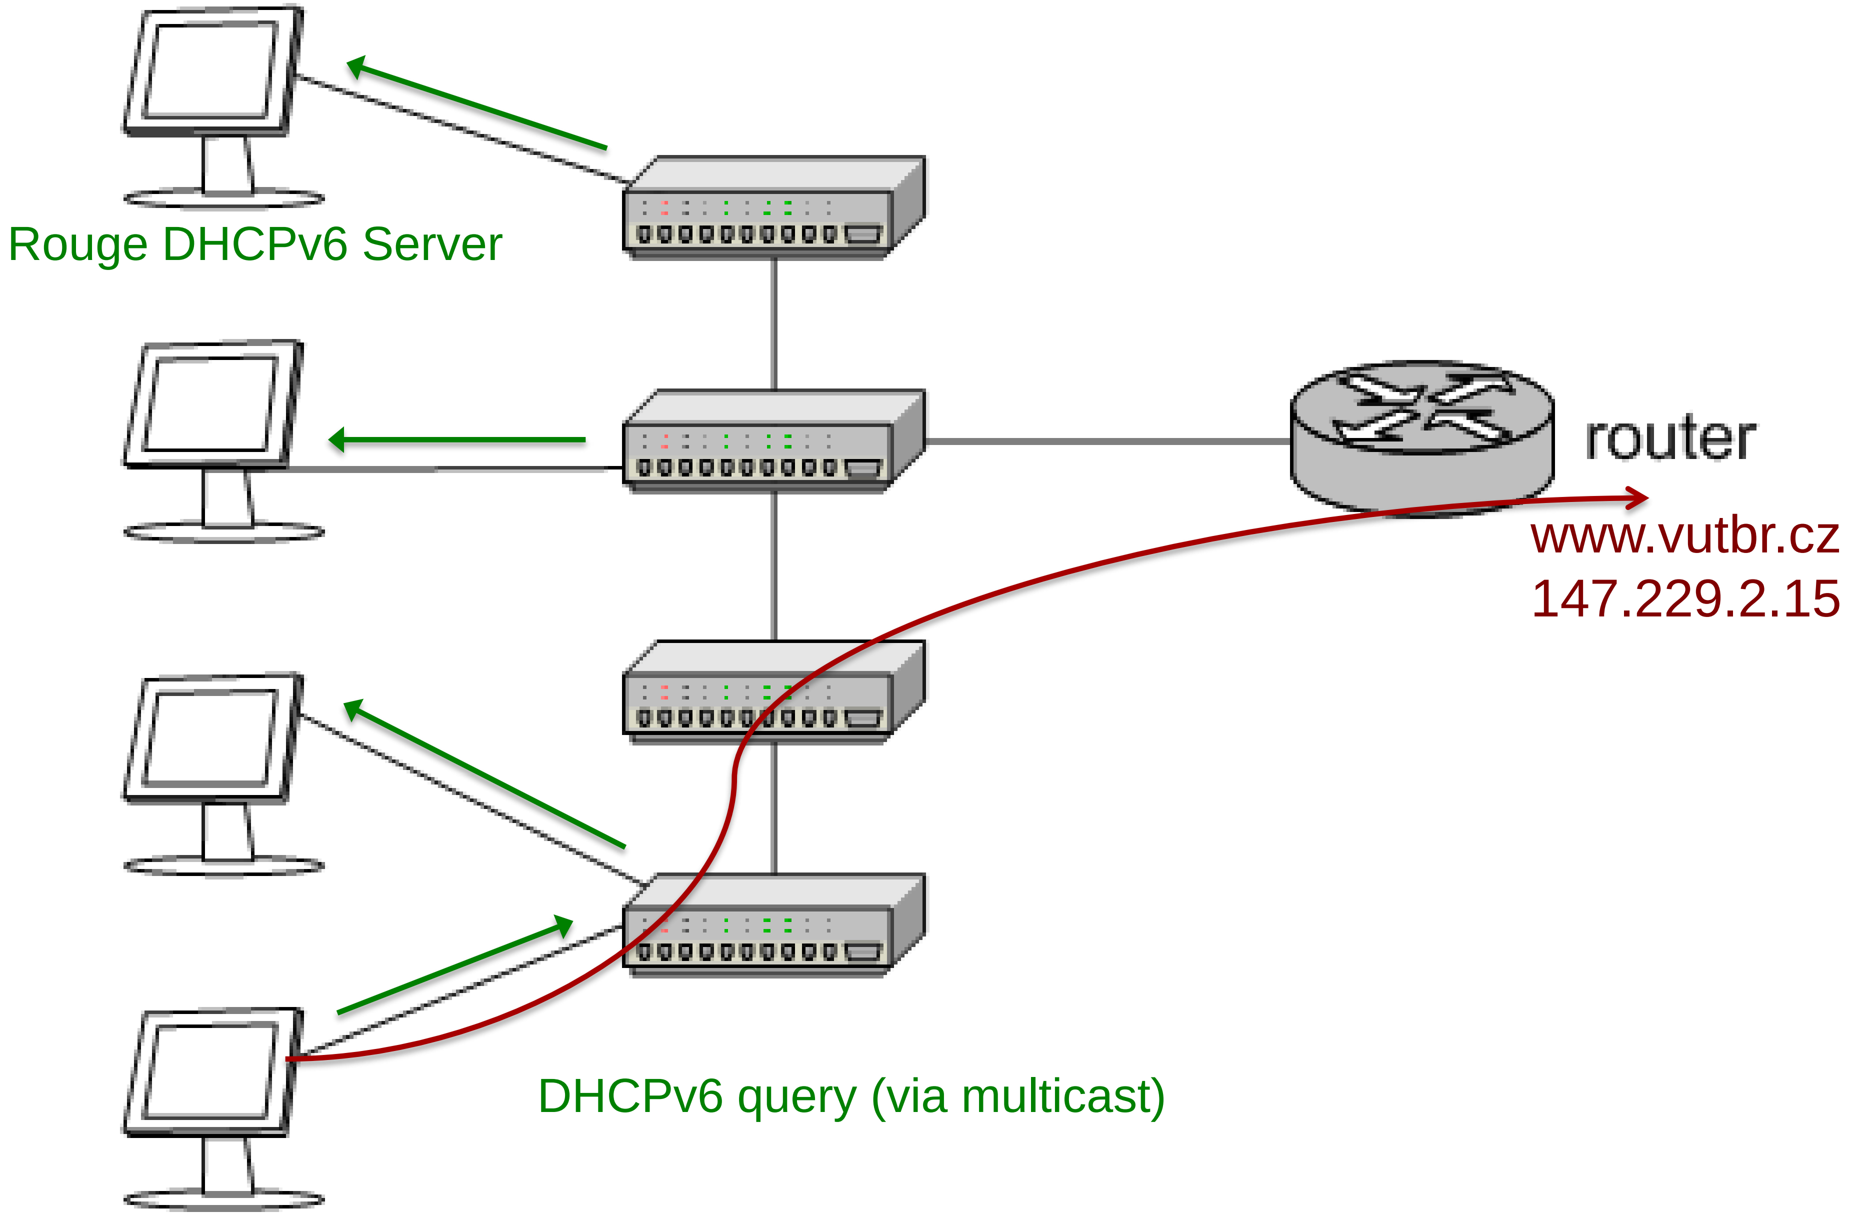

IPv4 traffic redirection using IPv6

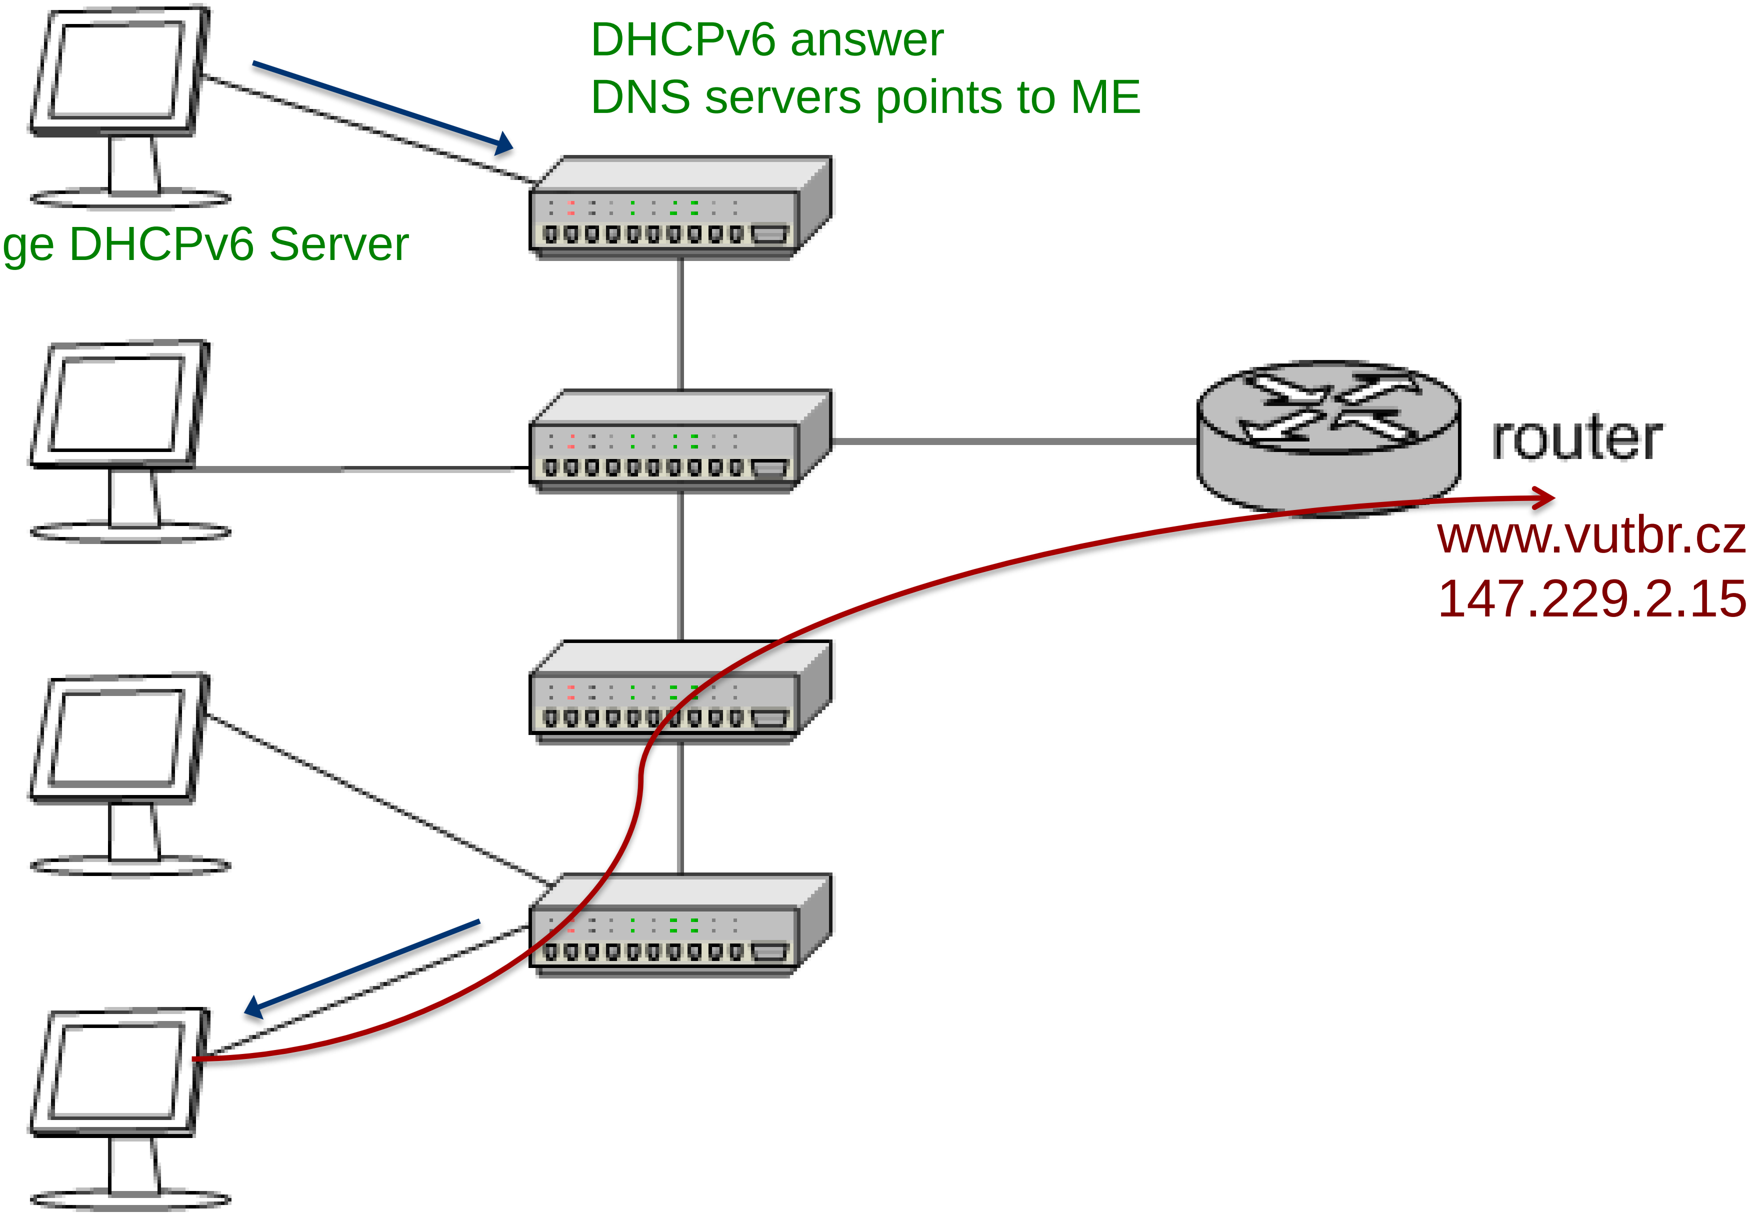

IPv4 traffic redirection using IPv6

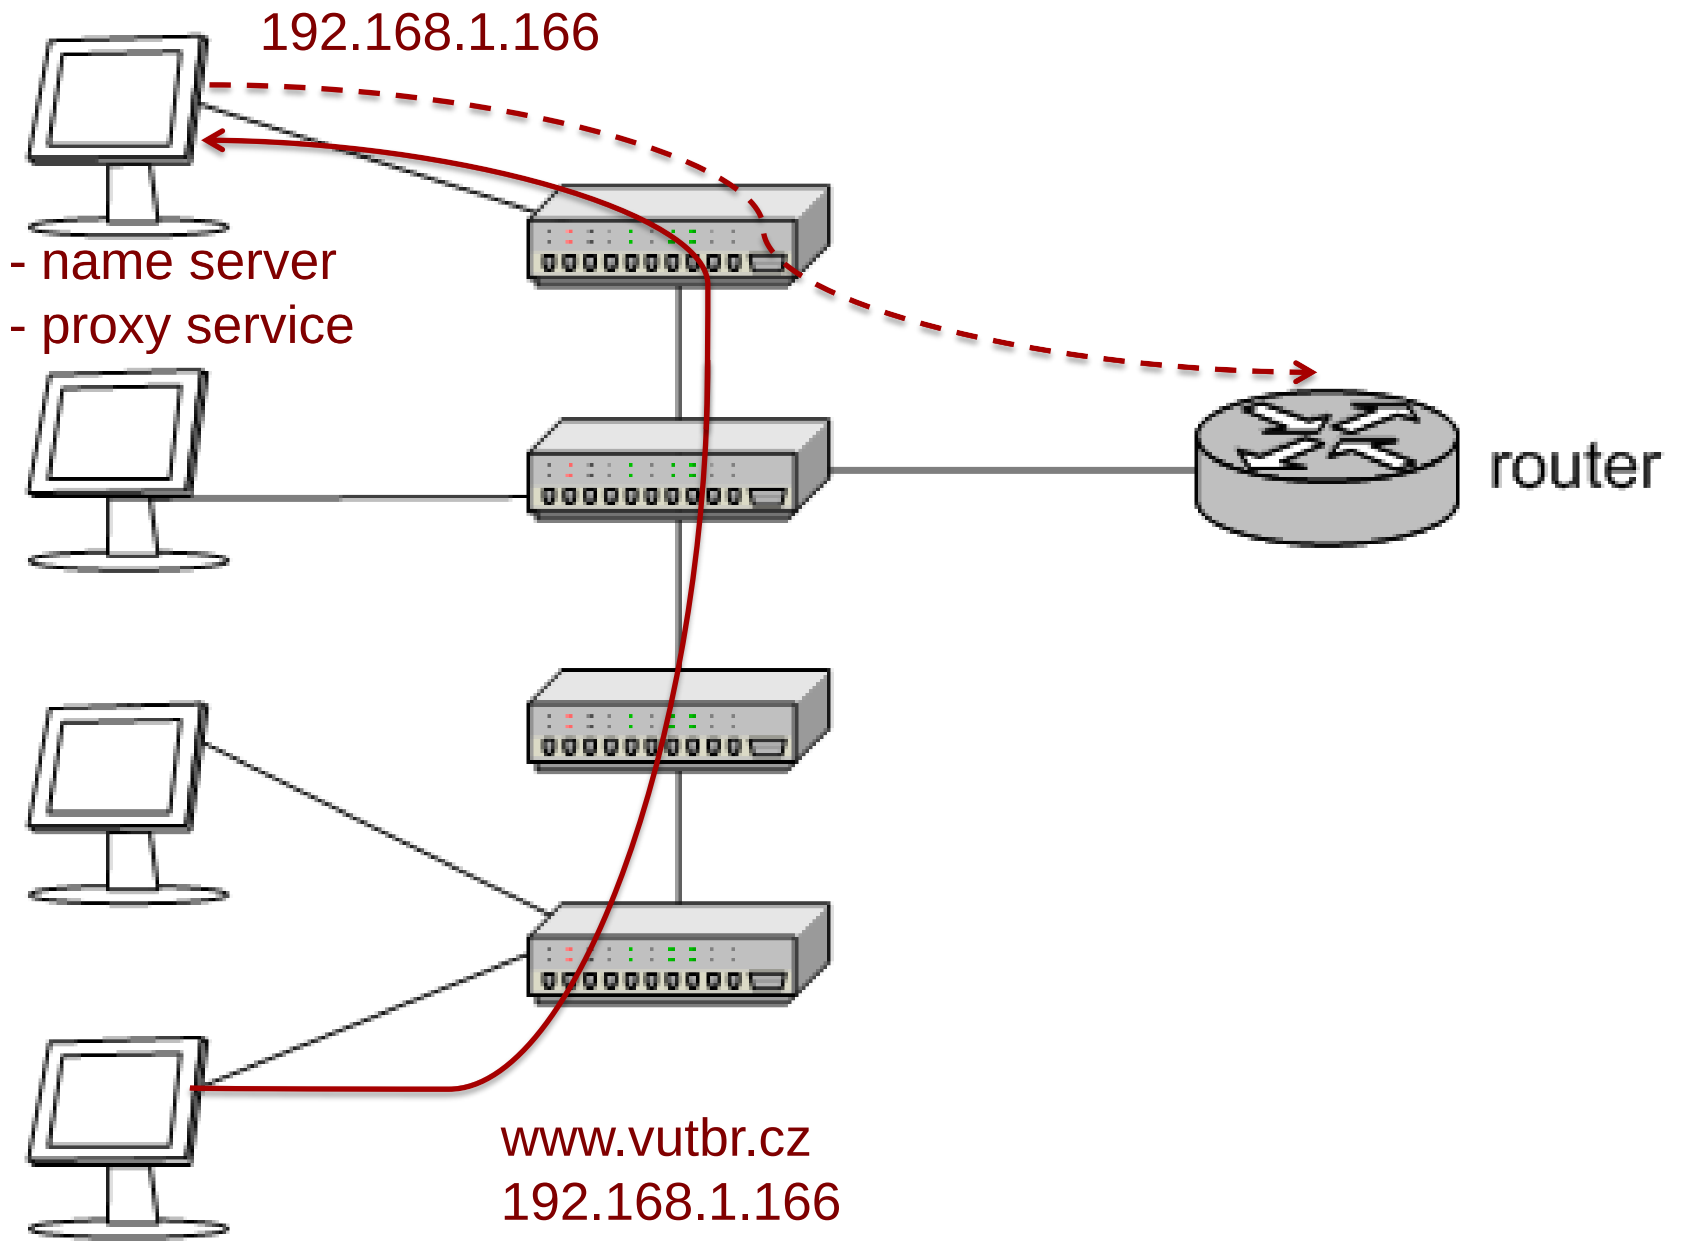

First hop security for IPv6

- SeND (RFC 3971, March 2005)
  - Based on cryptography CGA keys
  - Requires PKI infrastructure
    - How client obtains his own certificate?
  - Can **not** work with
    - Manually configured, EUI 64 and Privacy Extension addresses
