The pitfalls of IPv6 address choice

address allocation and RIPE policies

Gert Döring, RIPE address policy working group

devoteam

7th Nov. 2013

Table of Contents

Table of Contents

1. High level overview and introduction
2. IPv6 address distribution rules (“address policy”)
3. IPv6 case studies

DEVOTEAM
Consulting • Solutions • Expertise

هيئة الإتصالات وتقنية المعلومات
Communications and Information Technology Commission

Saudi Arabia IPv6 Task Force

IPv6 – enough addresses for everyone?

- IPv6 truly has „enough“ addresses:
  $2^{128} = 3.4*10^{38}$
- but still, sometimes network plans get refuted with „it cannot be done that way“ and „that would be against the rules“

- who makes „the rules“?
- what are „the rules“?

- … and given these rules, what are „best practices“ for typical usage scenarios?

introducing the speaker

- Dipl.-Phys. Gert Döring

- studied physics at the technical univ. Munich
- working as freelancer for SpaceNET AG, Munich
  - Internet-Service-Provider for mid-sized enterprises
- first contact with IPv6 in May 1997
- IPv6 allocation for SpaceNet in August 1999
- working with Devoteam on IPv6 in KSA since 2008
- since 1997 actively participating in RIPE community

- ... and still convinced that IPv6 is the right way to scale the Internet to the next decades!

address management: tree structure

- IP addresses need to be globally unique
- most easily achieved by a central registry
- or using a distribution tree, taking regional differences into account by regional registries
- root of the tree: ICANN/IANA
- registry for europe and the middle east: the RIPE NCC in Amsterdam

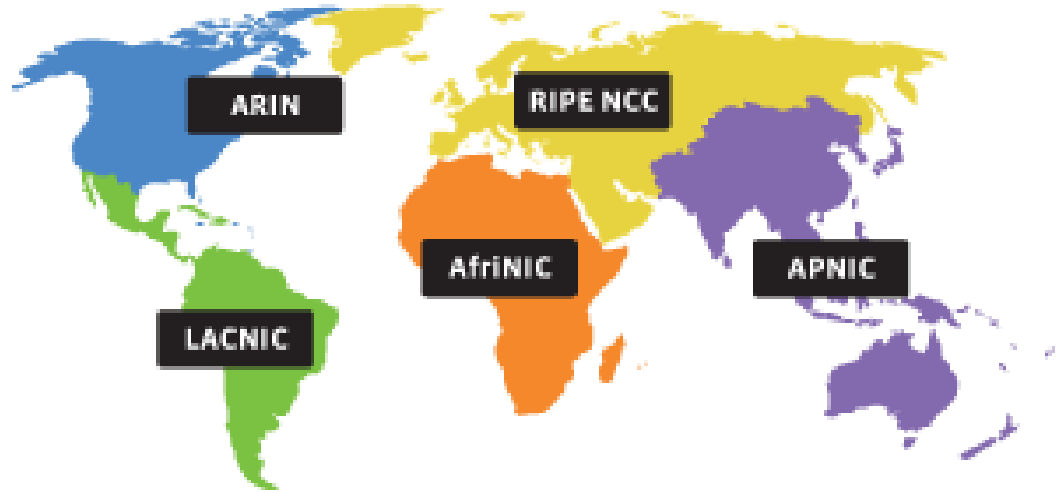

what is the RIPE NCC?

- RIPE „Network Coordination Center“
- non profit organisation, located in Amsterdam
- central role: neutral and independent „secretariat“ for all sort of „Internet coordination activities“
- distributes IPv4 and IPv6 addresses and AS numbers to members (usually ISPs) and to end users
- rules for address distribution („Address Policy“) are created by the RIPE *community* in a „bottom-up“ process, and are then implemented by the RIPE NCC
- RIPE community = all (interested) Internet users in the RIPE service region (Europe, middle east)

IPv6 addresses: provider aggregateable

- intended receipients: „Internet Providers“
  - connecting large numbers of end users to the Internet
  - thus needing large block of IP addresses to give to end users

