PA/PI - Unification of IPv6 Address Space

Elvis Daniel Velea - V4Escrow, LLC

Daniel Stolpe - Resilians AB

Olaf Sonderegger - Abraxas Informatik

RIPE Meeting, 16 October 2013

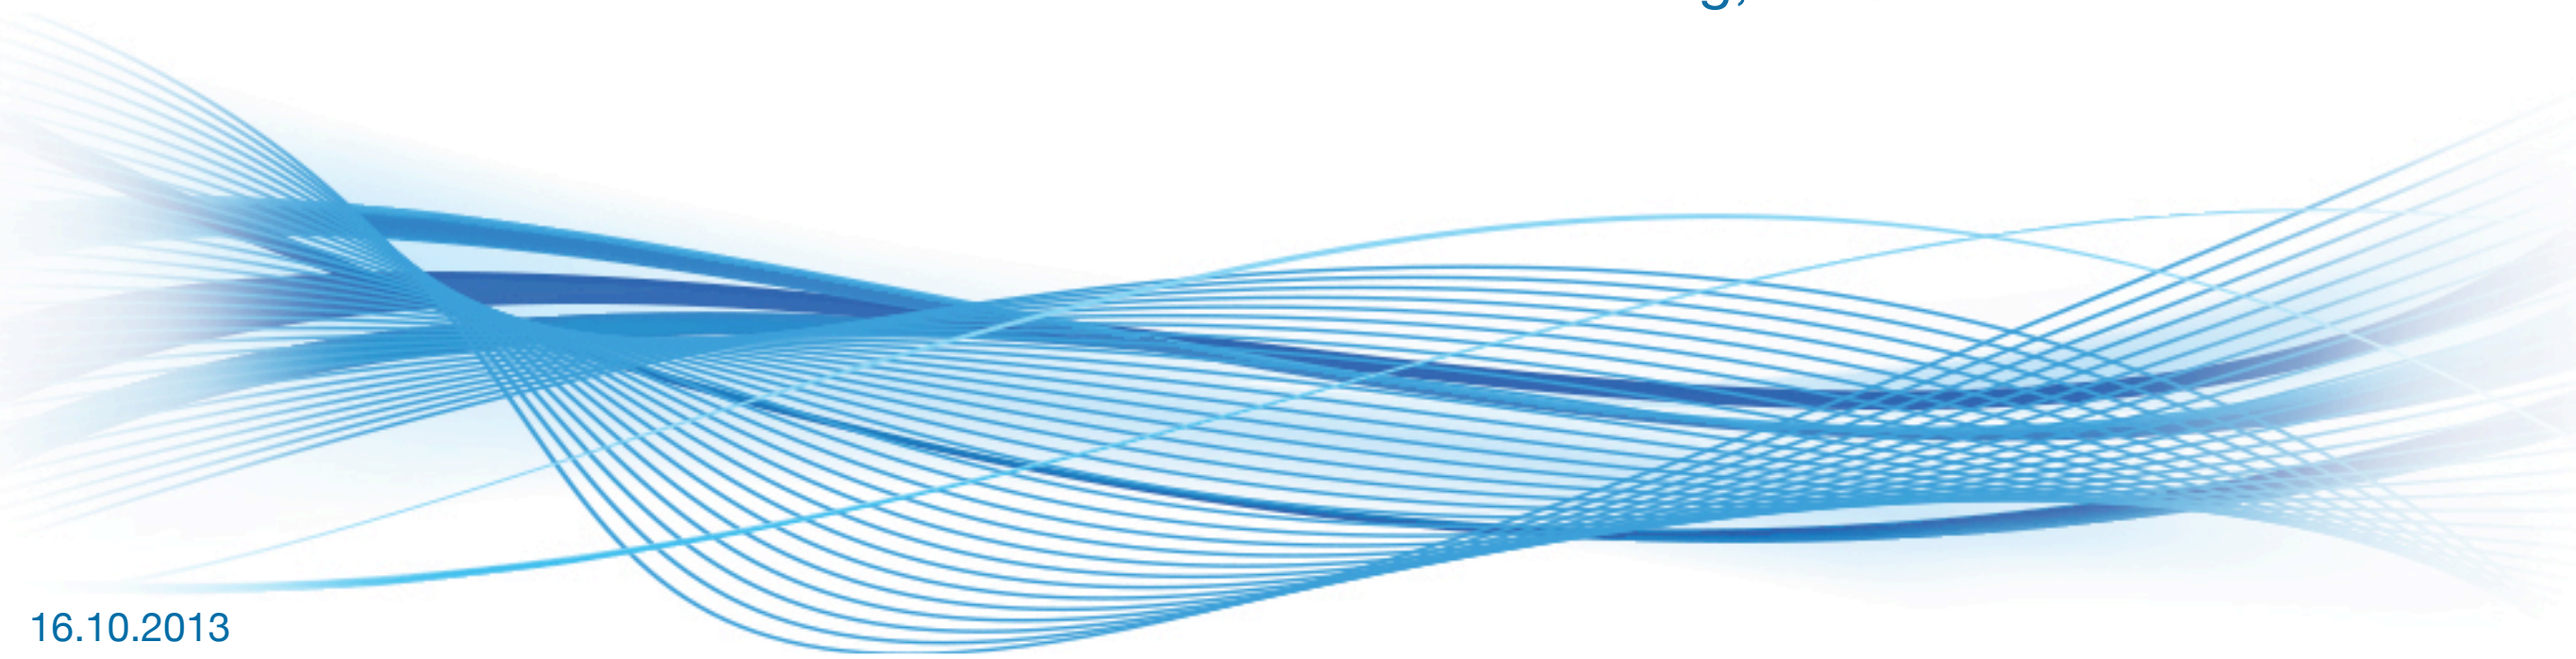

It all started with RIPE62

why is there a difference between PA and PI?

- in the end, it's "just some numbers" given out by the RIPE NCC to "consumers" of these numbers
- difference comes from intended use:
- PA
  - intended to aggregate (A) thousands or millions of end users into a single block, single routing table slot
  - assumed that "ISP" would be RIPE member anyway
  - liberal sizing, no strings attached
- PI
  - intended for a single independent (I) end-user network
  - not indended as "cheap replacement for RIPE membership"
  - specific purpose (BGP multhoming) ⇒ strings attached

It all started with RIPE62

address policy needs to balance...

