


ОСНОВЫ IPv6

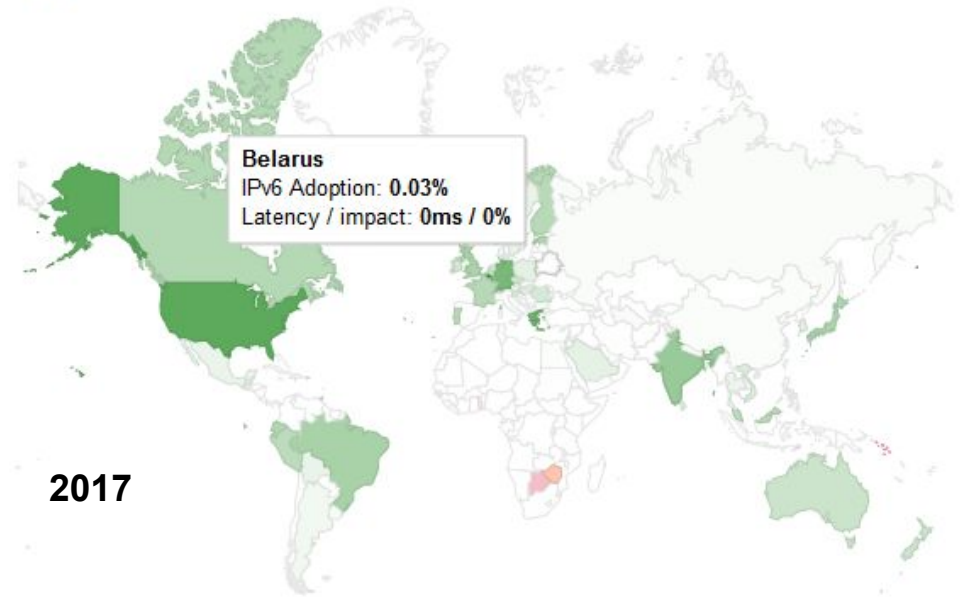
Иван Семерник
hoster.by
2017



RIR IPv4 Address Run-Down Model

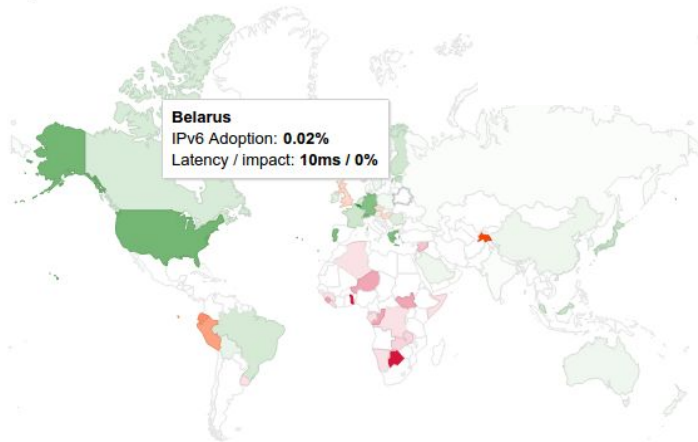


Per-Country IPv6 adoption




2017

Per-Country IPv6 adoption



2016

Native: 18.20% 6to4/Teredo: 0.05% Total IPv6: 18.25% | Jun 4, 2017



	IPv4	IPv6
Адрес	32 бита (4 байта) 212.98.163.254 4 294 967 296 адресов 1981 ~1 адрес на человека 8,5 адресов на 1км2 площади Земли	128 бит (16 байт) 2001:0db8:0006:0056:0000:0000:0000:0053 340 282 366 920 938 463 463 374 607 431 768 211 456 адресов 1998 4,7*10^28 адресов на человека 6,7*10^23 адресов на м2 площади Земли
MTU (минимальный)	576, фрагментация опционально	1280 байт без фрагментации
Фрагментация	Роутеры и хосты	Только хосты
DNS записи	A PTR in-addr.arpa	AAAA PTR ip6.arpa
Настройка адреса	Ручная или DHCP	Ручная, StateLess Address AutoConfiguration (SLAAC) и/или DHCPv6
Определение MAC	Broadcast ARP	Multicast Neighbor Solicitation
Broadcast	Да	Нет
Multicast \ Anycast	Да \ Да	Да \ Да
IPSec заголовки	Опционально	Обязательно
Маска подсети	Да	Нет