- 1000s or millions of customers show up in the global routing system under a *single* aggregated route
- requirement: RIPE membership (ca. 1800 EUR/yr)
- standard allocation size is /29.../32, „just ask for it“
- larger networks get more space, e.g. /19 for Germany‘s DTAG or France Telecom (needs to be justified)
- service provider is assumed to assign whole *networks* to end users, in the /48.../64 range (typical: /48 or /56)
- for most networks, allocations are so big that they „last forever“

IPv6 addresses: provider aggregateable (2)

- *liberal* distribution to end users explicitely *encouraged*:
  - /64, if the end user only has a single subnet
    (enough addresses for arbitrary number of machines)
  - /56, if the end user has „multiple networks“
    (this is 2^8 = 256 x /64 subnets)
  - /48 for „larger“ end user networks
    (this is 2^{16} = 65536 x /64, or 256 /56)
- end user assignment size (up to /48) is a local decision by ISP, without having to get approval by RIPE NCC
- individual assignments to not have to be registered in the RIPE database (privacy!)
- → network management gets much easier, less time consuming and more pleasant

IPv6 addresses: provider independent

- intended receipients: end users that are not able (or do not want) to use addresses from their provider‘s PA block
  - typical use case: multihoming to multiple ISPs with BGP
- PI addresses are given by RIPE NCC to end user, not to provider
  - typical size in IPv6: /48 network
- every single PI prefix used on the Internet shows up in the global routing table, impacting all routers world wide → some restrictions have to be met:
  - must not be used to number „third parties“
  - requires some contractual agreements to be in place
  - yearly fee of 50 EUR per PI prefix
  - ~~[must be used for BGP multihoming]~~ removed

confusing terms: „PA“ = „not independent“?

- at the distribution level below the RIPE NCC…
  - both „PA“ and „PI“-Space are just „a block of IPv6 addresses“
  - both are fully independent from any other ISP or RIPE member(!)

- PA *block* given by NCC to ISP is „independent“

- „independence“ of addresses becomes relevant when space is distributed all the way to the end user:
  - PA space is handed out by a particular Internet provider to their customers, and can not be taken away to the next provider (→ renumbering is needed)
  - PI space is „tied to“ the *end user*, and can be taken along to the next ISP