- routing table
  - 1 million routes will break it for everbody
- NCC costs
  - we need the NCC to have a solid financial basis
- end user costs
  - too expensive RIR cost will lead to creative workarounds
- usefulness
  - address space acquired must be useful for the purpose
- address space efficiency
- good stewardship: encourage /48.../64 to end users

It all started with RIPE62

more radical approach

- abandon distinction between PA and PI completely
- RIPE members (LIRs) go to RIPE NCC and ask for "numbers"
- numbers are then used to "number things"
- difference between "ISP like" users and "end users" could be taken into account by checkbox
  - ( ) I want to assign /56s to end users ⇒ /32 allocated
  - otherwise default is /48
  - larger than /48 or /32 if documented need
- "sponsoring LIR" model or "become a member"
- AGM and NCC board to re-balance the costs to make size of IPv6 allocation not relevant for "become LIR or not?" decision

And continued at RIPE63

more radical approach

- abandon distinction between PA and PI completely
- RIPE members (LIRs) go to RIPE NCC and ask for “numbers”
- numbers are then used to “number things”
- sounds easy...?
- the tricky bit is to get the details right
- draft at http://www.ripe.net/ripe/mail/archives/address-policy-wg/2011-October/006496.html

And continued at RIPE63

1. WHO gets address space?

- keep LIR (= RIPE member) and "sponsoring LIR" model
- all number blocks go from the RIPE NCC to a LIR
- then
  - either LIR uses it "for its own network"
  - or LIR passes on to customer that has signed appropriate contracts (keeping the requirements of 2007-01)

And continued at RIPE63

2. HOW BIG should a single block be?

- /48 “by default”
- larger than /48 for “end-sites with large networks”, if justified
- /32 (.../29) when planning to assign /48. . . /64 to 3rd parties
- larger than that: when documented need (as now)
- automatic consequence: “multiple blocks of numbers” to a single LIR will have to be accepted as “frequently seen usage case” (we’ll come back to this)

And continued at RIPE63

3. SPECIAL CASE networks

- currently: special case PI exists for IXP, Root DNS, Anycast DNS
- proposed implementation:
  - checkbox on the template “I want to use that for IXP/Root/Anycast DNS”
  - not used for evaluation (this is just numbers!) but used for selection of the address range to pick numbers from
  - people want to have the option to treat “special numbers” differently in their routing policy, and that’s easier if they are easily recognized
- also proposed to have a well-documented range for /33. . . /48 number blocks (smaller “minimum assignment size” range)

And continued at RIPE63

4. COSTS

- we don't decide costs
- but we can send recommendations to the AGM (and then vote)
- one possible model could be:
  - yearly base fee for LIR
  - per-piece yearly fee for each number block held
  - /48 = 50 EUR, /32 = 100 EUR, /31+ = 200 EUR /year?
  - (per-piece installation fee for each new number block?)

And continued at RIPE63

5. MULTIPLE BLOCKS per LIR

- "more than one block for a single LIR" is not possible today for PA (unless full), but would be needed for new "number blocks"
- "get any number of blocks you ask for"
  - not likely to get consensus
- proposal for compromise:
  one "block of numbers" per "network"
- definition of "network"?
  - interconnected nodes
  - operated by same entity
  - operated with a common routing policy
  - operated as a *layer 3* network
- goal is to be *reasonably* flexible here

And then RIPE64

L. On IPv6 PI/PA Unification

- long-term project to unify different IPv6 PA and PI policies into “just IPv6 addresses with no colours”
- presented at RIPE 62 and RIPE 63, and on the mailing list
- encouraging feedback from audience

- since RIPE 63:
- work on “formal policy document” draft has been done
- but not ready for wider publication yet
- also delayed to wait for 2011-02 and 2011-04 conclusion

- the plan: get policy into PDP well before RIPE 65

And RIPE65

H. update on ongoing policy projects

- IPv6 PA/PI unification policy (Gert Döring)
  - unify IPv6 PA and IPv6 PI into "address blocks"
  - get rid of unexpected restrictions on "addresses"
  - simplification of things can get surprisingly complicated
  - got stuck in IPv4 run out hectic, expect document "soon"

And, lastly, RIPE66

H. IPv6 PA/PI unification policy

- unify IPv6 PA and IPv6 PI into "address blocks"
- get rid of unexpected restrictions on "addresses"
- simplification of things can get surprisingly complicated
- got stuck in *(insert excuses here)*
- new approach: find volunteers to help

Now at RIPE67

- Volunteers found
  - **Elvis Daniel Velea** from V4Escrow,LLC
  - **Daniel Stolpe** from Resilians AB
  - **Olaf Sonderegger** from Abraxas Informatik
- Worked over several revisions between July 4th and September 19th
- Unfortunately, Olaf could not be here with us but he is watching the webcast :)

Major problems with current IPv6 policies

- the ISP is associated with the LIR/PA definition, the PI is used by non-ISPs
  - **no longer a clear distinction**

- IPv6 PI can only be used to number an internal network and not to offer services, not even to your cousin's server

- various restrictions on PI are slowing down IPv6 adoption

Minor problems with current IPv6 policies

- we have three documents for IPv6 policies:
  - ripe-589 IPv6 Address Allocation and Assignment Policy
  - ripe-451 IPv6 Address Space Policy For Internet Exchange Points
  - ripe-233 IPv6 Addresses for Internet Root Servers In The RIPE Region

- policy prevents assignments(registration) smaller than /64

- LIRs could not request multiple allocations when they had multiple disconnected networks

How did we try to fix these problems?

- Removed differences between PI and PA
- One single policy document + included special cases
- Included the definition of the Sponsoring LIR in Policy
- Removed the ASSIGNMENT
- Introduced the SUB-ALLOCATION
  - document real world structure of IP hierarchy
- Allowed additional allocation for routing purposes (for those organisations that have disconnected networks)

How did we try to fix these problems? (2)

- You want to make large sub-allocations
  - you can request a /32 or larger
- You do not want to make large sub-allocations
  - you can get a /48
- Applies to both members and non-members

- Current proposal does not clearly show this and after discussion between authors, it will be changed in v2

How did we try to fix these problems? (3)

- Need more than /48 for an end site -> request approval from RIPE NCC
- Need to sub-allocate more than a /40 -> request approval from RIPE NCC
- Current proposal says something else, after discussion between authors, it will be changed in v2

How did we try to fix these problems? (4)

- HD-Ratio - if anything within a /56 is registered, the whole /56 is considered to be in use.
- If there is nothing registered within a /48 sub-allocation, the whole /48 is considered used
- only /48 (or smaller) sub-allocations will be considered to be in use and count in the HD-Ratio calculation

- Current document is not clear and after discussion between authors, it will also be changed in v2

Questions raised on the mailing list

- **1**. Why limit ANYCAST, ENUM, IXP to a /48 if anyone else can get a /32?
  - we did not want to change too much so we kept these limitations as they were in the previous policy text
  - should we remove these limitations and allow ANYCAST, ENUM and IXP operators to request/receive a /32?
- **2**. Limit the Non-LIRs to a maximum /40? Or maximum /32?
  - isn't this change all about removing artificial limits?
  - what would be the right limit?

