RIPE NCC Regional Meeting Moscow 2010

IPv6 implementation AS8359

Andrey Alekseyev
Director IP service network
COMSTAR-UTS

Overview

- **Largest broadband customer base in Moscow**
- **Operates IPv4/IPv6 IP/MPLS backbone across Russia and Europe**
- **First Russian operator to launch commercial IPTV services in 2005**
- **LIR – ru.mtu (category LARGE)**
- **13 IPv4 allocations, largest is /14**
- **IPv6 2a02:28::/32 from RIPE NCC**
- **IPv6 native across entire IP/MPLS backbone**
- Currently Cisco-only network
- P routers – Cisco CRS-1 Series
- PE routers – Cisco 7600 Series, Cisco ASR 9000 Series
- Broadband customers in Moscow use routable IPv4 addressing
- Steadily running out of IPv4

Reasons for Operator

- Common issues with IPv4 exhaustion (growing customer base, fragmentation, inaccurate usage etc.)
- RIR tightening the policy
- What are the costs of delaying with implementation?
- Not implementing/testing it now means there will be problems with the transition in the future (administrative, technical, financial)
- Peering landscape is interesting – you can be a ‘transit-free’ one
- Big ones are doing that (Google, Facebook, Apple) – prepare for more before it’s too late
- Overcome psychological barriers/stereotypes – “Do or do not... there is no try” – but it still should be fun!
- Internet-of-things ahead (6loWPAN evolving etc.)
- IPTV deployments would benefit from it
- Certain regions of the world soon to become IPv6-only (APAC)?

What are the reasons for users

Hard to gather an opinion, but here are some excerpts from the forums:

- Torrents still the main driver after breakthrough in 2008
  - Could lead to faster downloads
  - Bypass anti-torrent measures
  - New IPv6 trackers facilitate IPv6 torrents traffic
- Tunnels provide static routable IPv6 addressing behind NAT
  - Own resources exposed globally
  - Uninterrupted connectivity in IPv4 session environment
- Selected applications work nicely (IRC, IM etc.)
- Anonymity (tentatively with tunnels, more with OnionCat-like apps)

Moscow network – key elements

- Access network in Moscow built and operated by MGTS (local PSTN)
- 250 access nodes (MGTS central offices), redundant 10GE to every node
- 5000 DSLAMs installed
- ADSL access 6-20 Mbps
- 80% of subscribers at 6 Mbps, 20% at 10+ Mbps
- PPPoE and RFC1483 (residential/business)