address hierarchy visualized

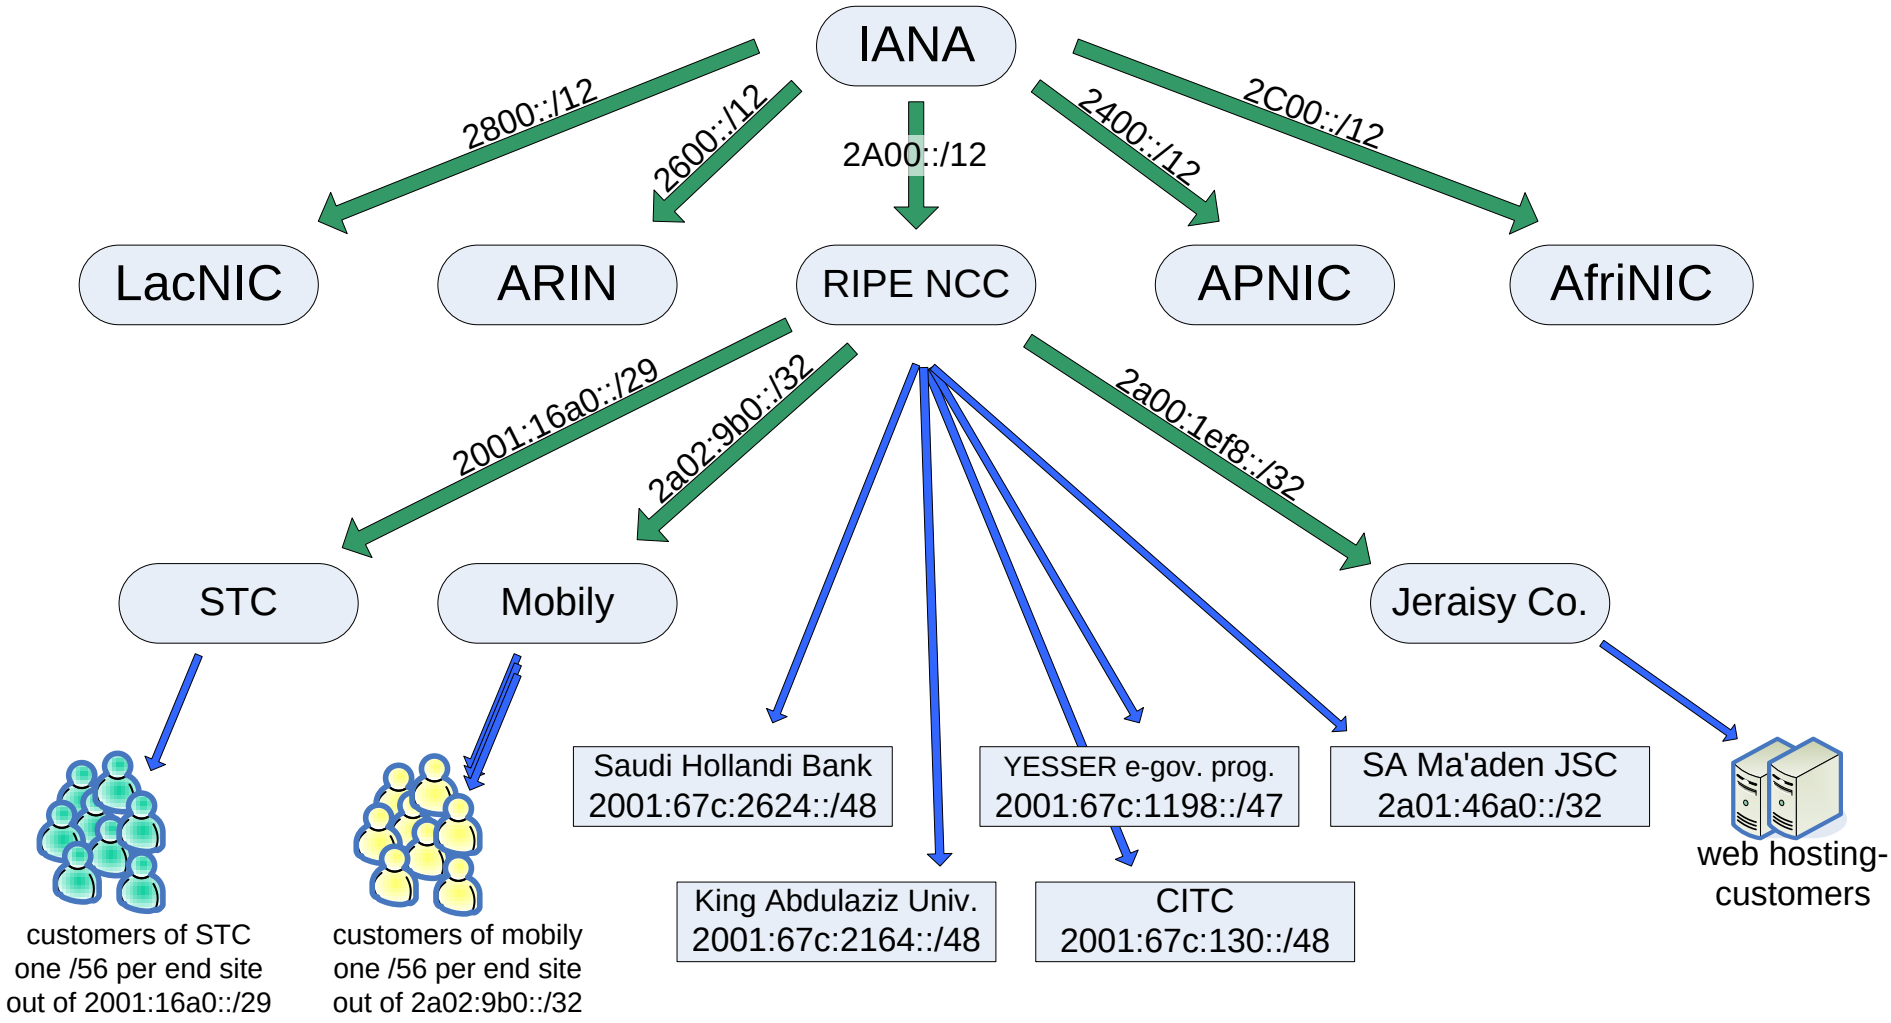

so, which address space is „right“?