Questions raised on the mailing list

- **3.** Remove all the definitions from the policy text
  - and just create a procedural document
  - or create a document defining all special terms
  - Richard Hartman said that “There’s already some initial work going on off-list”.
  - Once that document is approved, all policies can be updated again.

Questions raised on the mailing list

- **4.** Remove the PI and Sponsoring LIR concepts?
  - every address holder must become an LIR
  - if you have a problem with your LIR, you must renumber
  - the NCC may need to hire an army :-)
  - Sponsoring LIR concept solves more problems than it causes

Questions raised on the mailing list

- **5.** All bits to the left of /64 should be in scope of the policy
  - should we care what happens within each /64?
  - what if someone wants to use /112s for peering?
  - what if someone decides to use /128s to number customers?
  - should these organisations register the /112s and /128s?
  - we will change/update the sentence removed from 1.1 by mistake:

"There is no limit on how grained the registration in the RIPE Database can be"

- will be added in v2 of the policy proposal

Questions raised on the mailing list

- **6.** End Users (*)

An entity that uses IP address space for its network only and does not provide IP/ASN services to customers is called an End User. Strictly speaking, End Users are not part of the Internet Registry System. They do, however, play an important role with respect to the goals defined above.

In order to achieve the conservation goal, for example, End Users should plan their networks to use a minimum amount of address space. They must document their addressing and deployment plans to the LIR and furnish any additional information required by the LIR for making assignment decisions.

To achieve the aggregation goal, an End User should choose an appropriate LIR. End Users should be aware that changing ISPs may require replacing addresses in their networks.

End Users must provide and update registration data for the address space assigned to them in the RIPE Database.

(*) **http://www.ripe.net/ripe/internet-coordination/internet-governance/internet-technical-community/the-rir-system**

Questions raised on the mailing list

6. (continued) End Users

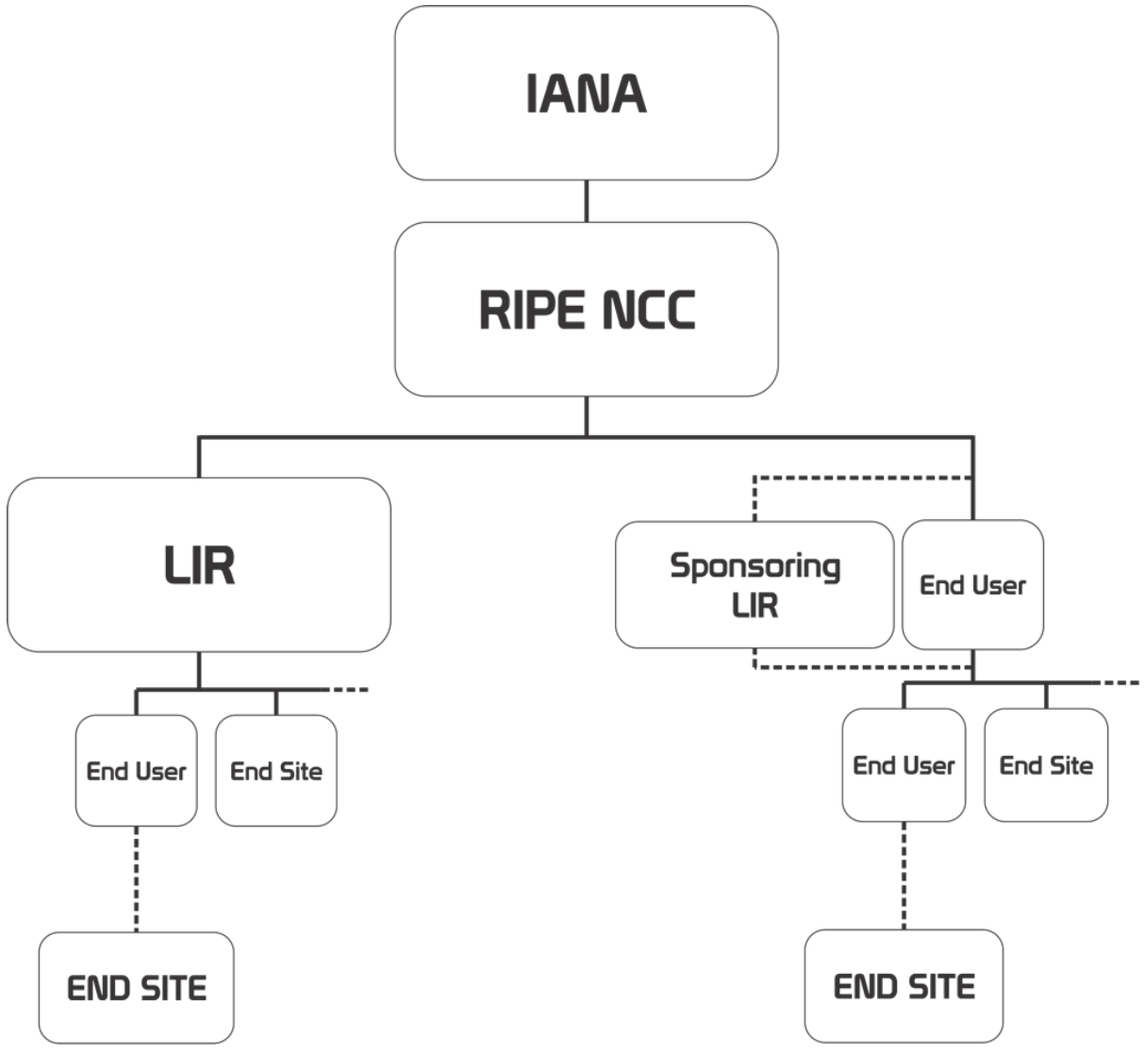

Questions raised on the mailing list

6. (continued) End Users (left side of the diagram)