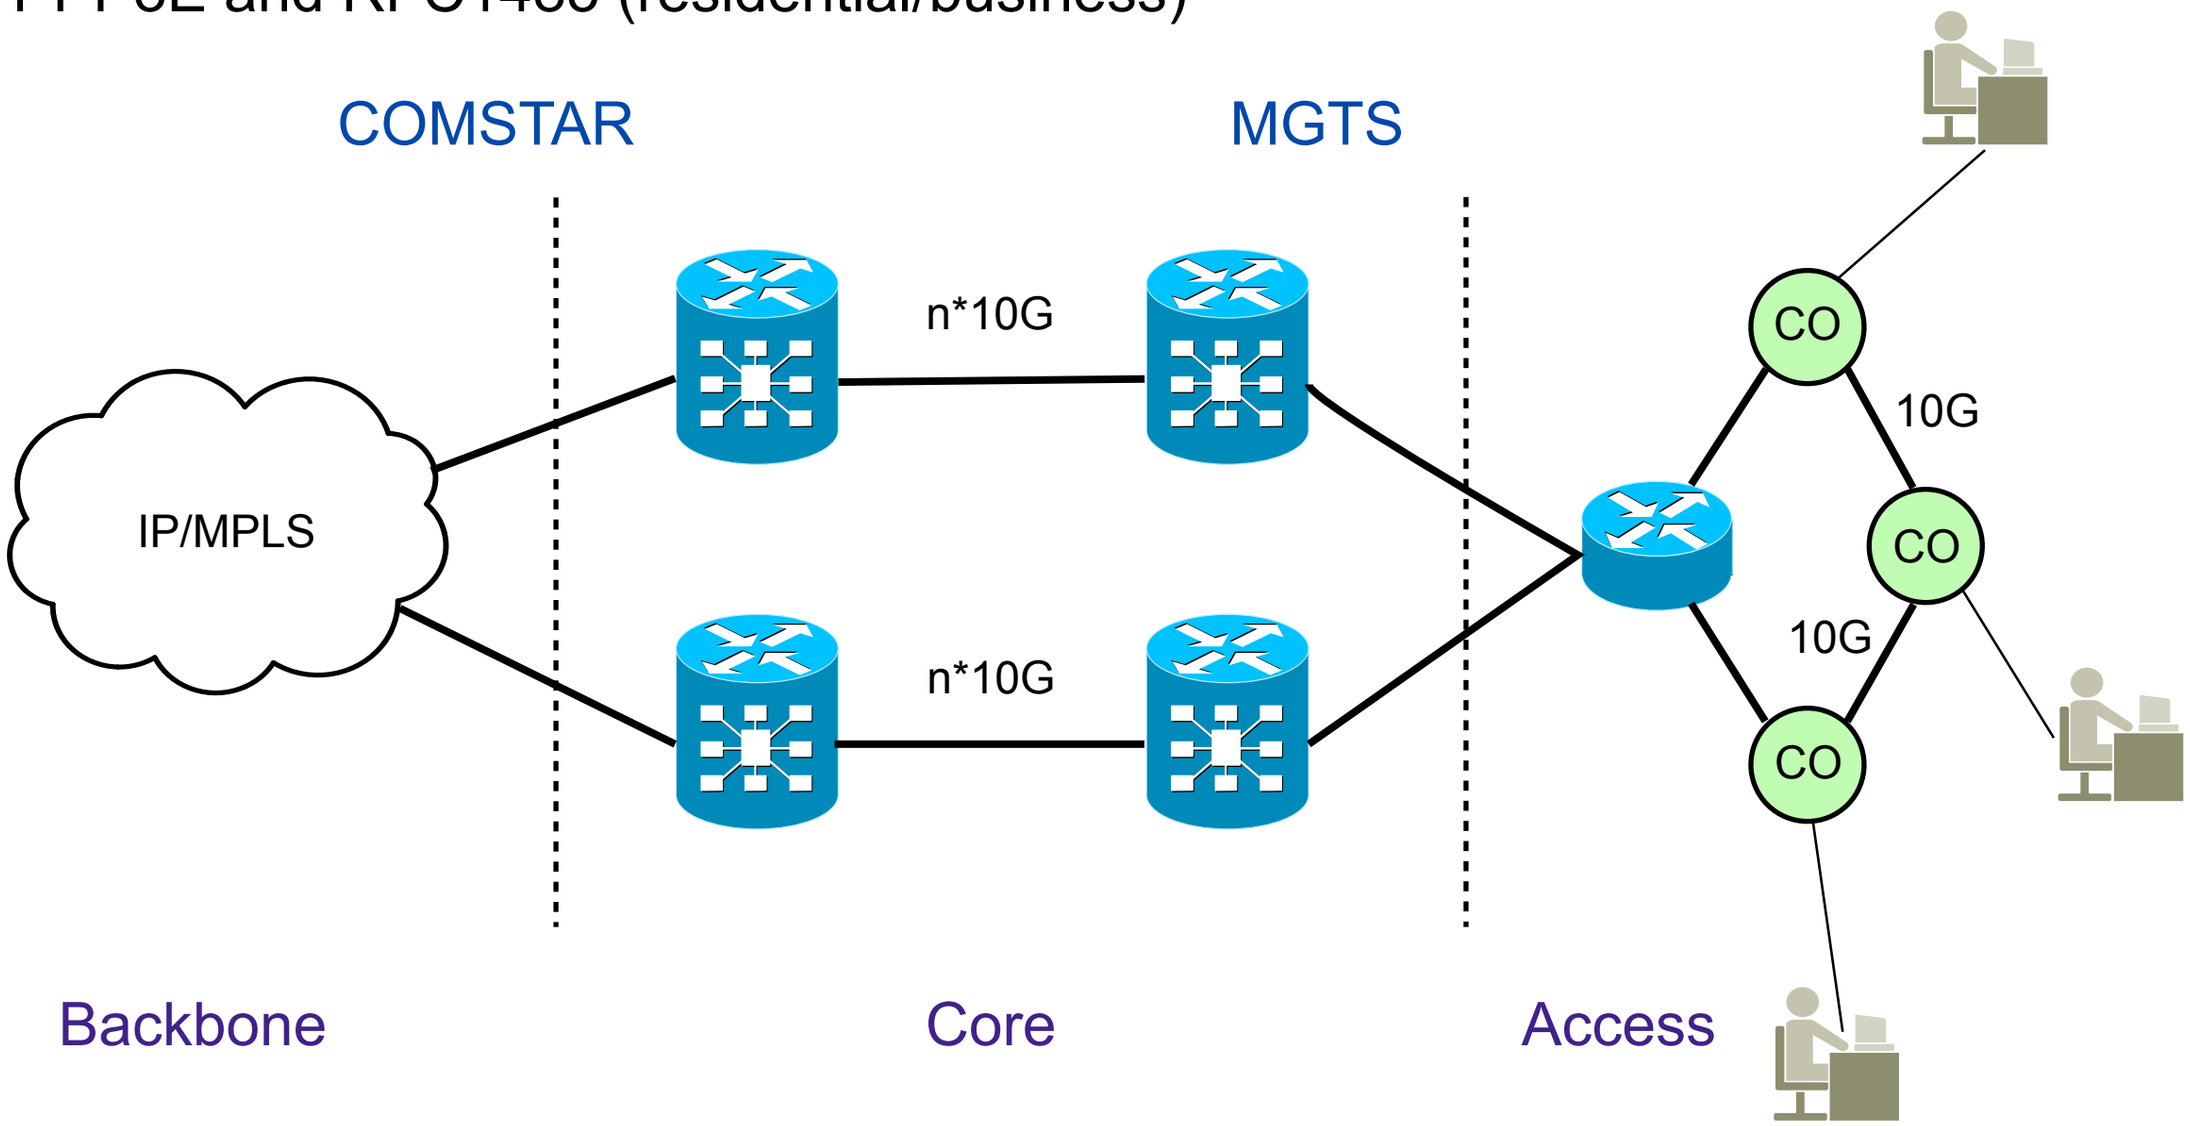

Subnetting

Point-to-point links inside the network /112	>	Assign 2a02:28:1:0::/64 Use /112 for a particular link e.g. 2a02:28:1:0::1:0/112, 2a02:28:1:0::2:0/112 etc. ^^^link #id
Loopbacks /128	>	Assign 2a02:28:1:1::/64 Use /128 for a loopback e.g. 2a02:28:1:1::1:1/128, 2a02:28:1:1::1:2/128 etc. ^^^router #id
Point-to-point IPv6 transit clients /112	>	Assign 2a02:28:1:2::/64 Use /112 for an edge link e.g. 2a02:28:1:2::1:0/112, 2a02:28:1:2::2:0/112 etc.
Server infrastructure /64	>	Assign 2a02:28:2::/48 Use /64 for an application, e.g. 2a02:28:2:1::/64, 2a02:28:2:2::/64 etc.