- the choices…:
  - PA space from one of the local providers
  - PA allocation from the RIPE NCC („I am an ISP, I need to give out PA“)
  - PI assignment from the RIPE NCC
  - private space (Unique Local Addresses, ULA)
- „right“ decision depends on individual network
  - cost for RIPE membership and BGP routing
  - flexibility in choosing and changing ISP
  - cost of address renumbering when changing ISP
- some case studies to help decide…

case study: Saudi Telecom (STC)

- „typical“ Internet Provider:
  - large number of end users (multiple million)
  - own network, routing and BGP expertise present
  - costs for RIPE membership not really relevant
  - renumbering (including customers!) not possible

- → „right“ decision is to become a RIPE member and ask for a (large) PA network block
  - STC received 2001:16a0::/29 from RIPE NCC
  - assuming an assignment size of /56 to end users…
  - /29 allocation has space for 2^{56-29} = 134 million /56 assignments to end users

case study: barber shop next door

- „typical“ end user
  - no network know-how on site
  - no externally visible servers in network
  - renumbering when changing ISP can be done in a fully automated fashion

- → „right“ decision is to use a PA assignment from one of the local ISPs
  - ISP assigns a /56 to customer router (CPE) using DHCPv6 prefix delegation
  - CPE assigns a /64 to each connected LAN segment
  - hosts use IPv6 autoconfig to configure addresses
  - printers etc. reached by „service discovery“ (mDNS)

case study: insurance company

- fairly demanding requirements in central office
  - usually complex internal network structure
  - travelling employees access CO network using VPN
  - renumbering network would be difficult and expensive
  - quite often multiple providers for higher resiliency
  - but not „typical“ ISP business
- → approach for central office: PI space from RIPE
  - stable /48 network, independent from providers
- → alternative approach: private addresses (ULA) and using proxy/NAT for Internet access using PA space from ISP
  - no impact on global BGP routing table

case study: insurance company (PI)