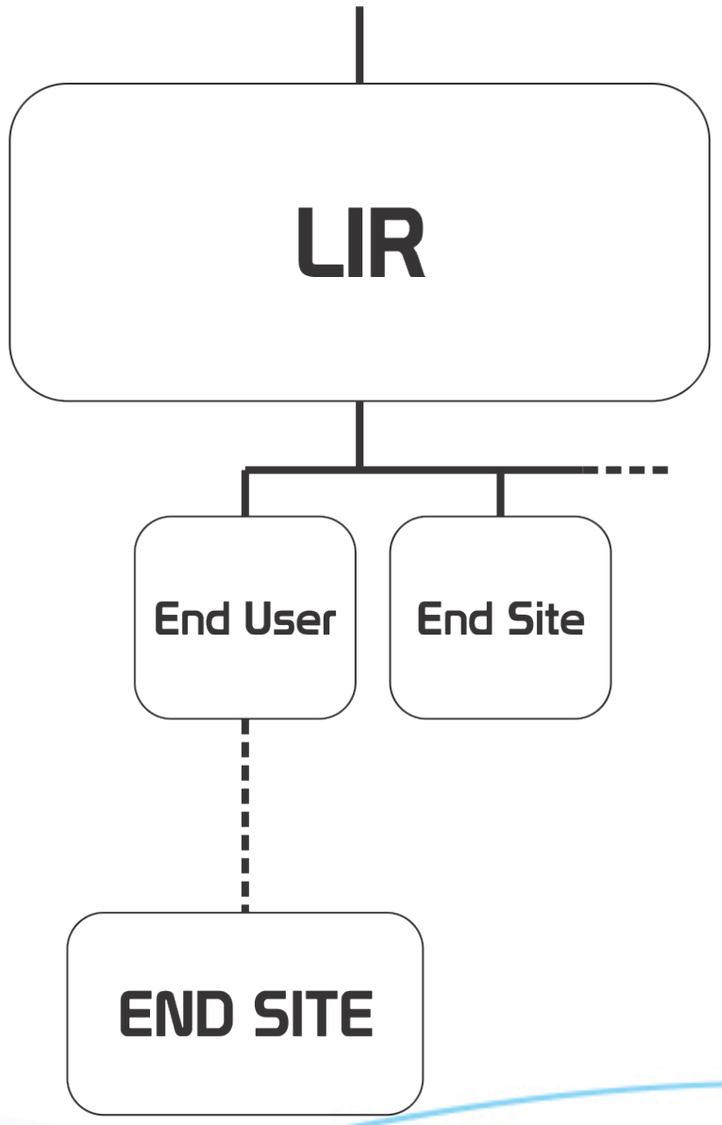

Questions raised on the mailing list

6. (continued) End Users (left side of the diagram)

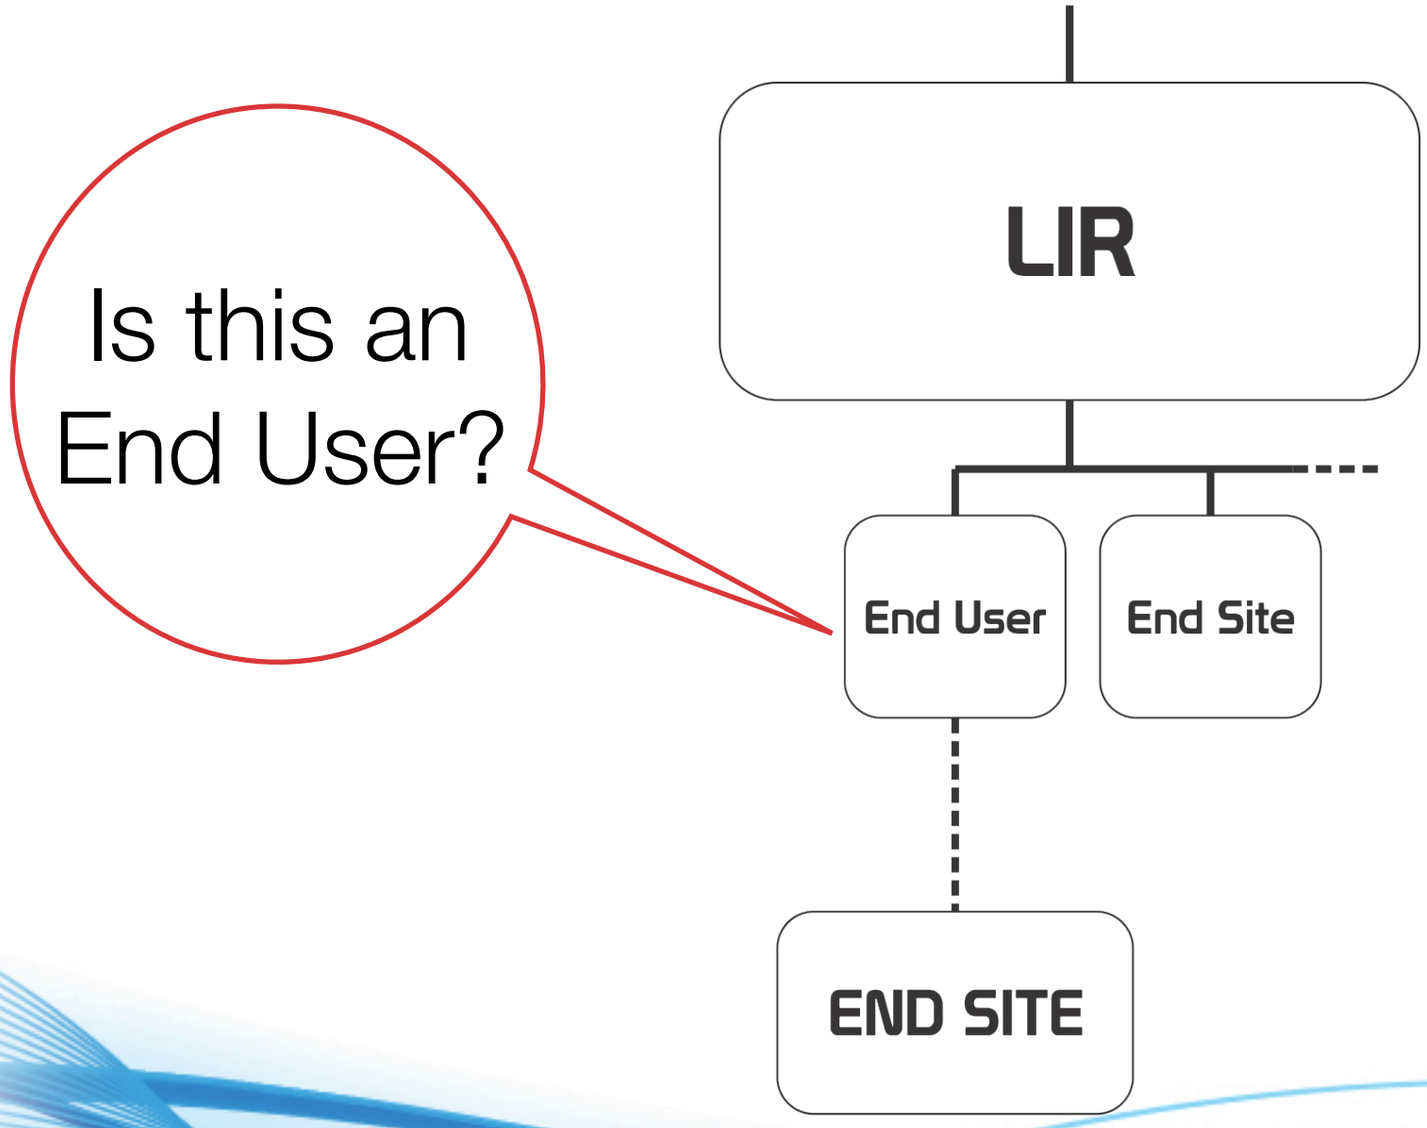

Questions raised on the mailing list

6. (continued) End Users (left side of the diagram)