Interface configuration

```
interface Port-channel3
 ip address 195.xx.xx.85 255.255.255.252
 ip ospf network point-to-point
 ipv6 address 2A02:28:1::21:2/112
 ipv6 enable
 ipv6 nd ra suppress
 no ipv6 redirects
 ipv6 ospf network point-to-point
 ipv6 ospf 8359 area 0

interface Loopback0
 ip address 195.xx.xx.254 255.255.255.255
 ipv6 address 2A02:28:1:1::1:5/128
 ipv6 enable

interface TenGigabitEthernet4/2
 ip address 195.xx.xx.145 255.255.255.252
 ipv6 address 2A02:28:x:x::x:1/112
 ipv6 enable
 ipv6 nd ra suppress
 no ipv6 redirects
```

< Point-to-Point

< Loopbacks

< IPv6 Transit

Routing

- IGP is for loopbacks
- IGP – OSPFv3
- Everything else is iBGP/eBGP AS8359
- Route-reflectors are the same for IPv4 and IPv6
- ‘Full-view’ v6 is around 3300+ prefixes (100 times less than IPv4 FV)

```
#sh bgp ipv6 unicast neighbors 2A02:28:x:x::x:2
advertised-routes | inc Total
Total number of prefixes 3325
```

- Currently far less client IPv6 prefixes to exchange on peerings

```
#sh bgp ipv6 unicast neighbors 2001:xxx::xxxx:0:1
advertised-routes | inc Total
Total number of prefixes 22
```

(compare this to IPv4)

```
#sh ip bgp neighbors 80.81.192.xxx advertised-
routes | inc Total
Total number of prefixes 4500+
```

Routing configuration

Relevant OSPF configuration excerpt

```
ipv6 router ospf 8359
 redistribute connected route-map v6-to-ospf
route-map v6-to-ospf permit 10
 match ipv6 address v6ospf
ipv6 access-list v6ospf
 permit ipv6 2A02:28:1:1::/64 any
```

Relevant BGP configuration excerpt

```
address-family ipv6
 redistribute connected route-map to-bgp6
 redistribute static route-map to-bgp6
 network 2A02:28::/32 route-map to-bgp6-8359:8359
 neighbor 2001:xxx:xx:xx::xx:36 activate
 neighbor 2001:xxx:xx:xx::xx:36 send-community
 neighbor 2001:xxx:xx:xx::xx:36 next-hop-self
 neighbor 2001:xxx:xx:xx::xx:36 maximum-prefix 50
ipv6 prefix-list 8359:localas-v6 seq 5 permit 2A02:28::/32 le 128
route-map to-bgp6 permit 10
 match ipv6 address prefix-list 8359:localas-v6
 set local-preference 200
 set community local-AS
route-map to-bgp6-8359:8359 permit 10
 set local-preference 200
 set community 8359:8359
```

Getting connectivity

- Options – IPv6 transit and/or peering
- IPv6 transit is mostly free (although some Tier-1 try to sell it to you)
- IPv6 transit usually comes as an additional value to IPv4 IPT
- In some cases you can get IPv6 IPT from a Tier-n without being an IPTv4 client
- If you peer extensively – no need for IPv6 transit (yet?)
- More and more IPv6 peers appear across Europe and US
- All major IXes are IPv6-enabled
- Adding new IPv6 peer every other day helps you get ‘full’ IPv6 connectivity without a need for IPv6 transit
- IPT services will change with transition to IPv6 (will we see that Tier-n landscape changing?)