- RA-Guard, PACL (RFC 6105, February 2011)
  - Dropping fake RA messages on access port (RA Snooping)
  - Can be easily avoided using fragmentation and extension headers (http://tools.ietf.org/html/draft-gont-v6ops-ra-guard-evasion-01)
- SAVI (draft-ietf-savi-*, divided into more drafts)
  - Complex solution including source addres validation
  - ND Inspection (Cisco devices) provides similar level of protection

First hop security for IPv6

Fragment reassembling or *undetermined traffic* option must be used (Cisco).

	RA-Guard	ACL/PACL	SAVI, ND inspection	~~SEND~~
Rogue DHCPv6 server		✓	✓	
Accidently rogue router advertise	✓	✓		~~✓~~
Intentional rouge router advertise				~~✓~~
ND cache poisoning			✓	~~✓~~
Source IPv6 address spoofing			✓	~~✓~~
DAD DOS attack			✓	
Neighbour cache overload			✓	
Source MAC address spoofing			✓	
Requires support on access switches	✓	✓	✓	
Requires support or configuration on client side				✓

Do we really need first hop security ?

- Fact: 15 years ago we did not have first hop security for IPv4 as well. We need more time while those features will be widely available for IPv6 (for good price).

but…

- Today we do not build networks that will be operated in 1998.
- Requirements on networks are extremely different than 15 years ago.
- **Today, we are building networks that should work for next 5 years!**

Costs of IPv6

Cost of first hop security for IPv6

- Real example: access switches for 150 users
  168 ports, 1Gb/s, non PoE

	price per port	RA guard	PACL	SAVI, ND inspection	DHCP snooping	ARP inspection, ARP protection	IP source guard, dynamic IP lockdown
1x HP 4208-96 vl Switch (J8775B) 3x HP 24-port Gig-T vl Module (J8768A)	**$58.90**				✔	✔	
1x HP 2910-48G al Switch (J9147A) 3x HP 2910-24G al Switch (J9145A)	**$96.23**				✔	✔	✔

Prices taken from http://www.amazon.com/ (list prices)