Формат IPv6 адреса

2a04:2e80:8:11:444e:5036:434e:a302

site prefix subnet ID

interface ID

site prefix + subnet ID = address prefix (/64)

Формат IPv6 адреса

2a04:2e80:8:11:444e:5036:434e:a302%17

Network Connection ID

Адресация

IPv6

```
$ sipcalc 2a02:bf0:6:56::53/64  
-[ipv6 : 2a02:bf0:6:56::53/64] - 0
```

[IPV6 INFO]

Expanded Address - 2a02:0bf0:0006:0056:0000:0000:0000:0053

Compressed address - 2a02:bf0:6:56::53

Subnet prefix (masked) - 2a02:bf0:6:56:0:0:0:0/64

Address ID (masked) - 0:0:0:0:0:0:53/64

Prefix address - ffff:ffff:ffff:ffff:0:0:0:0

Prefix length - 64

Address type - Aggregatable Global Unicast Addresses

Network range - 2a02:0bf0:0006:0056:0000:0000:0000:0000 -
2a02:0bf0:0006:0056:ffff:ffff:ffff:ffff

IPv4

```
$ ipcalc 212.98.163.242
```

Address: 212.98.163.242

Netmask: 255.255.255.0 = 24

Wildcard: 0.0.0.255

=>

Network: 212.98.163.0/24

HostMin: 212.98.163.1

HostMax: 212.98.163.254

Broadcast: 212.98.163.255

Hosts/Net: 254

```
11010100.01100010.10100011. 11110010  
11111111 .11111111 .11111111 . 00000000  
00000000.00000000.00000000. 11111111
```

```
11010100.01100010.10100011. 00000000  
11010100.01100010.10100011. 00000001  
11010100.01100010.10100011. 11111110  
11010100.01100010.10100011. 11111111
```

Class C

любое количество
нулей подряд, но
только один раз,
заменяются на ::

лидирующие нули
(слева) не указываются

Доступ к IPv6 адресу

```
wget http://[2a04:2e80:8:11::37]/
```

```
wget http://[2a04:2e80:8:11::37]:8080/
```

DNS

cloud.hoster.by. IN AAAA 2a0a:7d80:1:2::2

2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.2.0.0.0.1.0.0.0.0.8.d.7.a.0.a.2.ip6.arpa.

IN PTR cloud.hoster.by.

Выделение IPv6 префиксов

2000::/3 - $\frac{1}{8}$ of ::/0 (Global Unicast address space)


/23 - RIR (up to /12). 2a00:0000::/12 -> RIPE NCC

/32 - LIR, ISP / Large End Unit (up to /29). 2a0a:7d80::/29 -> hoster.by (наш LIR)

/48 - End Unit / Customer. 2a0a:7d80::/48 - > TUT.BY

/56 - End Unit / Home

/64 - End Unit L2 network



Типы IPv6 адресов

Unicast

Multicast

Anycast

Области IPv6 адресов

Assignable Global Unicast

2000::/3

>16 областей специального назначения:

loopback

::1/128

Unique Local Addresses (ULAs)

fc00::/7

Link-Local Addresses

fe80::/10

Multicast

ff00::/8

6to4

2002::/16

IPv4-mapped Address

::ffff:a.b.c.d/96

Teredo

2001:0000::/32

Documentation

2001:db8::/32

EUI-64

...