Getting connectivity

- Compare your ‘full view’ to a ‘true’ one (say, HE one)
- Check out an IPv6 route server (hint: route-server.het.net)
- There could be slight difference
- Analyze how many prefixes you’ve got from the major peers (HE gets you roughly half of IPv6 FV)
- Browse public statistics (hint: www.sixxs.net)
- Monitor reachability of foreign and internal resources
- Use external IPv6 tunnel brokers to check your network from the outside

IPv6 traffic levels (native)

- Native traffic still **extremely low** compared to existing IPv4 traffic
- Will not grow unless the clients will consume/generate it
- Difficulties with IPv6 introduction/migration in the access network
- Content players reluctant to modify/expose their legacy and new applications

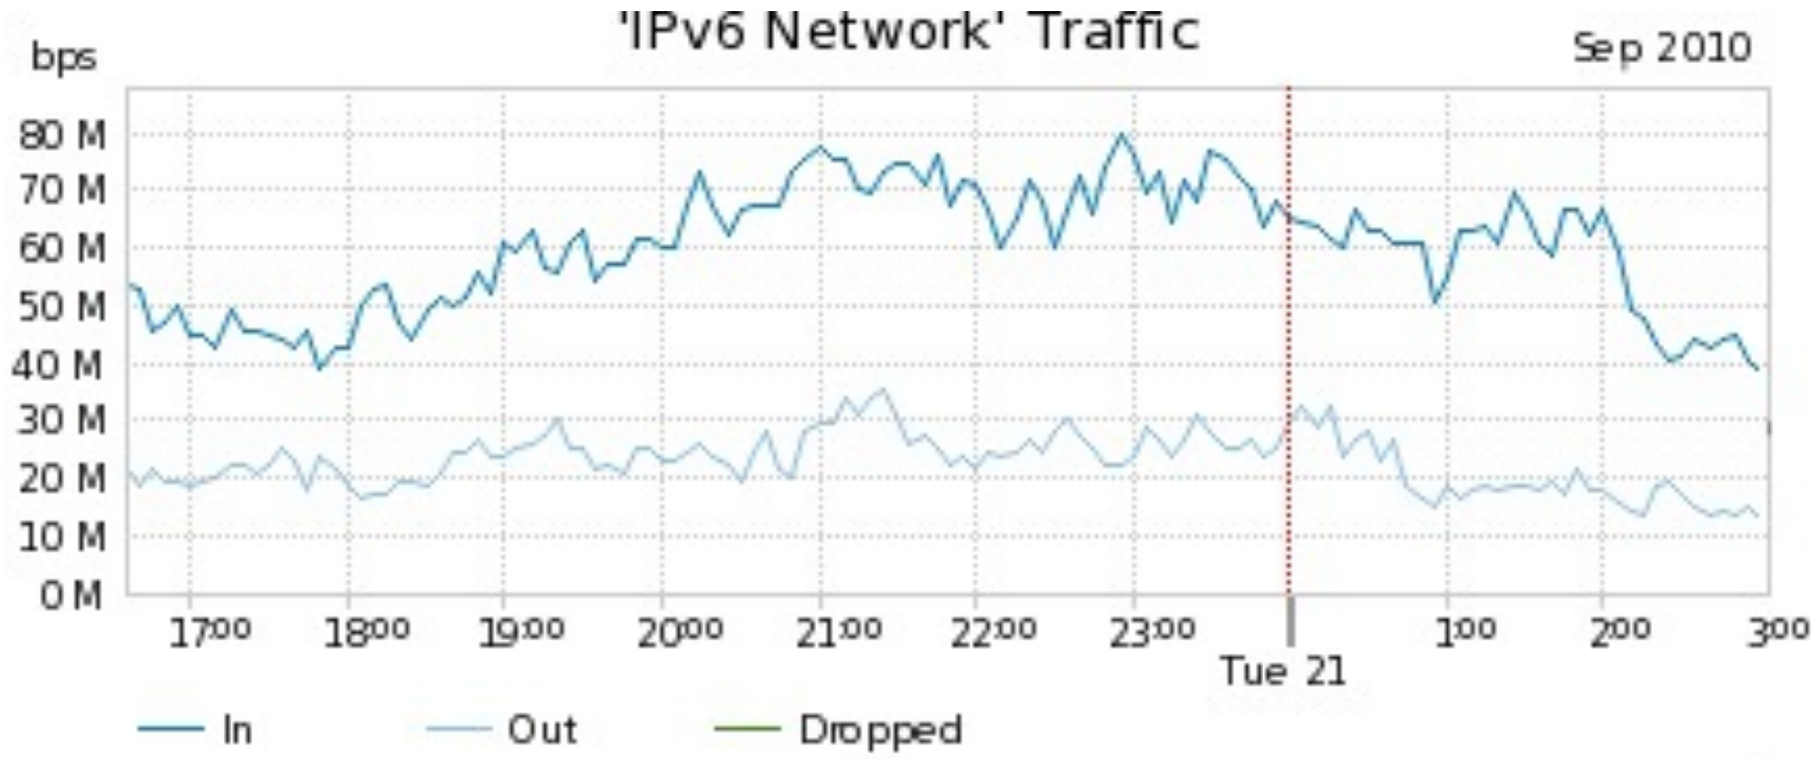

Maximum:	Average traffic:	95th Percentile:
115.23 Mbps	**63.53 Mbps**	**101.82 Mbps**

IPv6 traffic levels (tunnels)