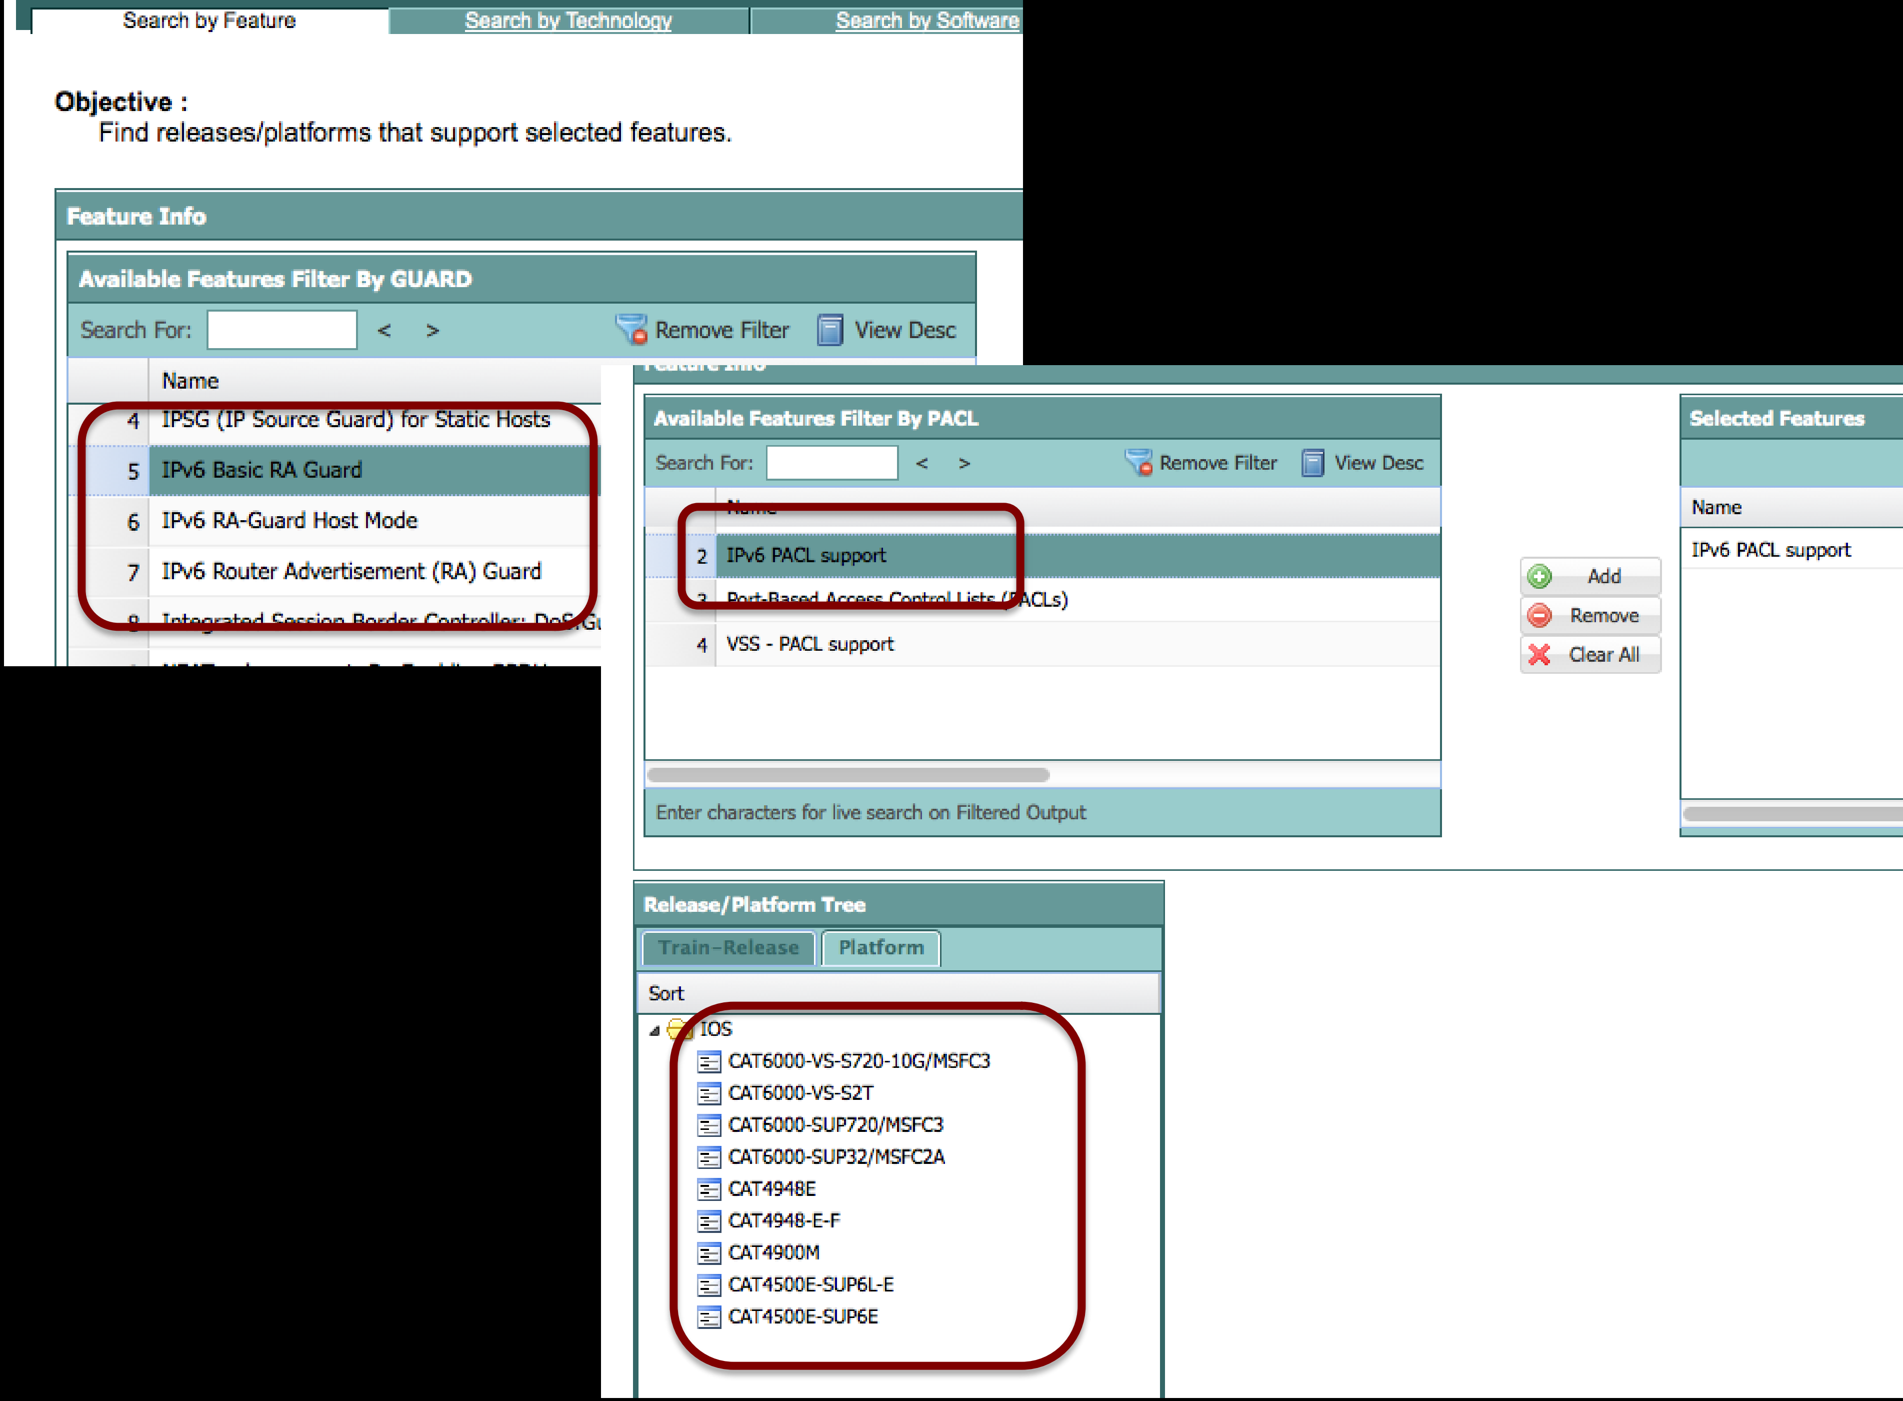

Objective :
Find releases/platforms that support selected features.

	Name
4	IPSG (IP Source Guard) for Static Hosts
5	IPv6 Basic RA Guard
6	IPv6 RA-Guard Host Mode
7	IPv6 Router Advertisement (RA) Guard
8	Integrated Session Border Controller: DoS/Gu

	Name
2	IPv6 PACL support
3	Port-Based Access Control Lists (PACLs)
4	VSS - PACL support

Name
IPv6 PACL support

- IOS
  - CAT6000-VS-S720-10G/MSFC3
  - CAT6000-VS-S2T
  - CAT6000-SUP720/MSFC3
  - CAT6000-SUP32/MSFC2A
  - CAT4948E
  - CAT4948-E-F
  - CAT4900M
  - CAT4500E-SUP6L-E
  - CAT4500E-SUP6E

Other costs related to IPv6

- Hardware upgrade/replacement
  - Netflow probes, switches
  - Many of them can be replaced as the part of “common upgrade” (but not all)
- Modification/upgrade of existing systems
  - Monitoring (new plugins, support for dualstack)
  - Administrative tools, network IS, …
  - RT system for security incidents
  - Documentation of the network
- Training
  - Administrators, operators (users)
  - Troubleshooting becomes more complicated

IPv6 and other priorities

- Second (resilient) data center
  - effects: *better services stability*
- Move servers to virtualization platform
  - effects: *optimalisation of energy consumption*
- Moving to MPLS
  - effects: *new L2, L3 VPN services for customers*
- Upgrade core of the network to 10Gb/s
