IPv6 Address Questions - what? how? why?

Gert Döring

SpaceNet AG, Munich

GUUG ECAI6 2007, Cologne, Germany

Typical Reactions...

- but this is *so* wasteful!
- yes, and people have said that 640K would be enough!
- where/how can I get IPv6 addresses?
- It's all RIPE's fault! RIPE is so evil!
- all these are frequently heard reactions from IPv4 admins
  $\Rightarrow$ goal of this talk: clear up some of the confusion

Overview

- IPv6 address distribution
- IPv6 address *math*
- special addresses: private, and provider independent
- who or what is RIPE?
- how to *change* RIPE policies

IPv6 addressing, vs. IPv4

- in IPv4, most important goal was *conservation*
  $\Rightarrow$ long-term IPv4 users really live this by heart
- in IPv6, we have enough addresses
- this means: priorities behind address distribution shifted
  - automatization and ease of use
  - ease of renumbering, ease of changing providers
  - limit routing table growth (aggregation)
  - and still, conservation

binary math recapitulation

- a single bit can have 2 values: 0, 1
- two bits can have 4 values: 00, 01, 10, 11
- for every additional bit, multiply # by 2 $\Rightarrow 2^n$ values

bits	#	values
1	$2^1 = 2$	0 + 1
2	$2^2 = 4$	00, 01 + 10, 11
3	$2^3 = 8$	000, 001, 010, 011 + 100, 101, 110, 111
4	$2^4 = 16$	0000, 0001, 0010, 0011, 0100, . . . , 1110, 1111
8	$2^8 = 256$	00000000, 00000001, 00000010, . . .
16	$2^{16} = 65536$	0000000000000000, 0000000000000001, . . .
32	$2^{32} = 4\ bio$	

Adressing on LAN networks

- every network is assigned a /64 prefix
- hosts can use this /64 plus a 64 bit host ID (normally based on Ethernet address) to *auto-config* 128 bit IPv6 address
- 64 bits for LAN is definitely wastive
  - that's $2^{64} = 18446744073709551616$ addresses!
- benefits considered significant enough:
  - *stateless* autoconfig, no need for DHCP server
  - no more misconfigured subnet masks
  - other uses (crypto) really want 64 bit host ID
- and: there are *enough* /64s (2^{64}) available

Adressing: Provider to Customer

- Customers ("end sites") will always* receive a /48
  - $\Rightarrow$ space for $2^{16} = 65536$ subnets
  - e.g.: a network of 200 locations with 200 subnets each
  - *large enough* for all but the biggest customers
- for a normal SoHo customer, this is clearly too big
- but is it *too* wastive?
  - there's $2^{45} = 35184372088832$ /48s in FP001
  - if we assume 10 billion people on earth, there's still 3500 /48 for every one of them $\Rightarrow$ *enough* space!

Addressing: Provider to Customer (2)

- actually, the policy is a bit more flexible
  - standard customer assignment size is **/48**
  - customers with only a single subnet get a **/64** (standard size for single subnets)
  - customers with only a single host get a **/128**
- but the current recommendation is "always assign /48"
- if a single customer has multiple non-interconnected locations, assigning one /48 *per site* is OK.
- larger netblocks can be assigned after getting RIPE approval
- main focus: *no discussions, ease of address management*

and to Providers?

- ultimate source for IPv6 (and v4) addresses is called IANA
- IANA distributes /12 blocks to 5 *regional* registries, RIRs
- regional registry for Europe is called *RIPE NCC*
  (other RIRs: ARIN, LacNIC, AfriNIC, APNIC)
- to get address space from RIPE, a provider needs to:
  - become a RIPE member, called LIR (Local Internet Registry)
  - sign a contract with the RIPE NCC
  - pay signup fee
  - ... and *ask* for an IPv6 block for his customers
- yearly RIPE fee is between 1300..5500 EUR, depending on size

RIPE to Providers (2)

- size of allocation to providers is variable
- default size is /32 ($2^{16} = 65536$ /48s)
- larger providers can get larger blocks
  - existing IPv4 customer base, or deployment plans
  - internal aggregation taken into account - HD-Ratio
- example: Deutsche Telekom: 2003::/19
  - enough space for $2^{48-19} = 536M$ /48s
  - multi-level internal aggregation necessary
- underlying idea: providers only need one network block, ever