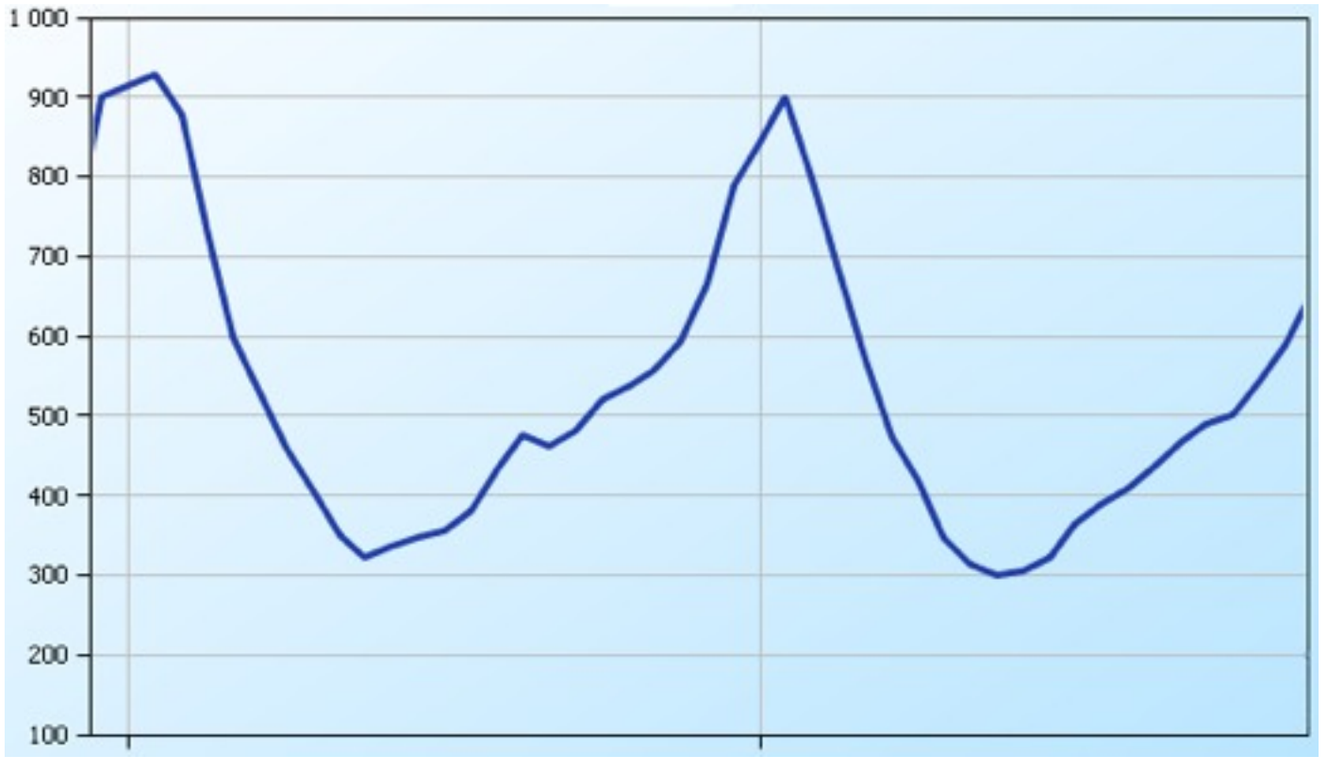

- IPv6 traffic from broadband users at around 1Gbps
- Peak activity aligned with IPv4 (9pm-11pm)
- µTorrent+Teredo?
- 35,000 concurrent sessions
- More analysis ahead

- Protocol 41 at around 15Mbps
- 6in4
- 6to4

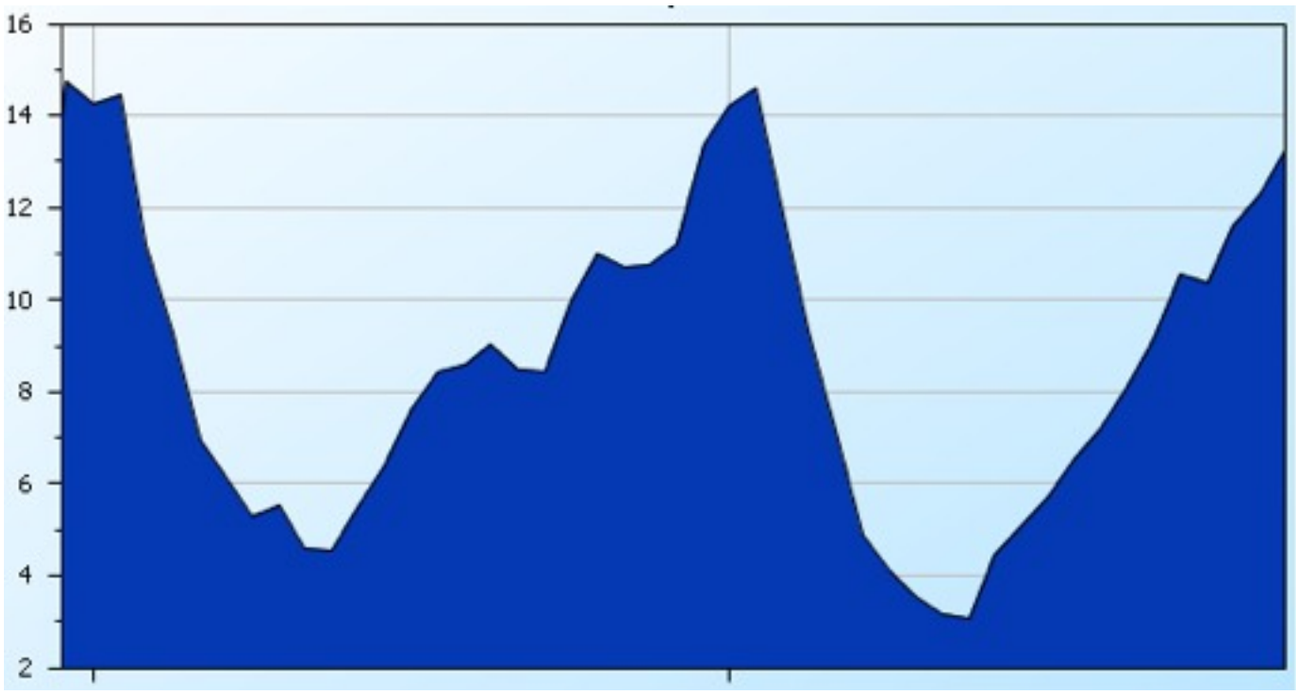

Resources to expose

- Resources that you can easily get up and running on IPv6
  - Authoritative DNS
  - Recursive DNS
  - E-mail (especially retrieval including web-access)
  - Company web sites

```
% dig @195.34.32.116 www.stream.ru aaaa
www.stream.ru. 600 IN AAAA  2a02:28:2::1076
% dig @195.34.32.116 umail.ru aaaa
umail.ru.      600 IN AAAA  2a02:28:2:3::101
```

