

Status of IPv6 Rollout at Swisscom

Martin Gysi, 22.10.2014
public

Status of IPv6 Rollout at Swisscom

Agenda

2

Remember IPv6? It's IP with longer addresses!

(Nothing more, nothing less.

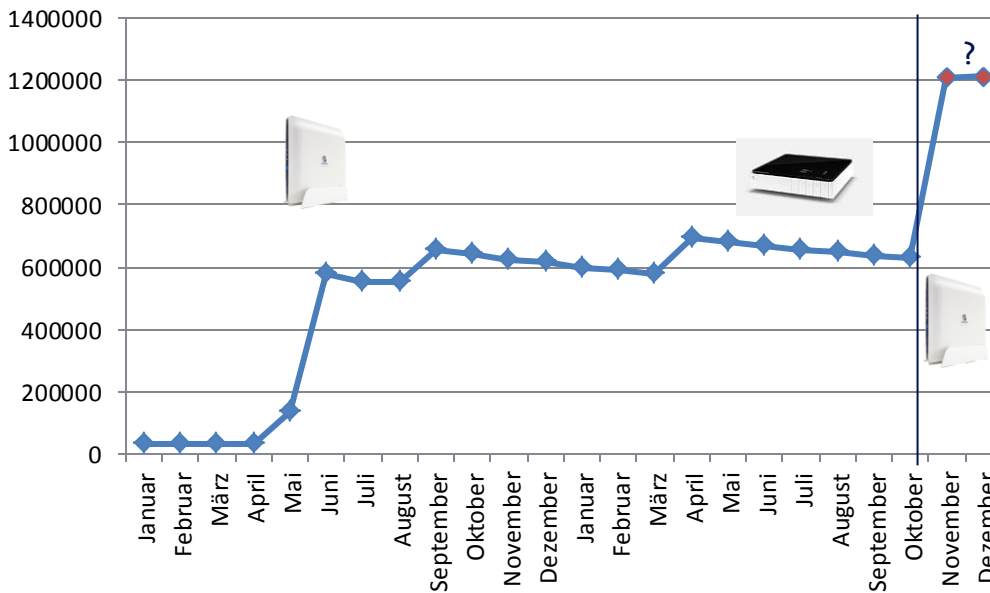
But it's crucial for the future of the Internet)

- Residential
- Datacenter
- Mobile
- Enterprise Customers

Residential Evolution of dual-stack customers

- Remote activation of IPv6 began in March 2013 for ADB «centro» routers

Residential IPv6 customers 2013-14




- IPv6 is *not* yet activated by default. So, due to customer churn, numbers decline after a rollout. New customers are activated periodically
- New firmware for older CPEs will allow turning on v6 *by default*. Rollout depends on VDSL vectoring, don't know when it is going to happen. Should roughly double the number of active users

Residential


Evolution of bandwidth for dual-stack users

- Average peak traffic per customer since start of rollout:
 - Total: +50%
 - IPv4: + 30%
 - IPv6: + 245%
- IPv6 traffic is 26% of total traffic