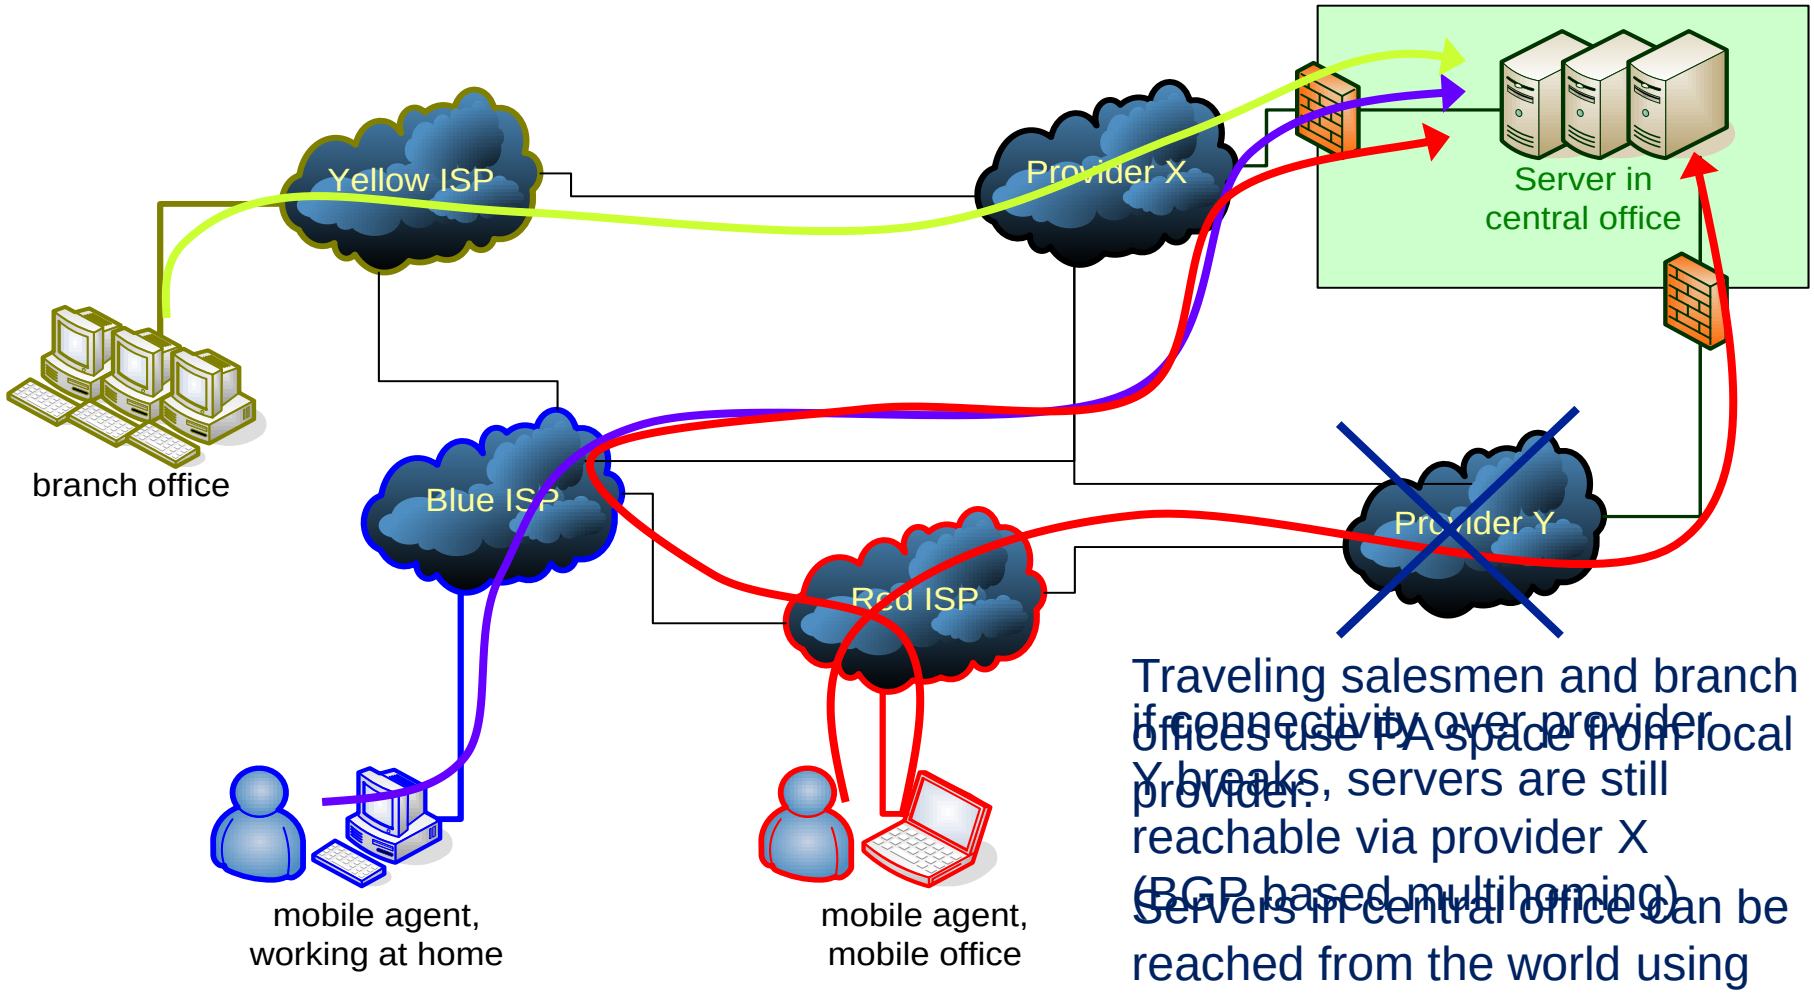

Traveling salesmen and branch offices use PA space from local provider.