Addressing: Summary

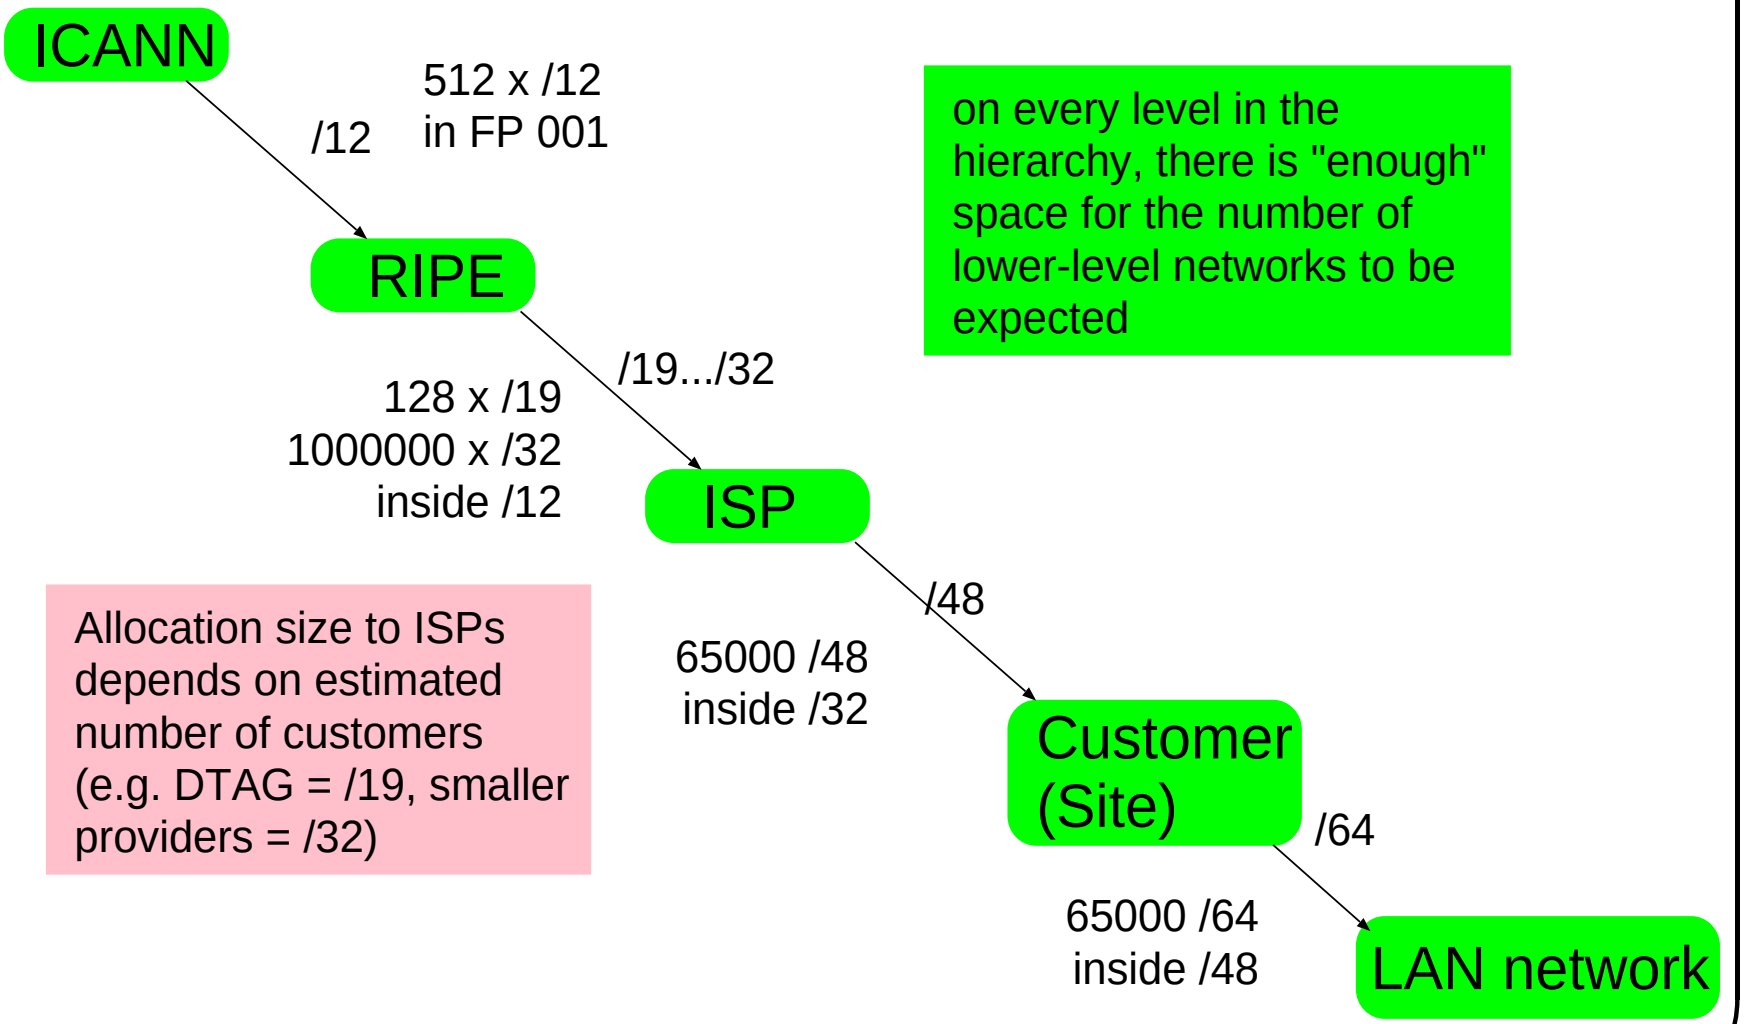

Private (RFC 1918) IPs?

- networks not connected to the Internet have no ISP
  ⇒ cannot get IPv6 addresses??
- on single-link networks (ad hoc), there is always the link-local address, fe80::/64 + host-id (autoconfig)
- as soon as multiple networks and routers come into picture, link-locals are not sufficient
- equivalent to IPv4's RFC 1918 networks is called ULA
  - Unique Local Addresses, RFC 4193
  - important goal: no address collisions with other networks
  - ULA-L: fd00::/8 + 40 random bits = /48 for site
  - ULA-C: fc00::/8 + centrally managed bits = /48
    (under discussion in IETF and RIRs, unclear outcome)

Provider Independent (PI)

- in the design of IPv6, *aggregation* was very important
- Provider Independent addresses are non-aggregateable, and every network needs a slot in the global routing table
- addressing policy space needs to balance
  - global cost (every ISP router would need to know this prefix)
  - stability of the routing system
  - convenience for the end site
- thus, no IPv6 PI space in RIPE land (yet)
- discussions are ongoing, though

what about multihoming?

- multihoming = having two providers for redundancy
- in IPv4, two common ways for multihoming exist
  - dual IPv4 uplinks, double NAT (small networks)
  - PI addresses, BGP (larger networks)
- in IPv6, it's different-but-similar
  - smaller networks get two /48 from two providers
    - PCs have two IPv6 addresses, pick "best" one
    - sessions survive outages with *shim6* protocol
  - large networks can become LIRs, get /32, use BGP
- this is a *very hot* political discussion right now