EUI-64

12:34:56:78:90:00

2001:db8:a:a:1034:56ff:fe78:9000

$12_{(16)} = 00010010_{(2)}$

$00010000_{(2)} = 10_{(16)}$

Используется в SLAAC

IPv6 Multicast scopes

FF0X::

- 0** - reserved
- 1** - Interface-Local scope
- 2** - Link-Local scope
- 3** - Realm-Local. Should be published in an RFC.
- 4** - Admin-Local scope
- 5** - Site-Local scope
- 6** and **7** are unassigned
- 8** - Organization-Local scope
- 9** to **D** are unassigned
- E** - Global scope
- F** - reserved

```
$ ping6 ff02::1%eth0
```

Address	Meaning	Scope
FF02::1	All nodes	Link-local
FF02::2	All routers	Link-local
FF02::9	All RIP routers	Link-local
FF02::1:FFXX:XXXX	Solicited-node	Link-local
FF02::4	DVMRP routers	Link-local
FF02::5	All OSPF routers	Link-local
FF02::6	OSPF designated routers	Link-local
FF02::A	All EIGRP routers	Link-local
FF02::D	All PIM routers	Link-local
FF05::101	All NTP routers	Site-local
FF05::1:3	All DHCP routers	Site-local

Solicited-node multicast address

fe80::96de:80ff:feba:776a	Мой link-local адрес
fe80:0000:0000:0000:96de:80ff:fe ba:776a	Полный
ba:776a	Последние 24 бита
ff02::1:ff00:0/104	Solicited-node multicast prefix
ff02:0000:0000:0000:0000:0001:ff00:0000/104	Полный
ff02:0000:0000:0000:0000:0001:ff	Первые 104 бита
ff02:0000:0000:0000:0000:0001:ffba:776a	Результат
ff02::1:ffba:776a	Сжатый результат

Solicited-node multicast address example

```
$ ip -6 maddr show dev eth0
```

```
2: eth0
```

```
inet6 ff02::1:ff1b:79af      <- 2a04:2e80:7:3:a174:a49d:ad1b:79af
```

```
inet6 ff02::1:ffba:776a     <- fe80::96de:80ff:feba:776a
```

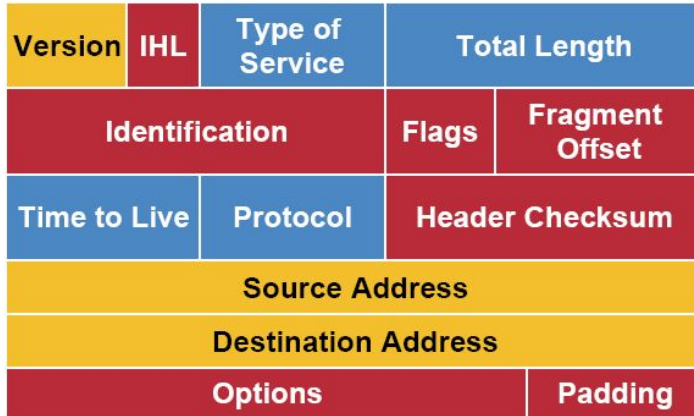
```
inet6 ff02::1
```

```
inet6 ff01::1
```