IPv4 average peak traffic per customer



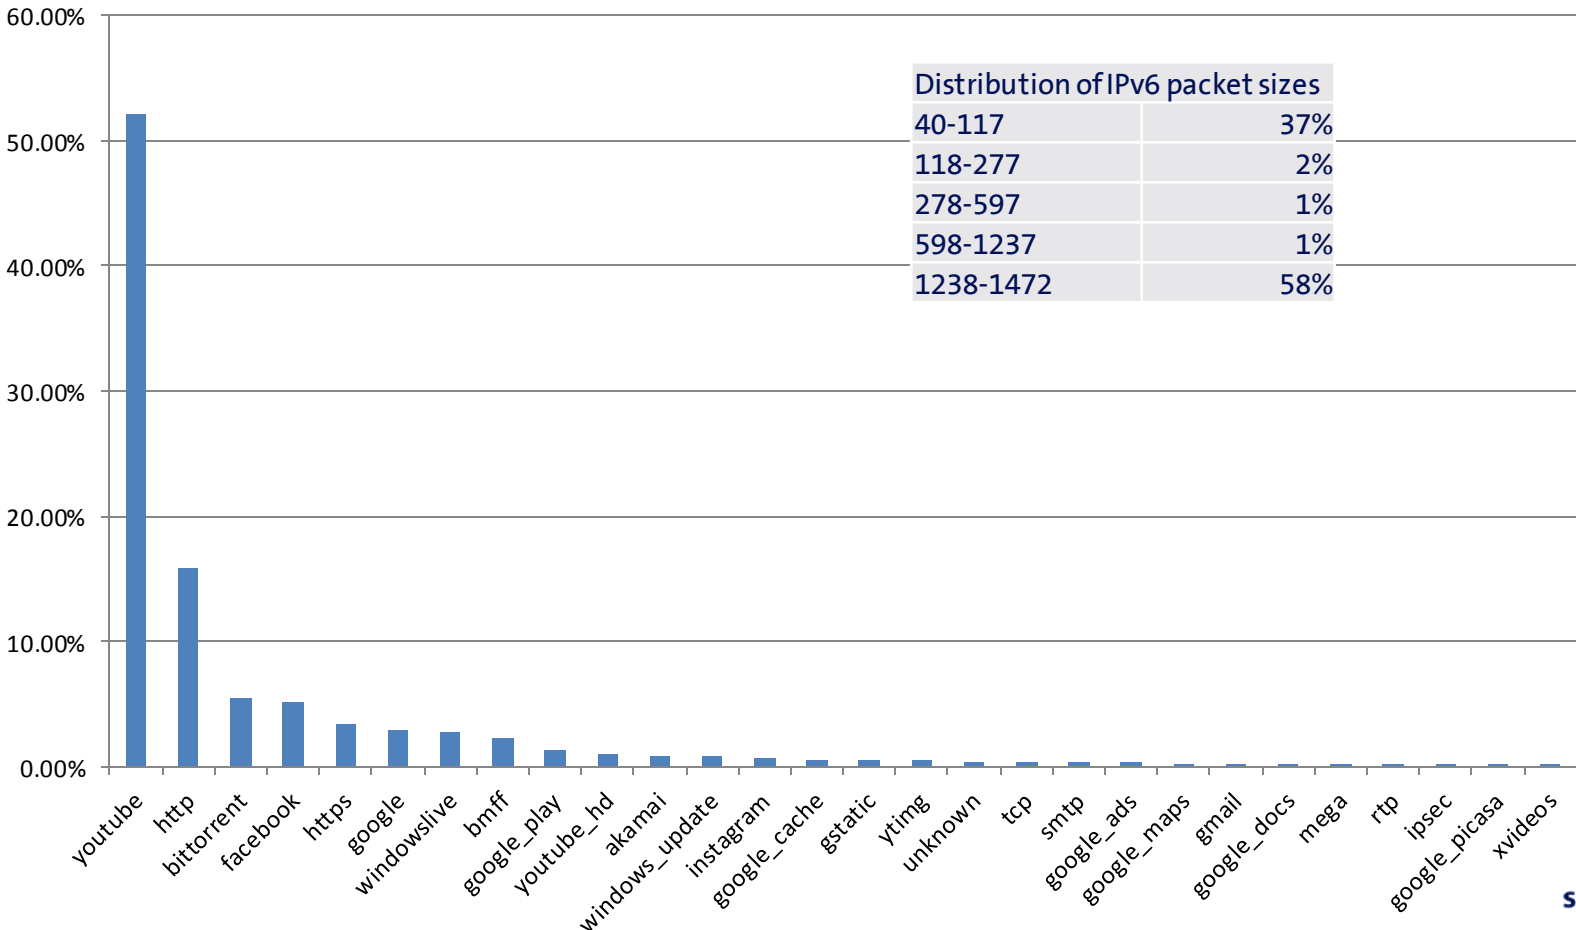
IPv6 average peak traffic per customer



- *Services that are available on IPv6 grow significantly faster than the average!*

Residential IPv6 traffic sources

60% from Google, 5% FB, 35% rest of Internet




IPv6 Rapid Deployment on IPv4 Infrastructures (RFC 5969)

6RD is a *Stateless* Tunnel Technology, Embedding the CE's IPv4 Address into the IPv6 Prefix.