IETF-style IPv6 multihoming

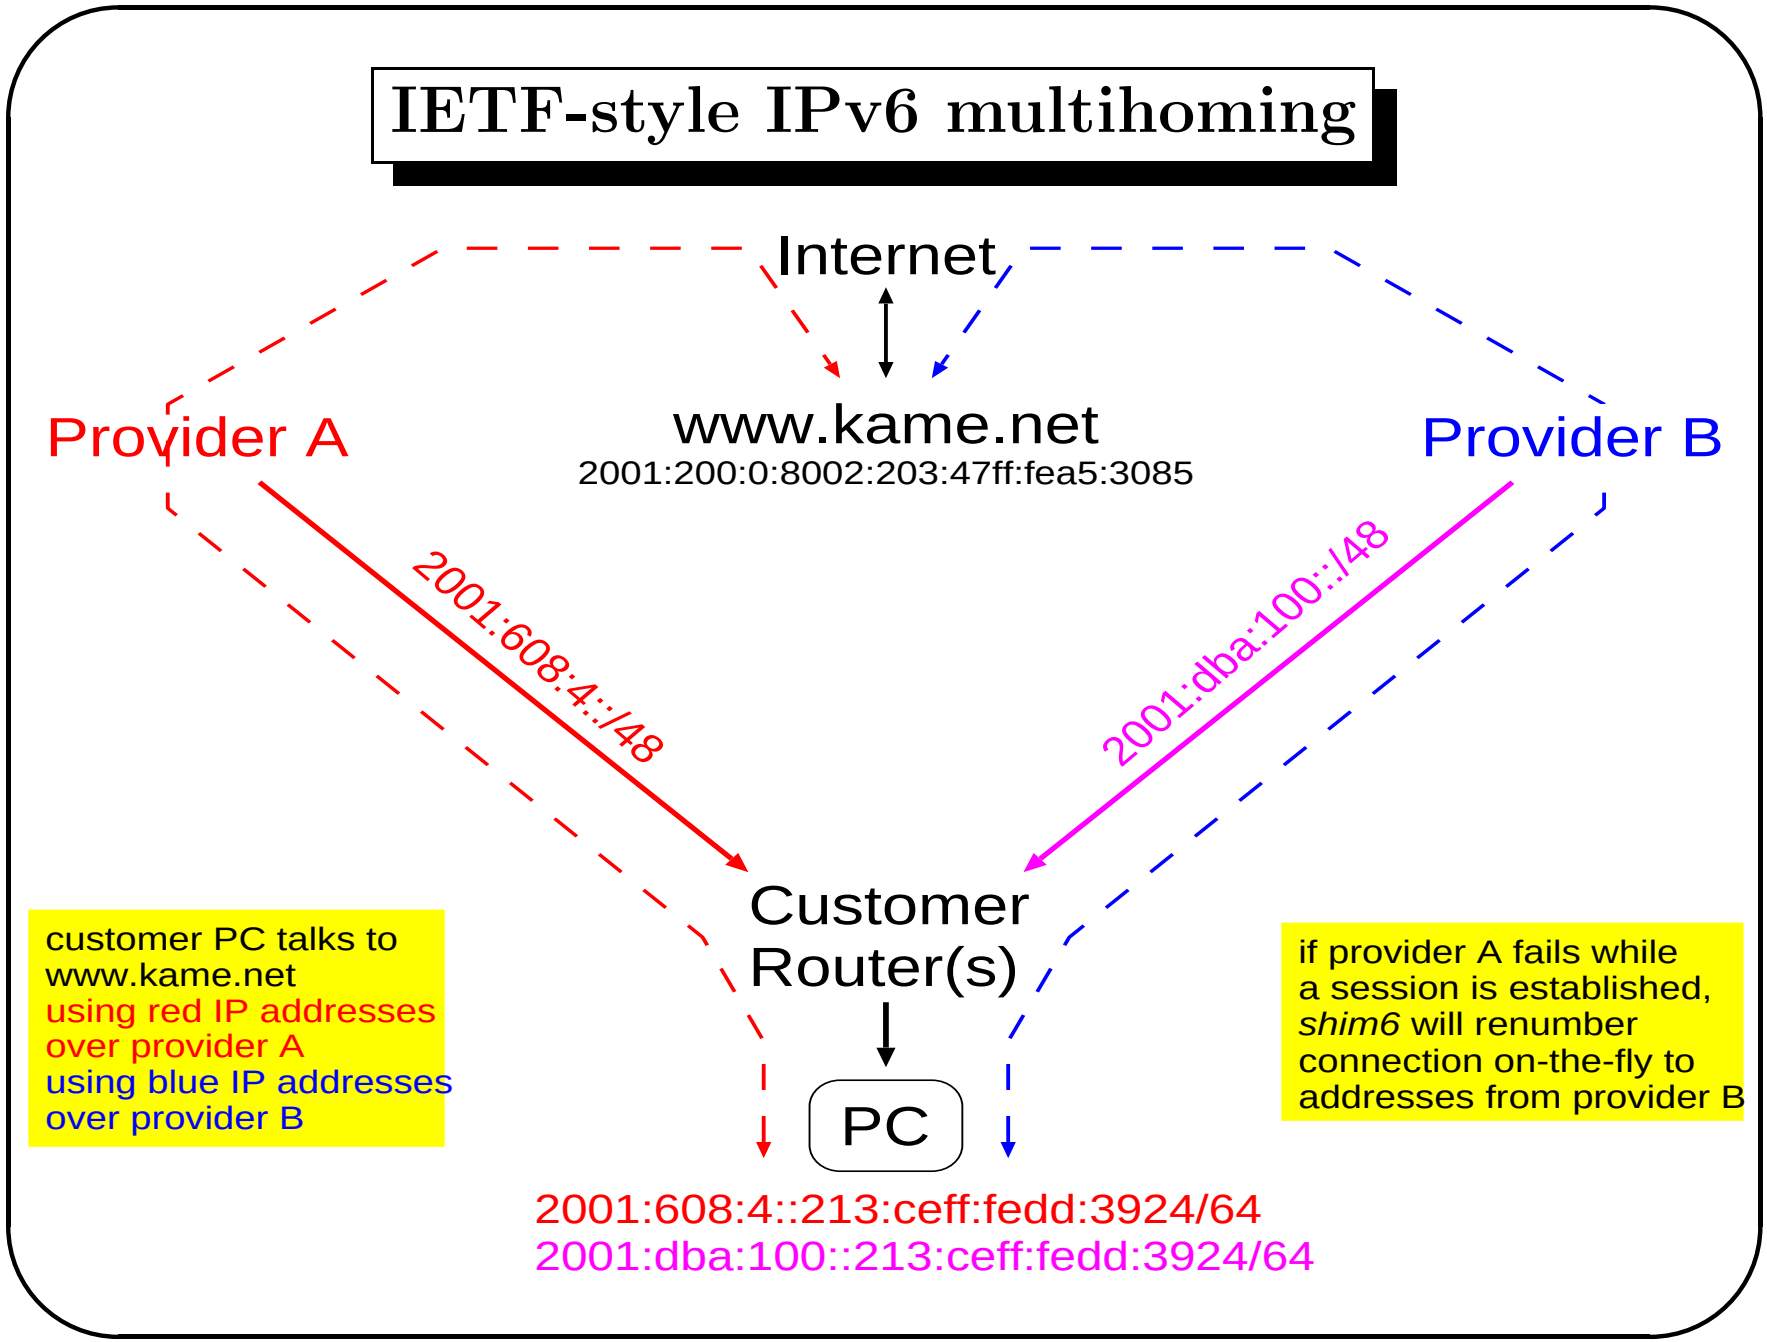

other IPv6 policies (for completeness)

- the normal ISP (PA) policy doesn't cover all special cases:
- IXP
  - Internet Exchange Points can get a /48
  - meant for the *exchange fabric*, not for services
- Root Name Servers
  - Root DNS operators can get a /48 for the root server
  - special, because the DNS root is *really* hard to renumber
- Anycast DNS
  - DNS server operators that run anycast DNS can get a /48
  - special, because a dedicated prefix is needed

IPv6 policies in other RIR regions

- provider independent (PI) / portable addresses
  - accepted as policy in ARIN, APNIC and AfriNIC regions
  - end sites can get portable /48 directly from RIR
  - non-ideal situation, imbalance between RIR regions
- critical infrastructure
  - some more "special cases" in ARIN and AfriNIC regions
  - widely varying definition of "critical"
- standard-size /48 for end sites
  - changed to /56-or-/48 in APNIC and ARIN region
  - under discussion in RIPE region

who or what is RIPE?

- RIPE NCC
  - non-for-profit company, located in Amsterdam
  - can be seen as a secretariat for European ISPs
  - make sure address-distribution is done in a *unique*, *impartial*, *fair* and *transparent* way
  - operation guided by RIPE policy documents
- RIPE *members* are ISPs paying for RIPE NCC services
- RIPE
  - is all of us (!)
  - the community of Internet users in the RIPE service region
  - community creates rules how the RIPE NCC works: *policies*

how can a RIPE policy be changed?