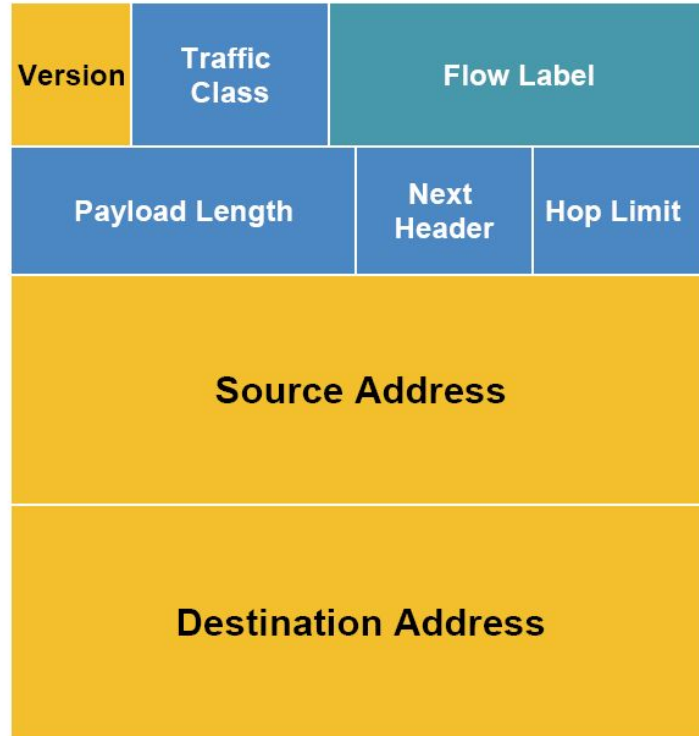
IPv4 BroadCast VS ipv6 MultiCast

Headers

IPv4 Header



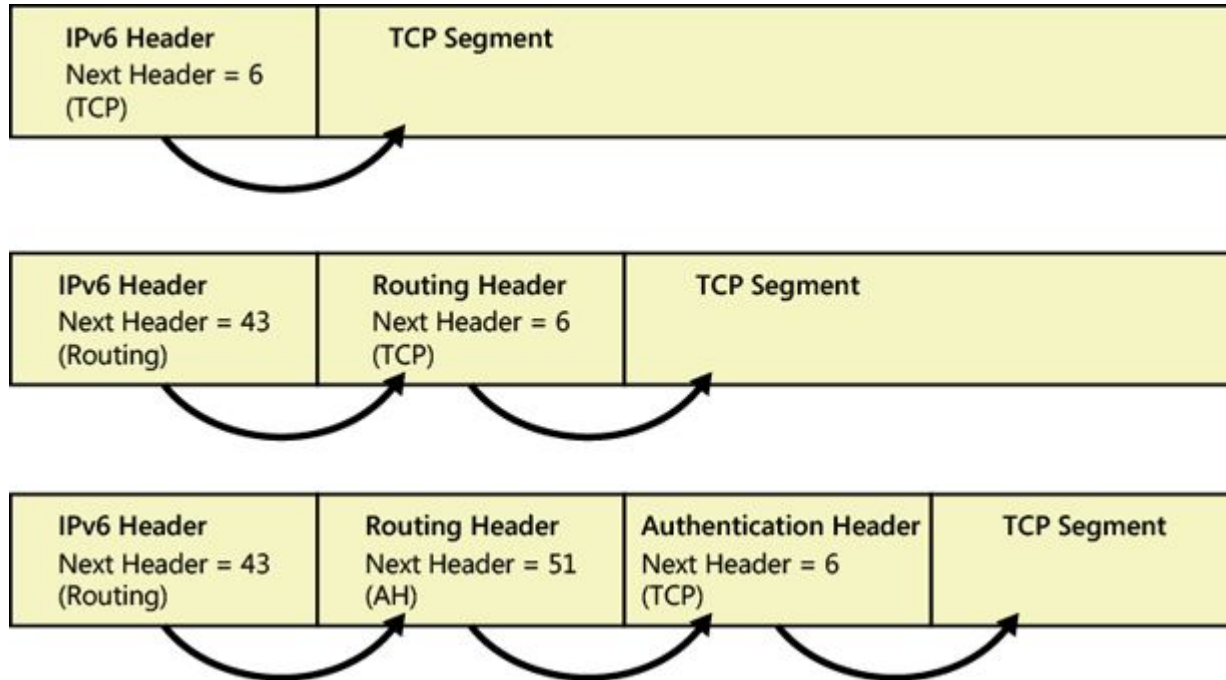
IPv6 Header



Legend

- Field's Name Kept from IPv4 to IPv6
- Fields Not Kept in IPv6
- Name and Position Changed in IPv6
- New Field in IPv6