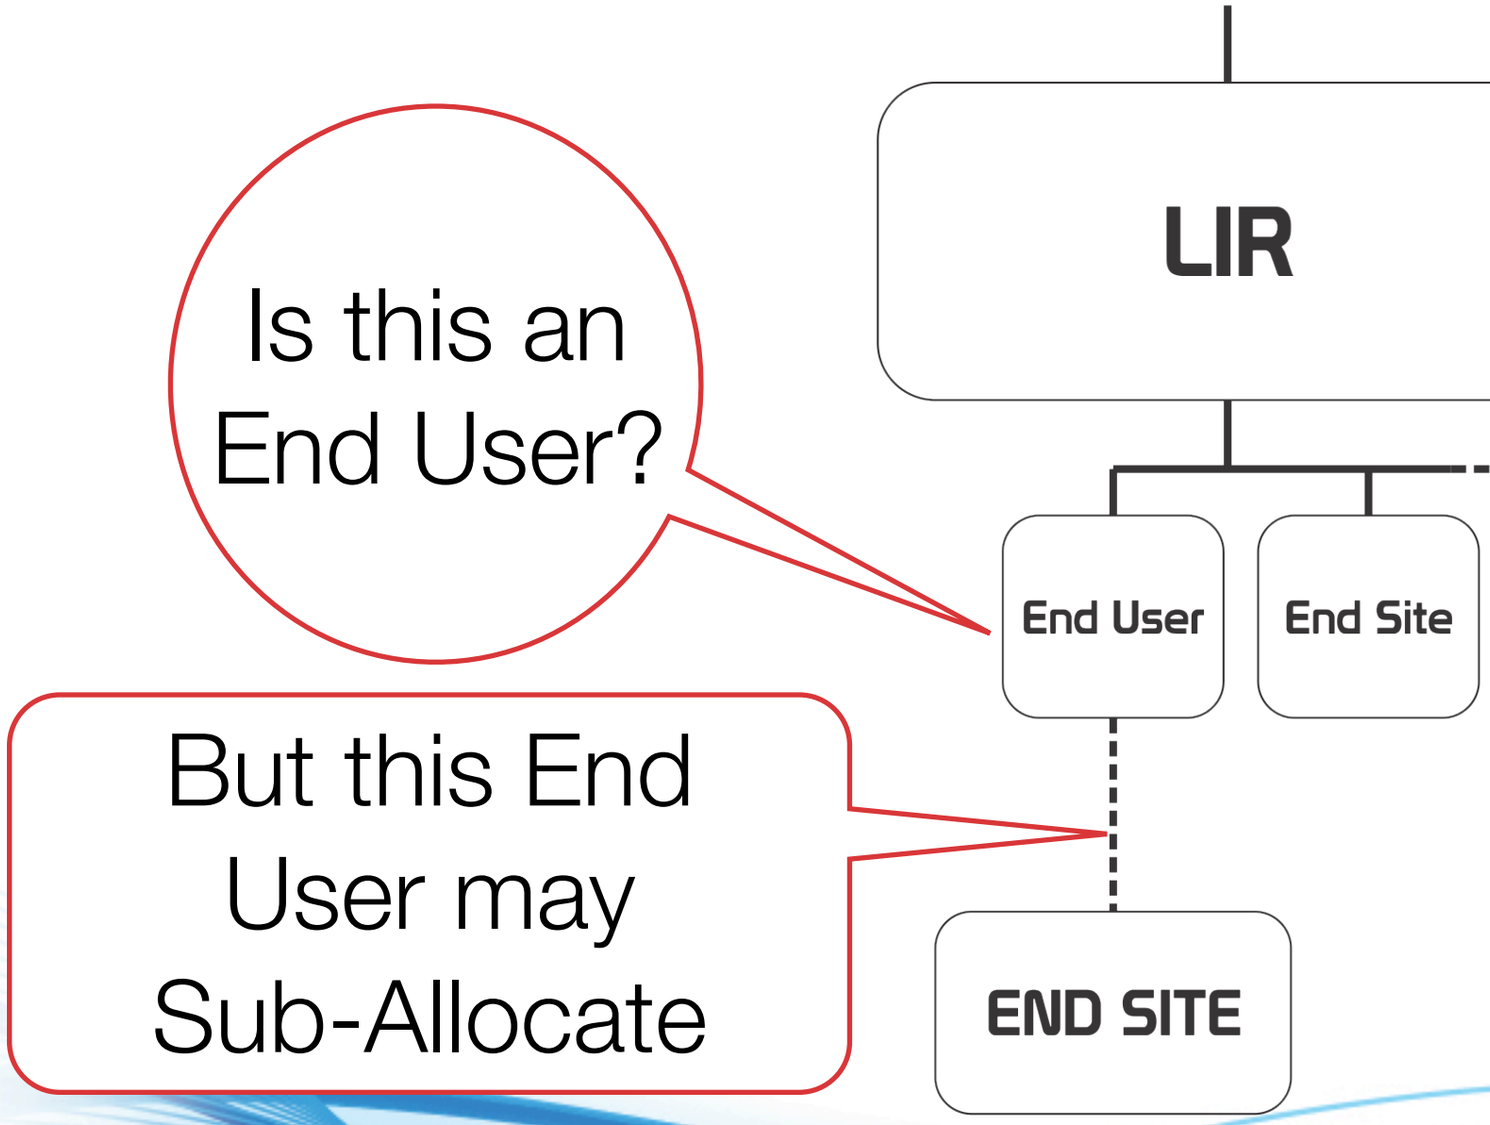

IF YOU'RE NOT

CONFUSED

YOU'RE NOT PAYING ATTENTION

Questions raised on the mailing list

6. (continued) End Users (left side of the diagram)