if connectivity over provider Y breaks, servers are still reachable via provider X (BGP based multihoming)

Servers in central office can be reached from the world using their PI addresses.

case study: insurance company (ULA)

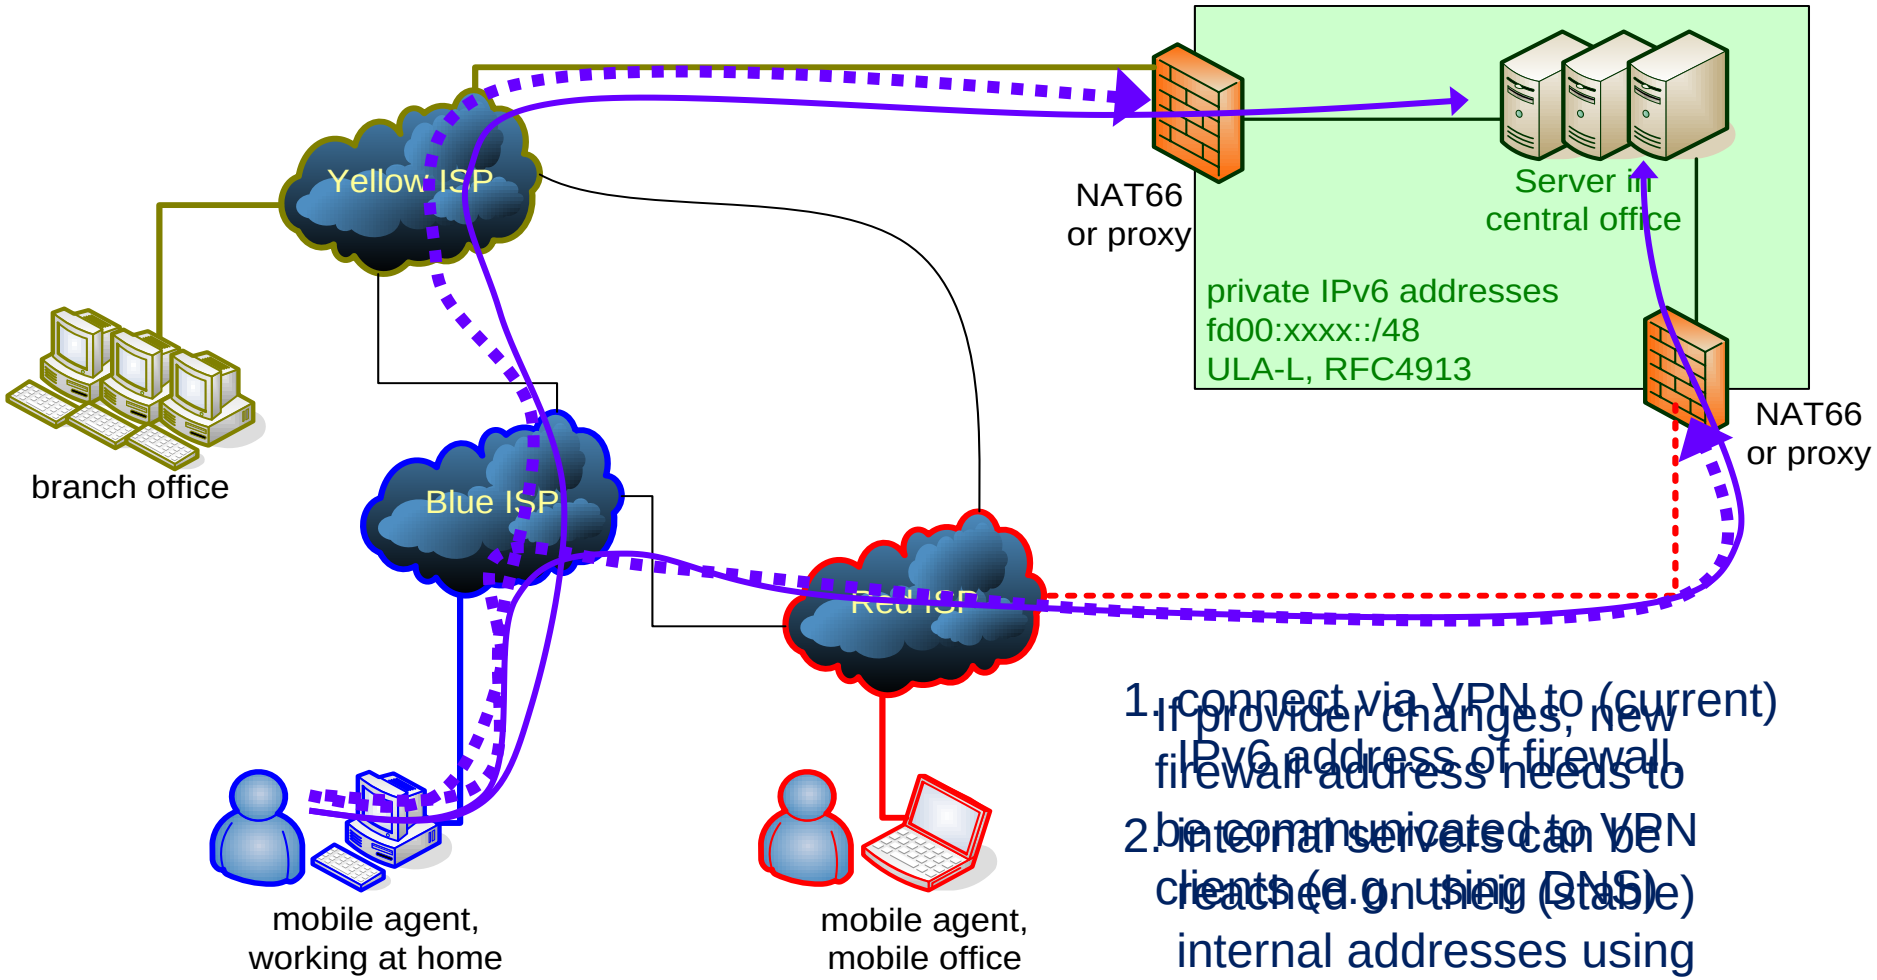

1. connect via VPN to (current) IPv6 address of firewall
2. internal servers can be reached on their (stable) internal addresses using VPN.

If provider changes, new firewall address needs to be communicated to VPN clients (e.g. using DNS)

case study: (small) web hosting provider

- not a „typical“ Internet Service Provider
  - no strong need for /32 PA network block, because in a /48 PI, there are enough IPv6 addresses to number all web hosting customers
  - quite often operating on limited budget: RIPE costs relevant
  - renumbering completely non-practical (as lots of customers would need to change their setups)

- → one would expect „IPv6 PI from RIPE“ to be the answer
  - /48, enough address space, independent
  - but: currently there is a restriction on using IPv6 PI space to number „third party“ devices („no sub-assignments permitted“), and web hosting customers are considered „third parties“ – so PI space gets denied then
  - option 1: RIPE membership, get their own /32 PA block from RIPE NCC
  - option 2: *change RIPE policy*

changing RIPE policies?

- every RIPE policy is a compromise
  - global routing, conservation, „good stewardship“
  - documentation
  - financing RIPE NCC, ...
- some of the policies are done with best intentions, but do not work out „right“ in practice – and sometimes this is only recognized years later, because *environment* has changed
- RIPE Address Policy Working Group is the forum for this:
  http://www.ripe.net/ripe/groups/wg/ap
  mailing-liste: address-policy-wg@ripe.net
- ongoing discussions in the address policy realm:
  - unification of IPv6 PA and PI to „just numbers“
  - clean up old stuff from IPv4 policy that is no longer relevant

QUESTIONS & ANSWERS


Thank you