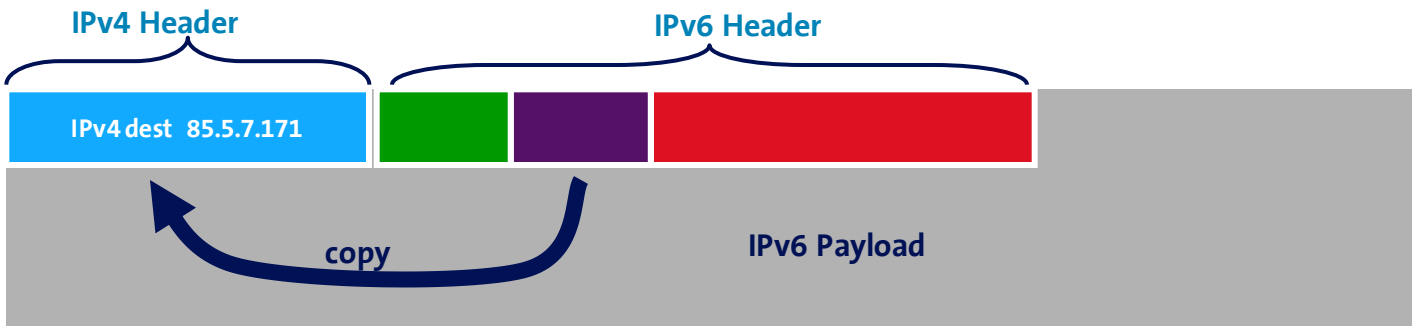
Network topology



IPv6 address format for 6rd




IPv4 header & encapsulated IPv6 packet (downstream)



Residential Border Relay (ASR 9010) technical details

- 6rd encapsulation and decapsulation is performed on «NPU» of line cards
- Three ports of a 24-Port Line Card share 1 NPU
- Service Blade for «exceptional traffic», e.g. ICMP (Ping, Traceroute, Path MTU Discovery)
- First step of rollout: Use every third port on LC.

```
service cgn SC_6rd
service-location preferred-active 0/0/CPU0
service-type tunnel v6rd 6rd1
path-mtu 1472
br
  ipv6-prefix 2a02:1200::/28
  source-address 193.5.29.1
  unicast address 2a02:120c:1051:d010::1
!
address-family ipv4
  interface ServiceApp4
!
address-family ipv6
  interface ServiceApp6
!
!
```



show command: #sh cgn tunnel v6rd 6rd1 statistics

```
Tunnel 6rd configuration
=====
Tunnel 6rd name: 6rd1
IPv6 Prefix/Length: 2a02:1200::/28
Source address: 193.5.29.1
BR Unicast address: 2a02:120c:1051:d010::1
IPv4 Prefix length: 0
IPv4 Suffix length: 0
TOS: , TTL: , Path MTU: 1472

Tunnel 6rd statistics
=====

IPv4 to IPv6
=====
Incoming packet count           : 2452764625
Incoming tunneled packets count : 2452712654
Decapsulated packets           : 2452694282
ICMP translation count         : 20353
Insufficient IPv4 payload drop count : 0
Security check failure drops   : 1483
No DB entry drop count         : 0
Unsupported protocol drop count : 0
Invalid IPv6 source prefix drop count : 16889

IPv6 to IPv4
=====
Incoming packet count           : 3174456512
Encapsulated packets count     : 3174445625
No DB drop count               : 0
Unsupported protocol drop count :

IPv4 ICMP
=====
Incoming packets count         : 51971
Reply packets count           : 1
Throttled packet count        : 21924
Nontranslatable drops         : 9693
Unsupported icmp type drop count : 0
```

```
IPv6 ICMP
=====
Incoming packets count         : 1346
Reply packets count           : 1325
Packet Too Big generated packets count : 2128
Packet Too Big not generated packets count : 0
NA generated packets count     : 0
TTL expiry generated packets count : 0
Unsupported icmp type drop count : 0
Throttled packet count        : 7434

IPv4 to IPv6 Fragments
=====
Incoming fragments count       : 0
Reassembled packet count      : 0
Reassembled fragments count   : 0
ICMP incoming fragments count : 0
Total fragment drop count     : 0
Fragments dropped due to timeout : 0
Reassembly throttled drop count : 0
Duplicate fragments drop count : 0
Reassembly disabled drop count : 0
No DB entry fragments drop count : 0
Fragments dropped due to security check failure : 0
Insufficient IPv4 payload fragment drop count : 0
Unsupported protocol fragment drops : 0
Invalid IPv6 prefix fragment drop count : 0

IPv6 to IPv4 Fragments
=====
Incoming ICMP fragment count   : 0
RP/0/RSP0/CPU0:ls5ic20p-brd001#
```