  - effects: *network performance for users and ser*
- Backup connectivity to provider
  - effects: *better stability for external c*
- ~~Deploy IPv6 in the network~~
  - ~~Let's do~~ something for ~~our~~ distant future

✓ done...
10% saved and
shifted to next year

Lesson we learned with deploying IPv6

Lesson we learned – technical view

- Some parts of IPv6 are not mature enough for a production yet
  - Lack of support in applications
    - Skype, ICQ, MSN, mysql, …
  - Lack of support in network devices
    - Access switches, Firewalls, Load balancers, IDS/IPS, VPN
  - Lack of support in standards (or not implemented yet)
    - DHCP failover, VRRPv3, routes over DHCPv6, PXE, …
  - Some features are only in experimental versions of firmware
  - IPv6 related problems are on the edge of the interest
    - Flooding RA (Win 8), Rogue routers,

Lesson we learned – management view

- Strong support from management is essential
  - Human resources, budget, …
  - Patience: firmware updates, more reboots, …
  - IPv6 can have an impact on existing infrastructure
- It is necessary to have good relationship with your hardware and software suppliers
  - There probably will be many features that have never been tested before
- Deploying IPv6 does not bring any short term economical benefits. We can switch IPv6 of and nobody would notice.

but …

We can we afford to ignore IPv6 ?