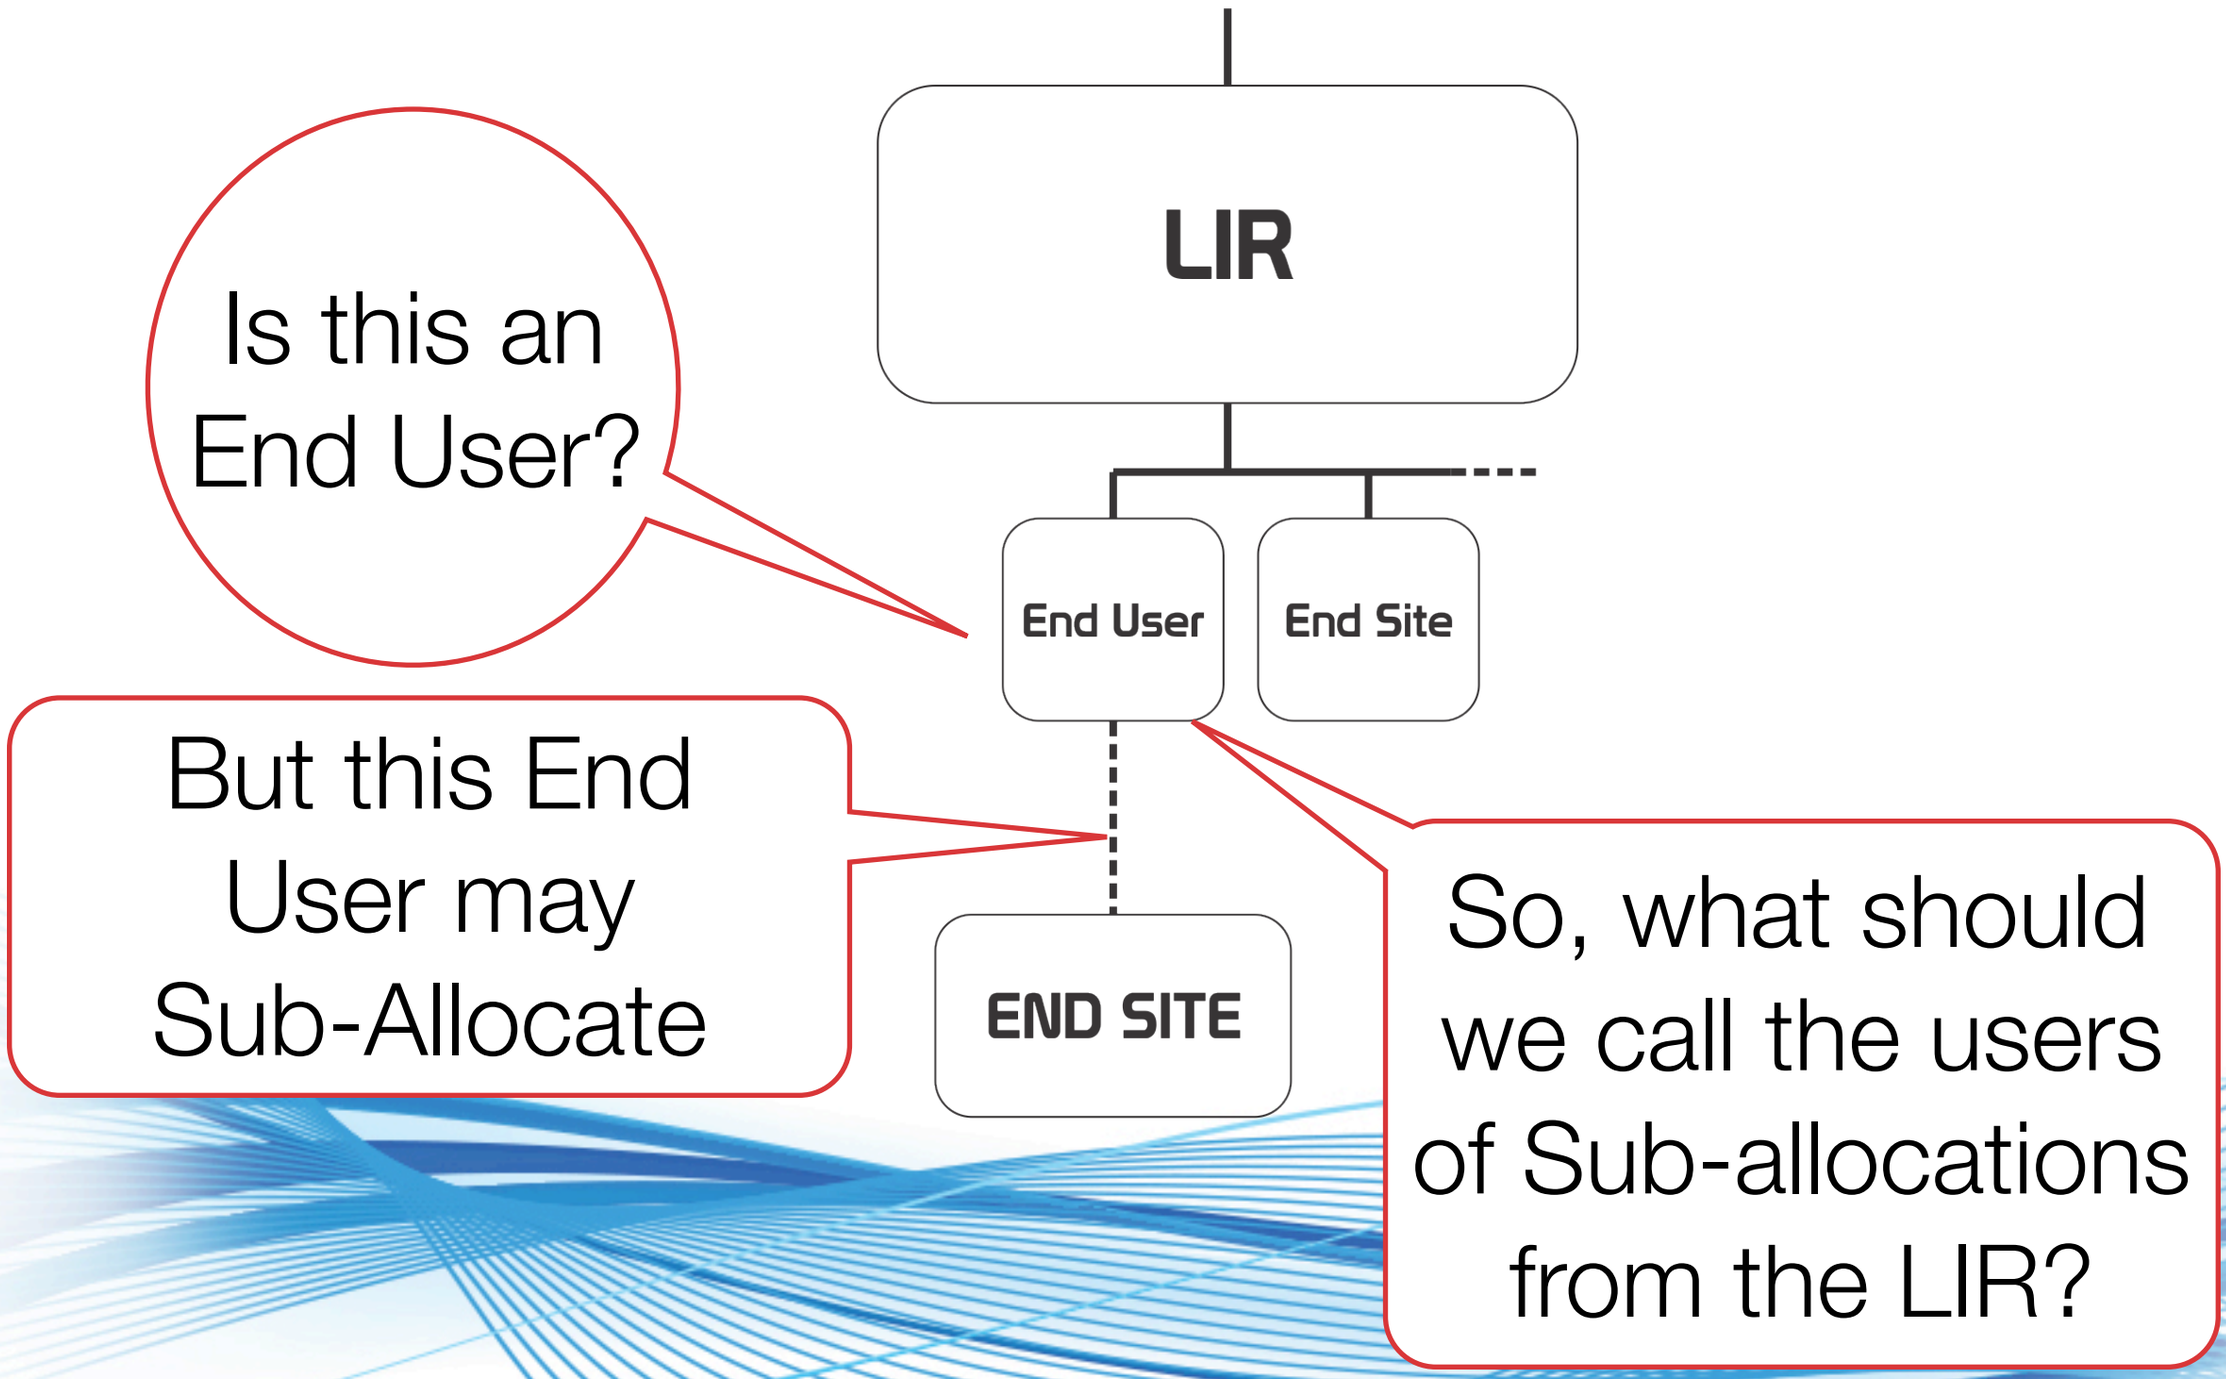

We could always use the Greek alphabet :)

Α Alpha	Β Beta	Γ Gamma	Δ Delta	Ε Epsilon	Ζ Zeta
Η Eta	Θ Theta	Ι Iota	Κ Kappa	Λ Lambda	Μ Mu
Ν Nu	Ξ Xi	Ο Omicron	Π Pi	Ρ Rho	Σ Sigma
Τ Tau	Υ Upsilon	Φ Phi	Χ Chi	Ψ Psi	Ω Omega

α Alpha	β Beta	γ Gamma	δ Delta	ε Epsilon	ζ Zeta
η Eta	θ Theta	ι Iota	κ Kappa	λ Lambda	μ Mu
ν Nu	ξ Xi	ο Omicron	π Pi	ρ Rho	σ Sigma
τ Tau	υ Upsilon	ϕ Phi	χ Chi	ψ Psi	ω Omega

Questions raised on the mailing list