- the RIPE community is organized in *working groups*
- for addressing: Address Policy Working Group
- policy making follows an open model
  - everbody can contribute
  - discussions are on a public mailing list
  - decisions are based on community consensus
  - face to face discussions at RIPE meetings
    (but that's *no requirement* for influencing policy changes)
- working group chair does not decide (!), just steers process

Policy Development Process (1)

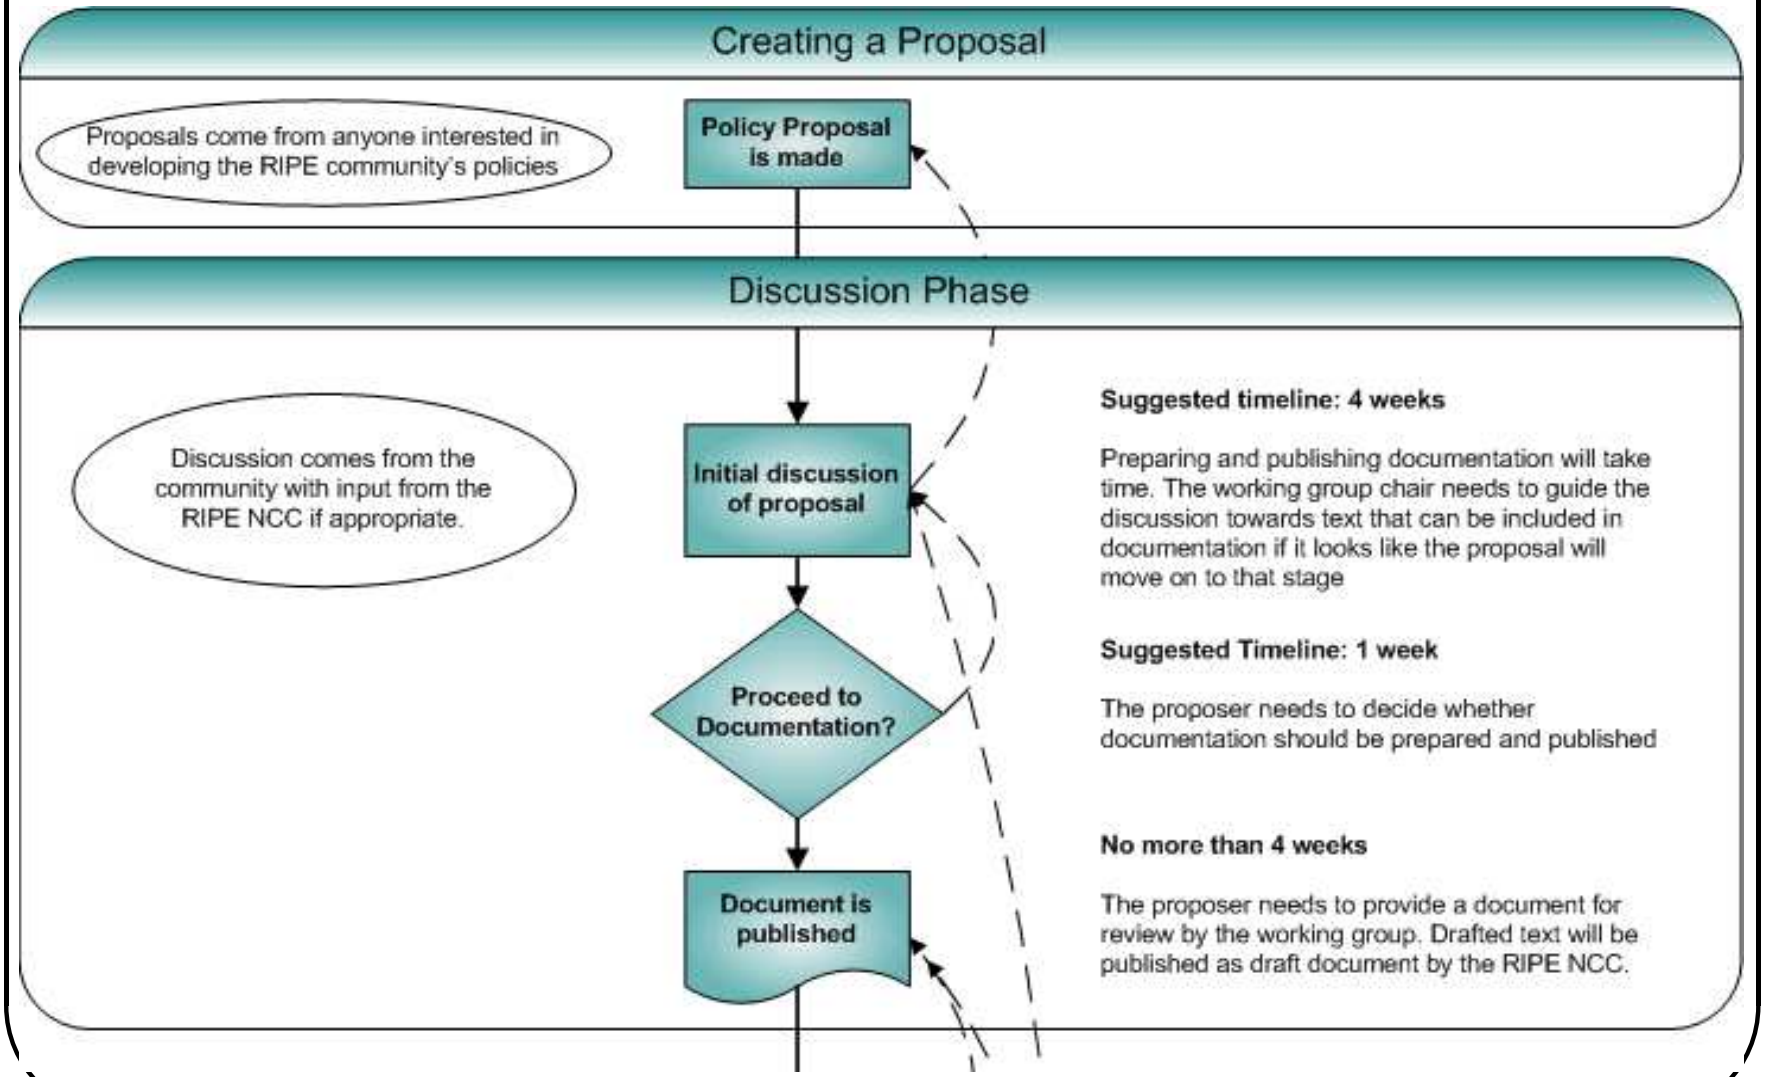

Policy Development Process (2)