IPv6 header extension

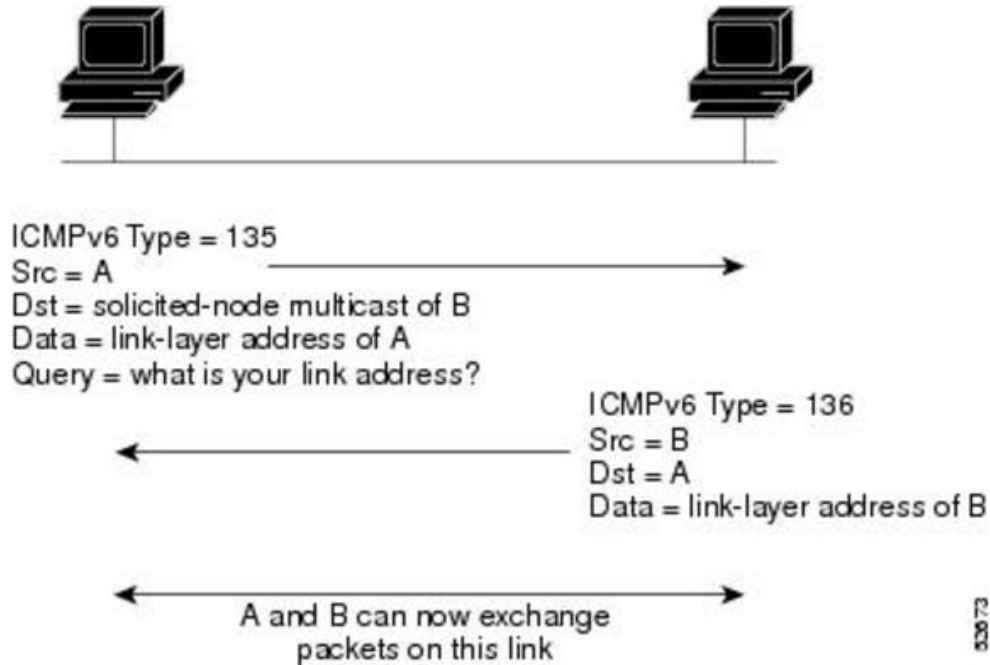


Коды ошибок			Информационные сообщения		
Тип	Описание	RFC	Тип	Описание	RFC
1	Destination Unreachable	RFC 4443	128	Echo Request	RFC 4443
2	Packet Too Big	RFC 4443	129	Echo Reply	RFC 4443
3	Time Exceeded	RFC 4443	130	Multicast Listener Query	RFC 2710 и RFC 3810
4	Parameter Problem	RFC 4443	131	Version 1 Multicast Listener Report	RFC 2710
100	Private experimentation		132	Multicast Listener Done	RFC 2710
101	Private experimentation		133	Router Solicitation	RFC 4861
127	Reserved for expansion of ICMPv6 error messages		134	Router Advertisement	RFC 4861
			135	Neighbor Solicitation	RFC 4861
			136	Neighbor Advertisement	RFC 4861
			137	Redirect	RFC 4861
			138	Router Renumbering	
			139	ICMP Node Information Query	
			140	ICMP Node Information Response	
			141	Inverse Neighbor Discovery Solicitation Message	RFC 3122
			142	Inverse Neighbor Discovery Advertisement Message	RFC 3122
			143	Version 2 Multicast Listener Report	RFC 3810
			144	Home Agent Address Discovery Request Message	RFC 3775
			145	Home Agent Address Discovery Reply Message	RFC 3775
			146	Mobile Prefix Solicitation	RFC 3775
			147	Mobile Prefix Advertisement	RFC 3775
			148	Certification Path Solicitation Message	RFC 3971
			149	Certification Path Advertisement Message	RFC 3971
			150	ICMP messages utilized by experimental mobility protocols such as Seamoby	RFC 4065
			151	Multicast Router Advertisement	RFC 4286
			152	Multicast Router Solicitation	RFC 4286
			153	Multicast Router Termination	RFC 4286
			200	Private experimentation	
			201	Private experimentation	
			255	Reserved for expansion of ICMPv6 informational messages	

ICMPv6

IPv6 NA

Figure 1 IPv6 Neighbor Discovery: Neighbor Solicitation Message



IPv6 RA

Figure 2 IPv6 Neighbor Discovery: RA Message



Router advertisement packet definitions:

ICMPv6 Type = 134

Src = router link-local address

Dst = all-nodes multicast address

Data = options, prefix, lifetime, autoconfig flag

▶ Frame 5: 174 bytes on wire (1392 bits), 174 bytes captured (1392 bits) on interface 0

▶ Ethernet II, Src: Hewlett-92:3a:8b (00:17:08:92:3a:8b), Dst: Giga-Byt_ba:77:6a (94:de:80:ba:77:6a)

▶ Internet Protocol Version 6, Src: fe80::217:8ff:fe92:3a8b (fe80::217:8ff:fe92:3a8b), Dst: fe80::96de:80ff:feba:776a (fe80::96de:80ff:feba:776a)

▼ Internet Control Message Protocol v6

Type: Router Advertisement (134)

Code: 0

Checksum: 0x8438 [correct]

Cur hop limit: 64

▼ Flags: 0x00

0... .. = Managed address configuration: Not set

.0... .. = Other configuration: Not set

..0... .. = Home Agent: Not set

...0... = Prf (Default Router Preference): Medium (0)

....0.. = Proxy: Not set

....0.. = Reserved: 0

Router lifetime (s): 30

Reachable time (ms): 0

Retrans timer (ms): 0

▼ ICMPv6 Option (Prefix information : 2a04:2e80:8:11::/64)

Type: Prefix information (3)

Length: 4 (32 bytes)

Prefix Length: 64

▼ Flag: 0xc0

1... .. = On-link flag(L): Set

.1.. .. = Autonomous address-configuration flag(A): Set

..0... = Router address flag(R): Not set

...0 0000 = Reserved: 0

Valid Lifetime: 3600

Preferred Lifetime: 1800

Reserved

Prefix: 2a04:2e80:8:11:: (2a04:2e80:8:11::)

▼ ICMPv6 Option (Route Information : Medium ::/0)

Type: Route Information (24)

Length: 3 (24 bytes)

Prefix Length: 0

▼ Flag: 0x00

...0 0... = Route Preference: Medium (0)

000..0000 = Reserved: 0

Route Lifetime: 3600

Prefix: :: (::)

▶ ICMPv6 Option (Recursive DNS Server 2a02:bf0:1:2281::1 2a02:bf0:1:2282::1)

▶ ICMPv6 Option (Source link-layer address : 00:17:08:92:3a:8b)

Linux: **radvd**, Windows: **wradvs**
Routers: **native**

```
root@gw:~# cat /etc/radvd.conf
interface vlan11
{
    AdvSendAdvert on;
    MinRtrAdvInterval 3;
    MaxRtrAdvInterval 10;
    prefix 2a04:2e80:8:11::/64 {
        AdvOnLink on;
        AdvAutonomous on;
        AdvRouterAddr off;
        AdvValidLifetime 3600;
        AdvPreferredLifetime 1800;
    };
    route ::/0 {
        AdvRouteLifetime 3600;
    };
    RDNSS 2a02:bf0:1:2281::1 2a02:bf0:1:2282::1 {
    };
    clients {
        fe80::96de:80ff:feba:776a;
        fe80::485f:43dc:9a92:d1a2;
    };
};
```

apt-get install rdnsd

Автоконфигурация IPv6

RA

ICMPv6 RA Flags

SLAAC (Stateless)