6. (continued) End Users (right)

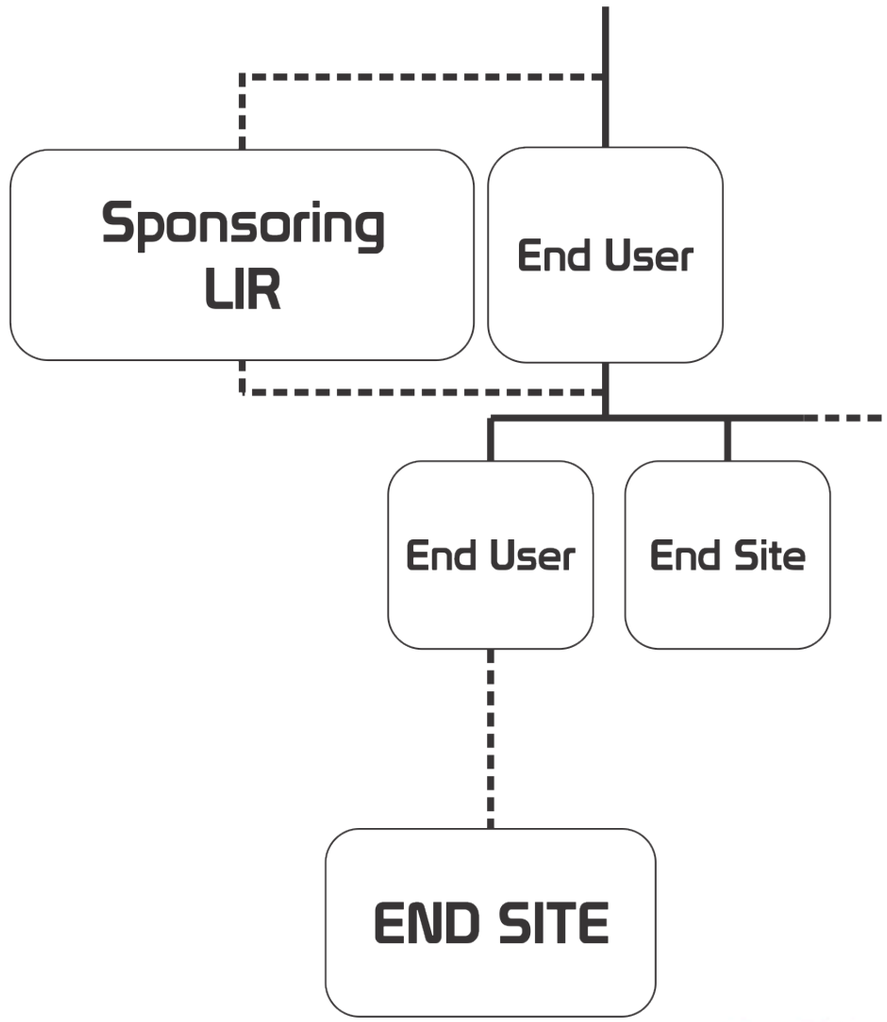

Allocation via the Sponsoring LIR

- This one is complicated
  - the entity receives the allocation from the RIPE NCC
    - can make sub-allocations
    - can use it for their infrastructure or for their customers
    - must register every sub-allocation in the RIPE Database
    - must keep a record of what was sub-allocated to whom
- It acts like an LIR but it's not an LIR
    - what can we call it?