Residential 6rd vs. CG-NAT



- As we don't have enough public IPv4 addresses, a number of (low-end) customers get RFC 6598 addresses combined with NAT.
- → Restricted IPv4 access.
- Of course, their IPv6 /60 prefix is public!
- 6rd uses IPv4 as transport between CPE and Border Relay, in a completely stateless fashion
- RFC 6598 addresses are routed (non-natted) to the BR.
- *Saving IPv4 addresses and deploying IPv6 are two separate, distinct problems that must both be solved!*

Reverse DNS in IPv6 – how?

- In IPv4, ISPs assign a reverse DNS entry for every IP address that they assign to their customers
 - `243.109.192.178.in-addr.arpa. IN PTR 243-109.192-178.cust.bluewin.ch.`
- In IPv6, that's a hard thing to do: Swisscom has 2^{100} addresses for residential customers alone. No DB or flat file could hold so many entries. (See Lee Howard from Time Warner Cable: draft-howard-isp-ip6rdns-06)
 - **No entry:** Common. But not best practice
 - «**No such domain (NXDOMAIN)**» entry: Gives you no answer, but at least you know that somebody owns this space
 - «**Wildcard**» entry: Return the same answer for all requests within a certain space. That's what Swisscom does. Disadvantage: Forward and reverse entries will never match.
 - `*.0.2.1.2.0.a.2.ip6.arpa. IN PTR dynamic.wline.6rd.res.cust.swisscom.ch.`
 - **Generated-on-the-fly** entry: More CPU intensive (especially with DNSsec). Advantage: Matching forward end reverse entries
 - **Dynamic DNS:** ISP's DNS delegate to Residential Gateways, hosts automatically register their names there.

Datacenter Overview

- «Product IT networks» (PIN) datacenter:
 - Routing (IPv6 on MPLS) has been in place for 40 weeks. Services are tested end-to-end in lab and work fine.
 - Load balancers *are still nok*, due to software issues not related to IPv6. So no live services are running dual-stack today, no DNS, no bluewin.ch, no mail.
 - Next IOS release for LB will be available in November. → another round of IPv6-testing.
- Swisscom IT Services datacenter:
 - swisscom.ch: Running dual-stack since «World IPv6 Launch», using two dedicated F5 load balancers

Datacenter

Some learnings and caveats

- PIN runs on old hardware (c7609 with SUP720 and PFC3B/C line cards)
- Routes: Limited to max 256 k
 - 32 k IPv6 routes by default (an IPv6 route counts as two IPv4 routes)

```
c6500#show mls cef maximum-routes
FIB TCAM maximum routes :
=====
Current :-
-----
IPv4 + MPLS      - 192k (default)
IPv6 + IP Multicast - 32k (default)
```

```
c6500(config)#mls cef maximum-routes ?
ip          number of ip routes
ip-multicast number of multicast routes
ipv6       number of ipv6 routes
mpls      number of MPLS labels
```

- ACLs: 128 Bit TCAM. So not enough room for IPv6 address + port, ACL exceeding 128 Bits are treated in software (*bad idea!*)
 - Activate «ACL compression»:

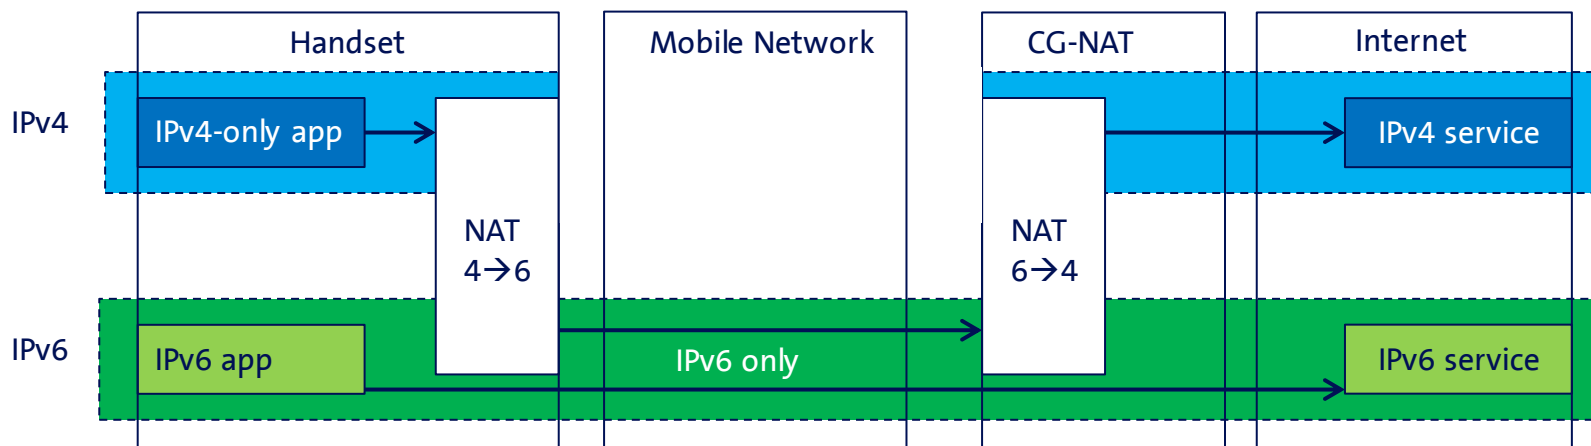
```
mls ipv6 acl compress address unicast
```

- Verify ACL is not punted to CPU (watch for «punt» in output):

```
sh tcam interface gigabitEthernet 1/10 acl in ipv6
```