- Dual-stack for DNS and E-mail
- IPv6 reverse proxy for web sites
- We use
  - Unbound
  - CommuniGate Pro (cluster)
  - Varnish

Hints and findings from operations

- Legacy equipment is painful
  - Our load-balancers did not honor neighbor discovery
  - Static “arp” may help

```
ipv6 neighbor 2A02:28:2:3::101 Vlan13 0000.7f00.0395
```

- Autoconfiguration for servers is not that easy
  - Many die-hard admins are used to memorize v4-ip’s and keep logging in by IP rather than the hostname
  - With IPv6 it’s a trap (mac-derived IPv6 addressing is tough)
  - For such folks we have a trick to offer – use the last octet from host’s IPv4 address

```
dns1.mtu.ru.      600 IN  A      212.188.8.37
dns1.mtu.ru.      600 IN  AAAA   2a02:28:2:1::37
rr1.mtu.ru.       600 IN  A      195.34.32.116
rr1.mtu.ru.       600 IN  AAAA   2a02:28:2:3::116
```

- Do not forget to align OSPF costs for IPv6 with those for IPv4 – it’s a separate thing to have in router configuration, otherwise your IPv4 and IPv6 traffic flows may differ

Hints and findings from operations

- Check out for inconsistencies
  - Having an IPv6 address for a DNS server but no DNS IPv6 socket on it (or no IPv6 interface on requestor) may lead to funny resolver behavior
  - Misconfigurations could be hard to find (for instance, same IPv6 address on both sides of a point-to-point link)
- Configure ACLs for:
  - Network equipment and management access in particular

```
ipv6 access-list deny-any
 deny ipv6 any any
line vty 0 4
 ipv6 access-class deny-any in
```

  - Servers and applications
- Add assigned address space to RIR DB
- Monitor every vital component of your IPv6 implementation
  - Network elements (router loopbacks, p-p, uplinks and peers)
  - Server and applications infrastructure
  - Connectivity

Migrating the users

- Tunnels, 6rd and other ‘non-native’ exposure of your broadband clients to IPv6 will only delay the painful transition
- In your own (and/or well-controlled) network the choice should be NATIVE
- Challenges are:
  - Lack of IPv6 support for broadband technologies at the edge
  - Lack of CPE support for PPP(oe)/native mode
  - Lack of provisioning mechanisms (TR-069 is barely working even in IPv4 scenarios despite being ‘widely supported’)
  - Clients OSes may really vary in flavour, setup, accuracy of configuration, bad applications/drivers, IPv6 implementation issues etc.
  - Mobile devices rarely support IPv6 in a cell environment (Wi-Fi can work just fine, though)

Essential URLs to check

http://www.sixxs.net/tools/grh/

http://www.sixxs.net/misc/usage/

http://www.bgpmon.net/

http://ipv6.he.net/bgpview/bgp-page-complete.html

http://www.tunnelbroker.net/

http://test-ipv6.com/

http://www.getipv6.info/index.php/IPv6_Addressing_Plans

http://www.getipv6.info/index.php/First_Steps_for_ISPs

http://tools.ietf.org/html/rfc3627

http://tools.ietf.org/html/rfc4472

http://technet.microsoft.com/en-us/network/cc917486.aspx

http://en.wikipedia.org/wiki/Teredo_tunneling

http://en.wikipedia.org/wiki/6in4

http://en.wikipedia.org/wiki/6to4

http://version6.ru/ (RU-only)

COMSTAR

Q & A