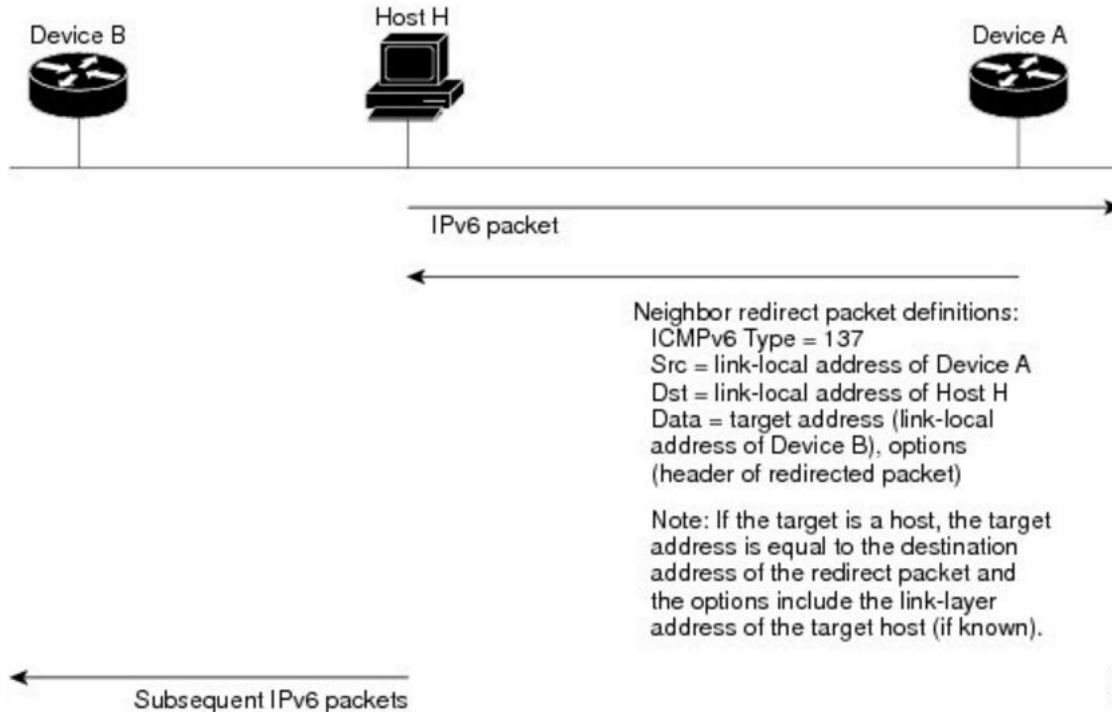
or/and

DHCPv6 (Stateful)

M bit	O bit	A bit	Prefix	Host Address	Other Config/Options
0	0	0	Static		Manual
0	0	1	RA Message	Autonomous	Manual
0	1	0	Static		DHCPv6
0	1	1	RA Message	Autonomous	DHCPv6
1	0	0	DHCPv6 (Stateful)		DHCPv6
1	0	1	DHCPv6 (Stateful) and/or autonomous		DHCPv6
1	1	0	DHCPv6 (Stateful)		DHCPv6
1	1	1	DHCPv6 (Stateful) and/or autonomous		DHCPv6

IPv6 redirects

Figure 3 IPv6 Neighbor Discovery: Neighbor Redirect Message



IPv6 DAD



ICMP type = 135

Src = ::

Dst = FF02::1:FF00:1

Data = 2000::1

NS



ICMP type = 136

Src = 2000::1

Dst = FF02::1

Target address = 2000::1

NA



Ручная настройка IPv6 (Linux)

/etc/sysconfig/network-scripts/ifcfg-interface

IPV6INIT="yes"

IPV6ADDR="2a02:f680:1:1100::afe/64"

IPV6ADDR_SECONDARIES="2a02:f680:1:1100::aff/64"

IPV6_DEFAULTGW="2a02:f680:1:1100::1"

CentOS/RHEL

/etc/sysconfig/network

NETWORKING_IPV6=yes

/etc/network/interfaces

iface eth0 inet6 static

address 2a04:2e80:1::157

netmask 64

gateway 2a04:2e80:1::

Debian/Ubuntu

Тюнинг IPv6 (Linux, sysctl)

```
net.ipv6.conf.eth0.disable_ipv6 = 0
net.ipv6.conf.eth0.accept_dad = 1
net.ipv6.conf.eth0.accept_ra = 1
net.ipv6.conf.eth0.accept_redirects = 1
net.ipv6.conf.eth0.autoconf = 1
net.ipv6.conf.eth0.dad_transmits = 1
net.ipv6.conf.eth0.forwarding = 0
net.ipv6.conf.eth0.hop_limit = 64
net.ipv6.conf.eth0.max_addresses = 16
net.ipv6.conf.eth0.mc_forwarding = 0
net.ipv6.conf.eth0.mtu = 1500
net.ipv6.conf.eth0.proxy_ndp = 0
net.ipv6.conf.eth0.temp_prefered_lft = 86400
net.ipv6.conf.eth0.temp_valid_lft = 604800
net.ipv6.conf.eth0.use_tempaddr = 2
```