Mobile

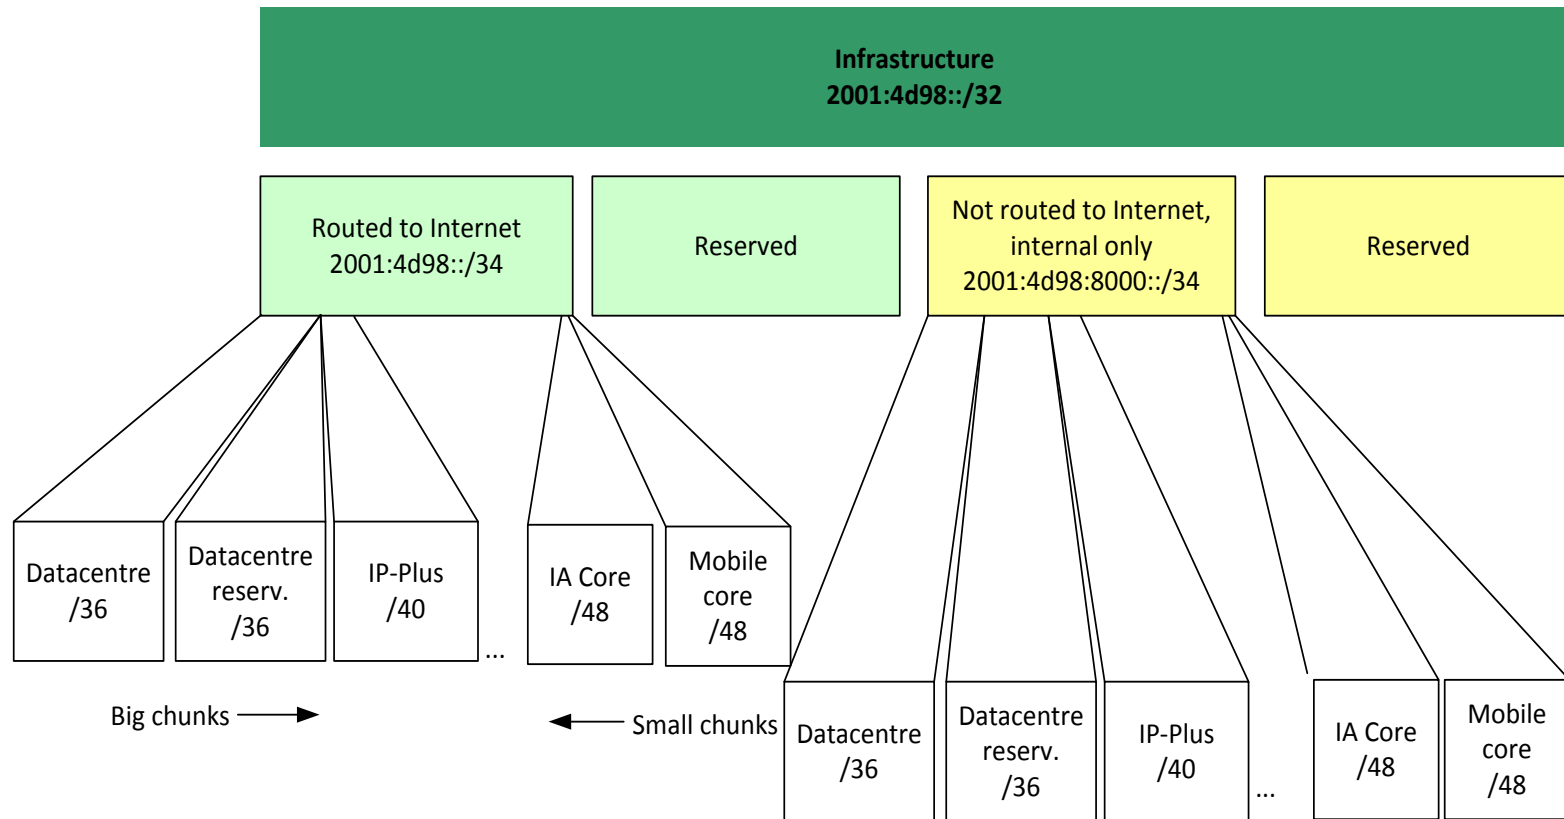
- VoLTE (VoIP over LTE) will use an IPv6-only APN (~VRF in mobile). Launch announced for mid 2015.
- No launch date for IPv6 on the Internet APN. But should be achievable with the experience gained from the VoLTE rollout.
- We plan on using 464XLAT (RFC 6877) to get rid of IPv4 in the mobile network, but still enable IPv4 connectivity for applications that need it.



Enterprise customers

- Currently working on a dual-stack offer for SME customers.
- IPv6 will be part of the «All-IP» product bundles (“my KMU office”), to be launched mid-2015.
- Mid-range offer will include
 - Static /48 subnet
 - CPE supporting DHCPv6 prefix delegation, making it possible to chain routers
- IPv6 Internet Services from IP-plus have been available for more than ten years, on a best-effort basis ...
- LAN-I (L3-VPN): On customer request. We currently have one pilot site running (Alte Kantonsschule Aarau).

Addressing the infrastructure 2001:4d98::/32



- Organized by service, zone, network, device type (e.g. mgmt) etc. Default /48 per service or zone. Larger allocations require justification by an address concept.
- Geographical significance (if any) only within assigned block, local responsibility
- Every service or zone is assigned with an Internet-routed (x_{16}) and an internal-only ($x_{16} + 8000_{16}$) block. Use according to need.

Deploy IPv6, help put Switzerland back on top!

- Belgium: 27.1%
- Luxembourg: 11.32%
- Germany: 11.15%
- Switzerland: 10.57%



IPv6. It's not an option!

Contact information

17

Swisscom AG
Martin Gysi
Network Development
Binzring 17
CH-8045 Zürich

martin.gysi@swisscom.com
www.swisscom.ch