Definitely not …

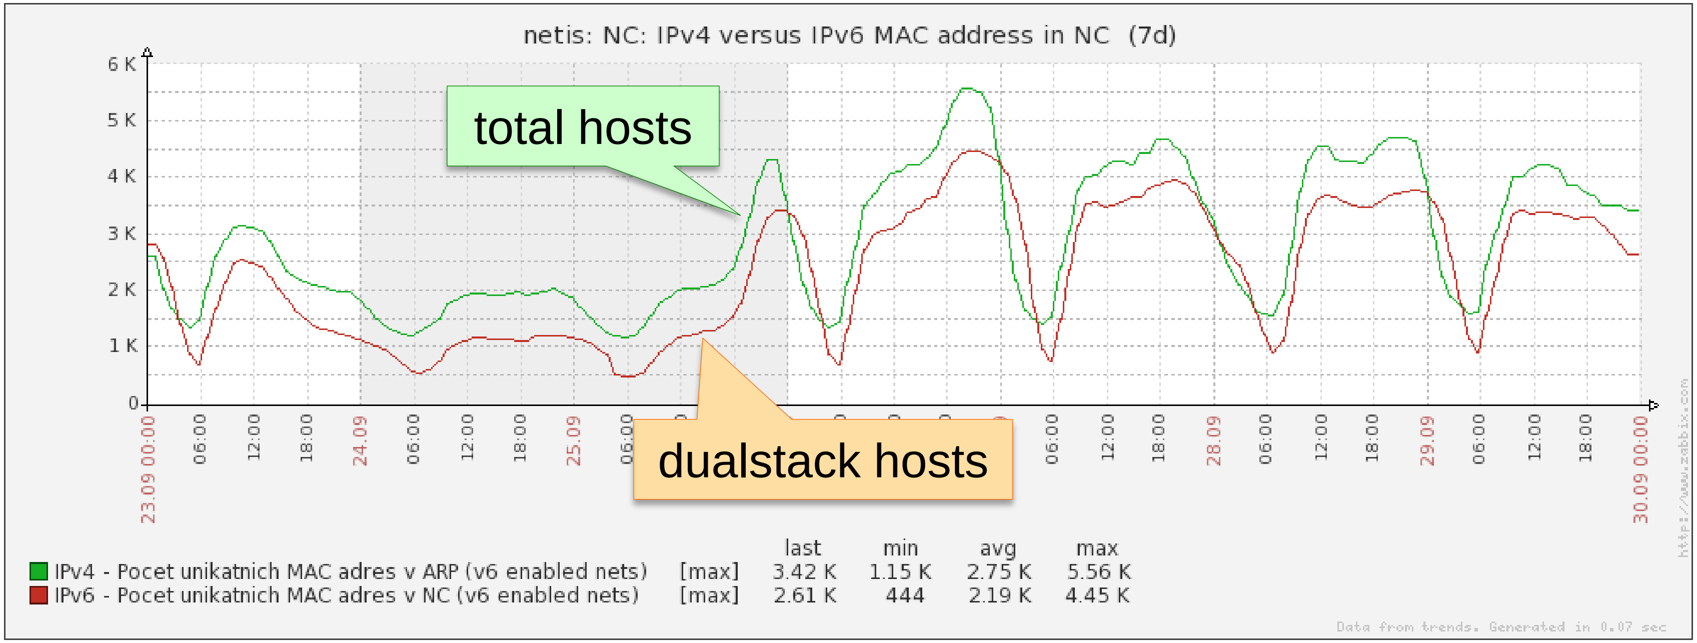

total hosts

dualstack hosts

		last	min	avg	max
🟩 IPv4 - Pocet unikatnich MAC adres v ARP (v6 enabled nets)	[max]	3.42 K	1.15 K	2.75 K	5.56 K
🟥 IPv6 - Pocet unikatnich MAC adres v NC (v6 enabled nets)	[max]	2.61 K	444	2.19 K	4.45 K

- There is almost 80% of hosts on the network with IPv6 enabled by default *(measured at students dormitory)*
- Ignoring existence of IPv6 can cause more troubles.

What can we do about it ? (if we can)

- Start using IPv6 immediately
  - We have been waiting for perfect IPv6 more than 15 years - it does not work
  - Until IPv6 is used we will not discover any problem
- Prefer native IPv6 connectivity (anywhere you can)
  - Native IPv6 is more secure than unattended tunneled traffic !
- Ask vendors and creators of standards to fix problems
  - More requests escalate troubles on the vendor side
  - Standardization of IPv6 is not enclosed process. Anyone can contribute or comment the standards
- Stop pretending that IPv6 does not have any troubles
  - IPv6 has got many problems
  - Problems can not be solved by covering them
  - Unreliable information led to broken trust amongst users. The naked truth is always better than the best dressed lie