IPv6 services

nginx:

```
server {  
    listen 443 ssl;  
    listen [::]:443 ssl;
```

Apache:

```
Listen [2001:db8::2]:80  
<VirtualHost 1.2.3.4:80 [2001:db8::2]:80>  
    ServerName example.com  
  
    ...  
</VirtualHost>
```

haproxy:

```
frontend app1_frontend  
    bind 1.2.3.4:80  
    bind 2001:db8::2:80  
    default_backend app1_backend
```

sshd:

```
ListenAddress 0.0.0.0  
ListenAddress ::
```

IPv6 security (server side)

Manual IPv6 configuration

Random Interface ID: 2001:2b8:0:0:**cf0:224c:0c3d:1bee**

`ip6tables -I INPUT 1 -p ipv6-icmp -j ACCEPT (обязательно!)`

`net.ipv6.conf.eth0.accept_ra = 0`

`net.ipv6.conf.eth0.accept_redirects = 0`

`net.ipv6.conf.eth0.autoconf = 0`

PHP \$_SERVER

<code>\$_SERVER["SERVER_SOFTWARE"]</code>	nginx/1.9.3
<code>\$_SERVER["REMOTE_ADDR"]</code>	2a04:2e80:1::66
<code>\$_SERVER["REMOTE_PORT"]</code>	50088
<code>\$_SERVER["SERVER_ADDR"]</code>	2a04:2e80:1::55
<code>\$_SERVER["SERVER_PORT"]</code>	443

Ваш IP: 0.0.0.102

IPv6 security (office)

Stateful firewall

DHCPv6 or IPv6 neighbour table logging

Как проверить

<http://internet.yandex.ru>

[**http://test-ipv6.com**](http://test-ipv6.com)

<http://ipv6test.google.com>

Портирование приложений из IPv4 в DualStack

AF_INET -> AF_INET6

sockaddr_in -> sockaddr_in6, dest_sin -> dest6_sin, sockfd -> sockfd6,

Константы похожи, размеры структур больше (32->128 бит)

gethostbyname() -> getaddrinfo()

<https://tools.ietf.org/html/rfc3493>

<http://www.lugod.org/presentations/ipv6programming/>

IPv6 в hoster.by и tut.by

ЦОДы:

IPv6 DualStack (manual configuration)

DNS TLD (.by .бел)

I.root-servers.net

TUT.BY - статика

Cloud

Secure Hosting

Офис (~500 devices):

SLAAC

30% офисного трафика - IPv6

Минск: МТС и Деловая Сеть

dnsX.tld.tutby.com

```
$ for i in {1..7}; do \  
dig +short dns$i.tld.tutby.com aaaa ; \  
done
```

2a04:2e80:7::2

2a04:2e80:7:1::2

2a05:4800:1:100::

2001:1af8:4010:a049:1::

2a00:c98:2200:a012:1::

2604:9a00:2010:a013:2::

2400:6180:0:d0::3d:b001

RIPE NCC Trainings



RIPE NCC

RIPE NETWORK COORDINATION CENTRE

Basic + Advanced IPv6 курсы

academy.ripe.net



RIPE NCC
RIPE NETWORK COORDINATION CENTRE

hoster.by



RIPE NCC Regional Meeting
Eurasia Network Operators Group

ENOG 14 / RIPE NCC Regional Meeting will take place
in Minsk, Belarus on 9-10 October 2017

enog.org

Спасибо за внимание!

Хакатон: IPv6 для Isee.org
2a0a:7d80:1:1::1bee:0/112 (65536 адресов)