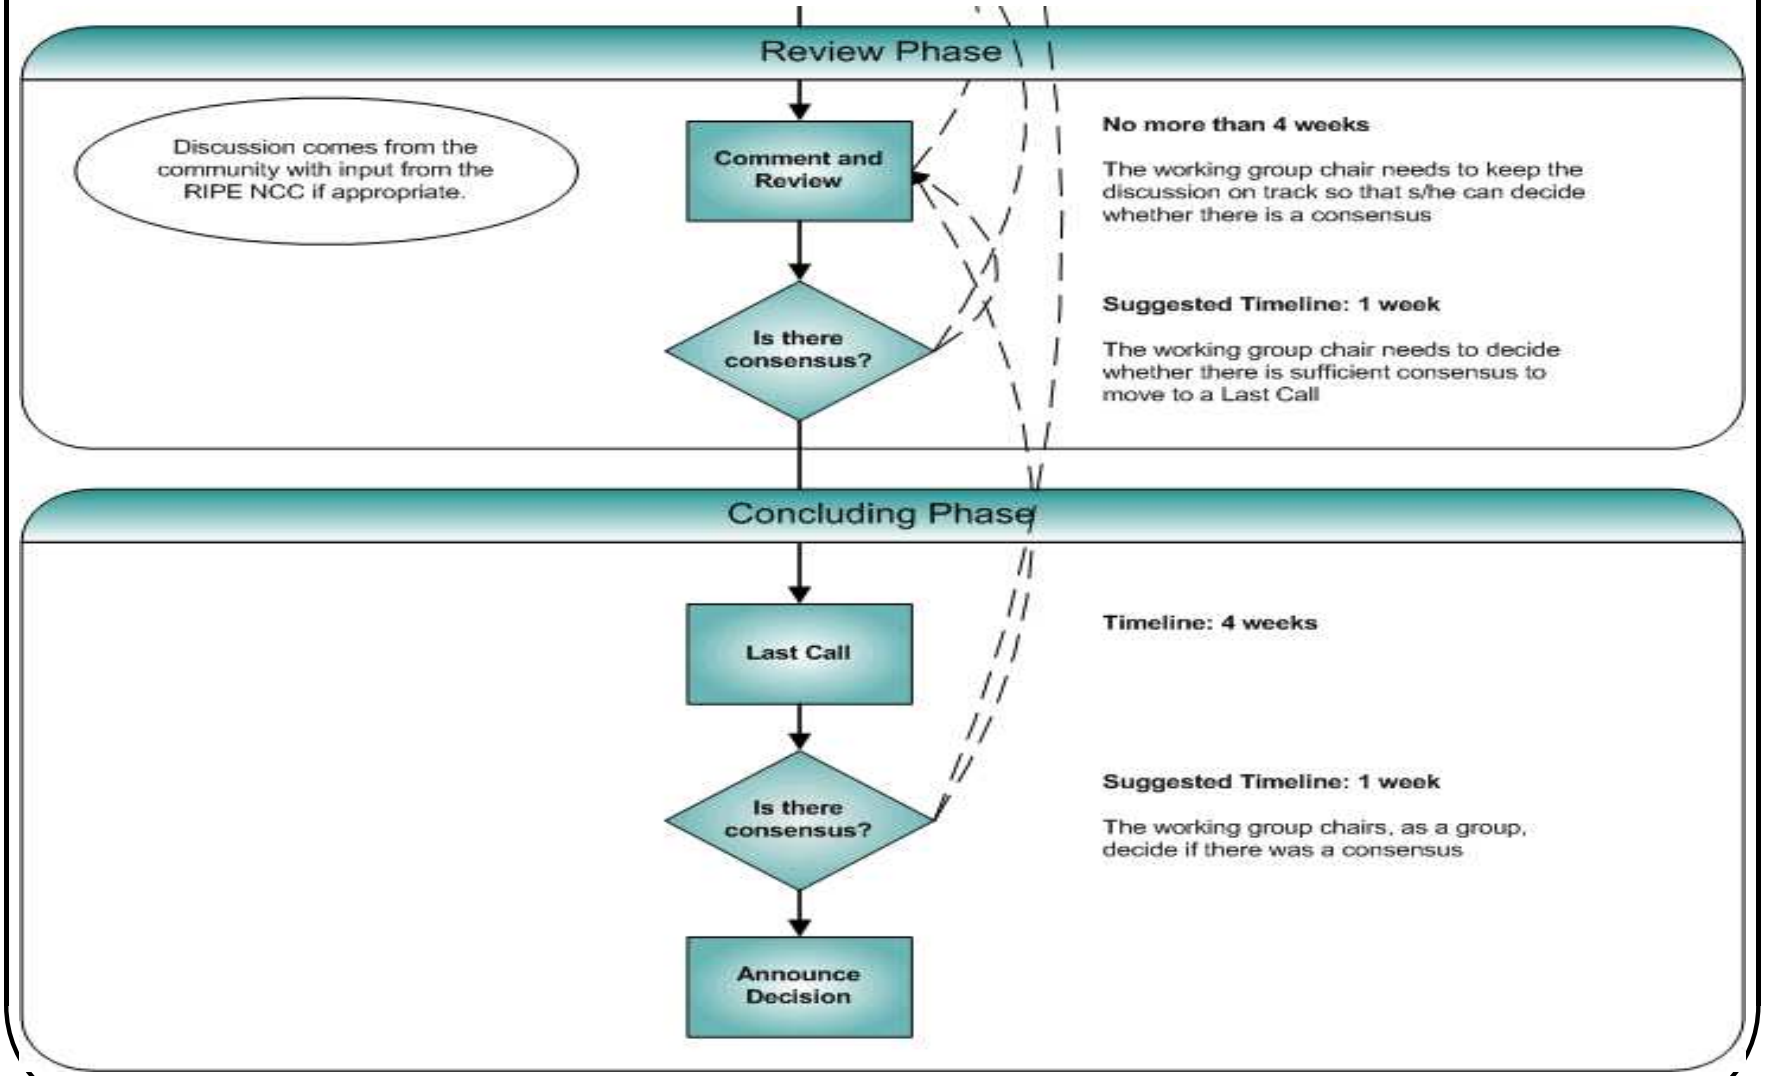

References & Questions?

- http://www.ripe.net/
- http://www.ripe.net/ripe/policies/
- APWG archive:
  http://www.ripe.net/ripe/wg/address-policy/
- IPv6-ops list:
  http://lists.cluenet.de/mailman/listinfo/ipv6-ops/

- questions and comments? ⇒ gert@space.net