Questions raised on the mailing list

6. (continued) End Users (right)

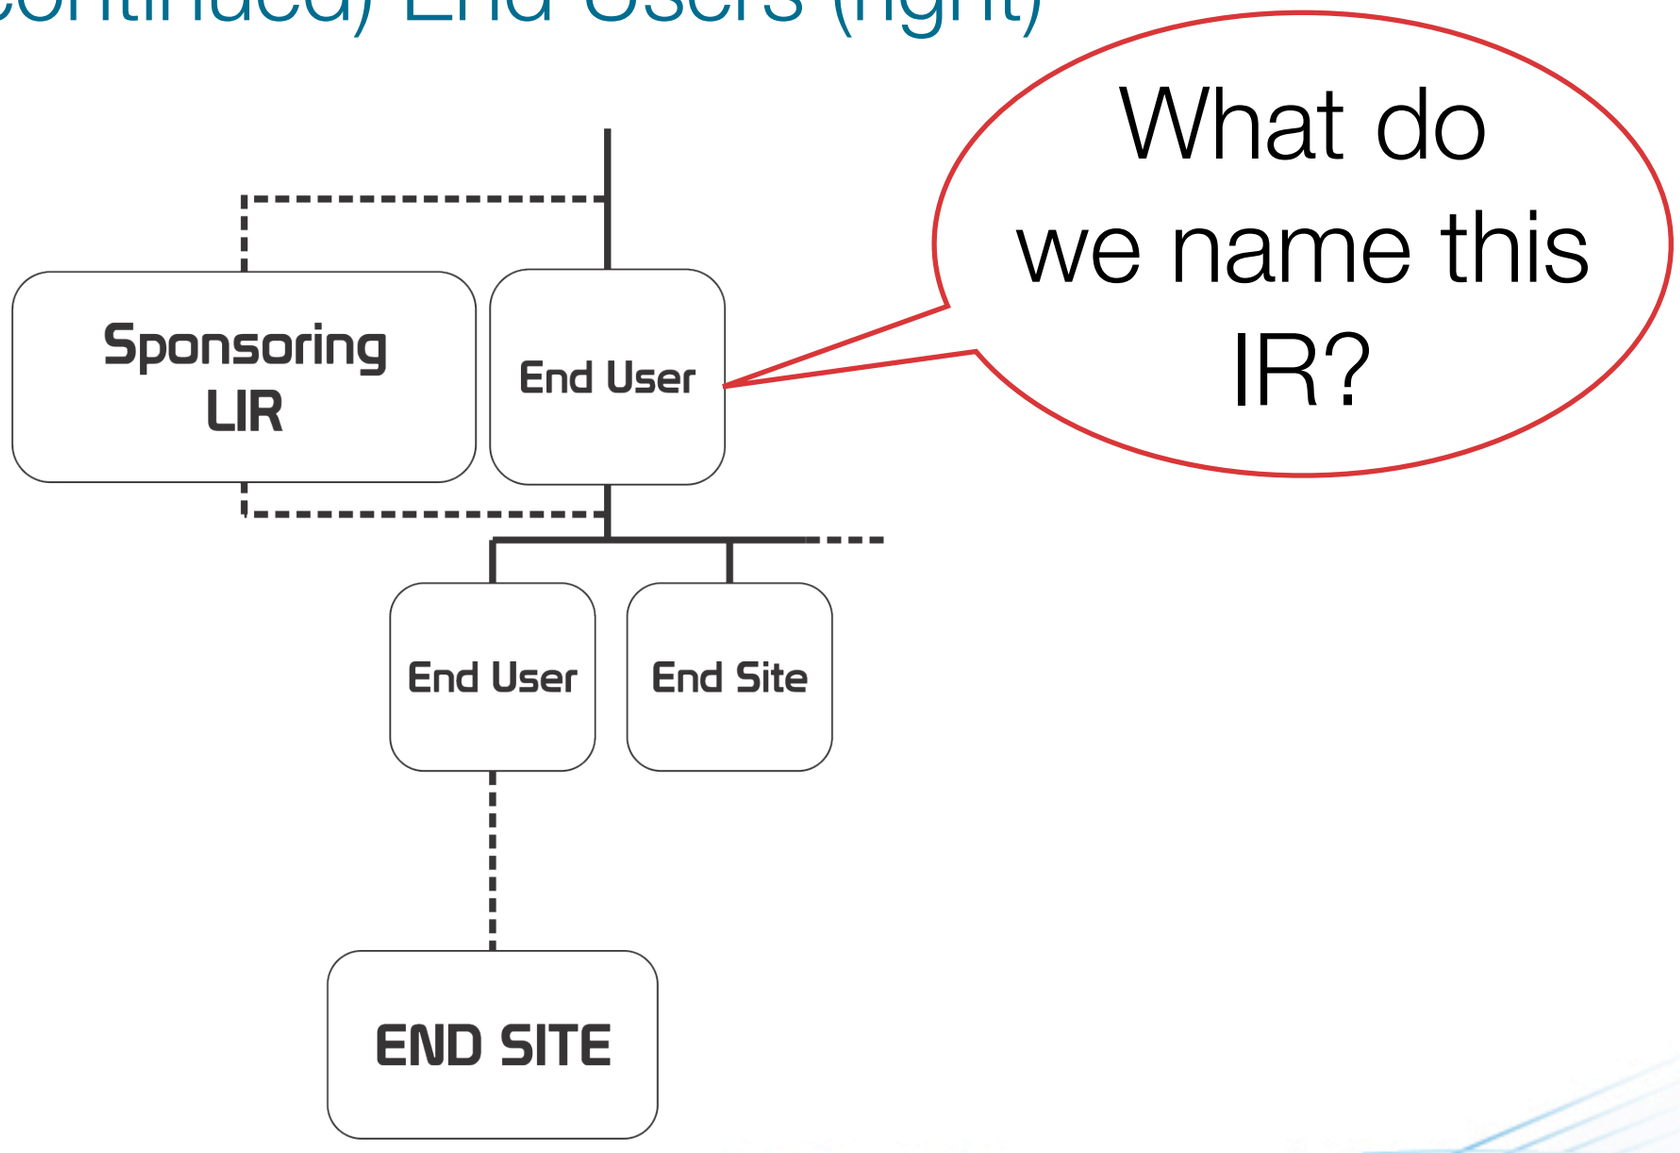

Questions raised on the mailing list

6. (continued) End Users (right)

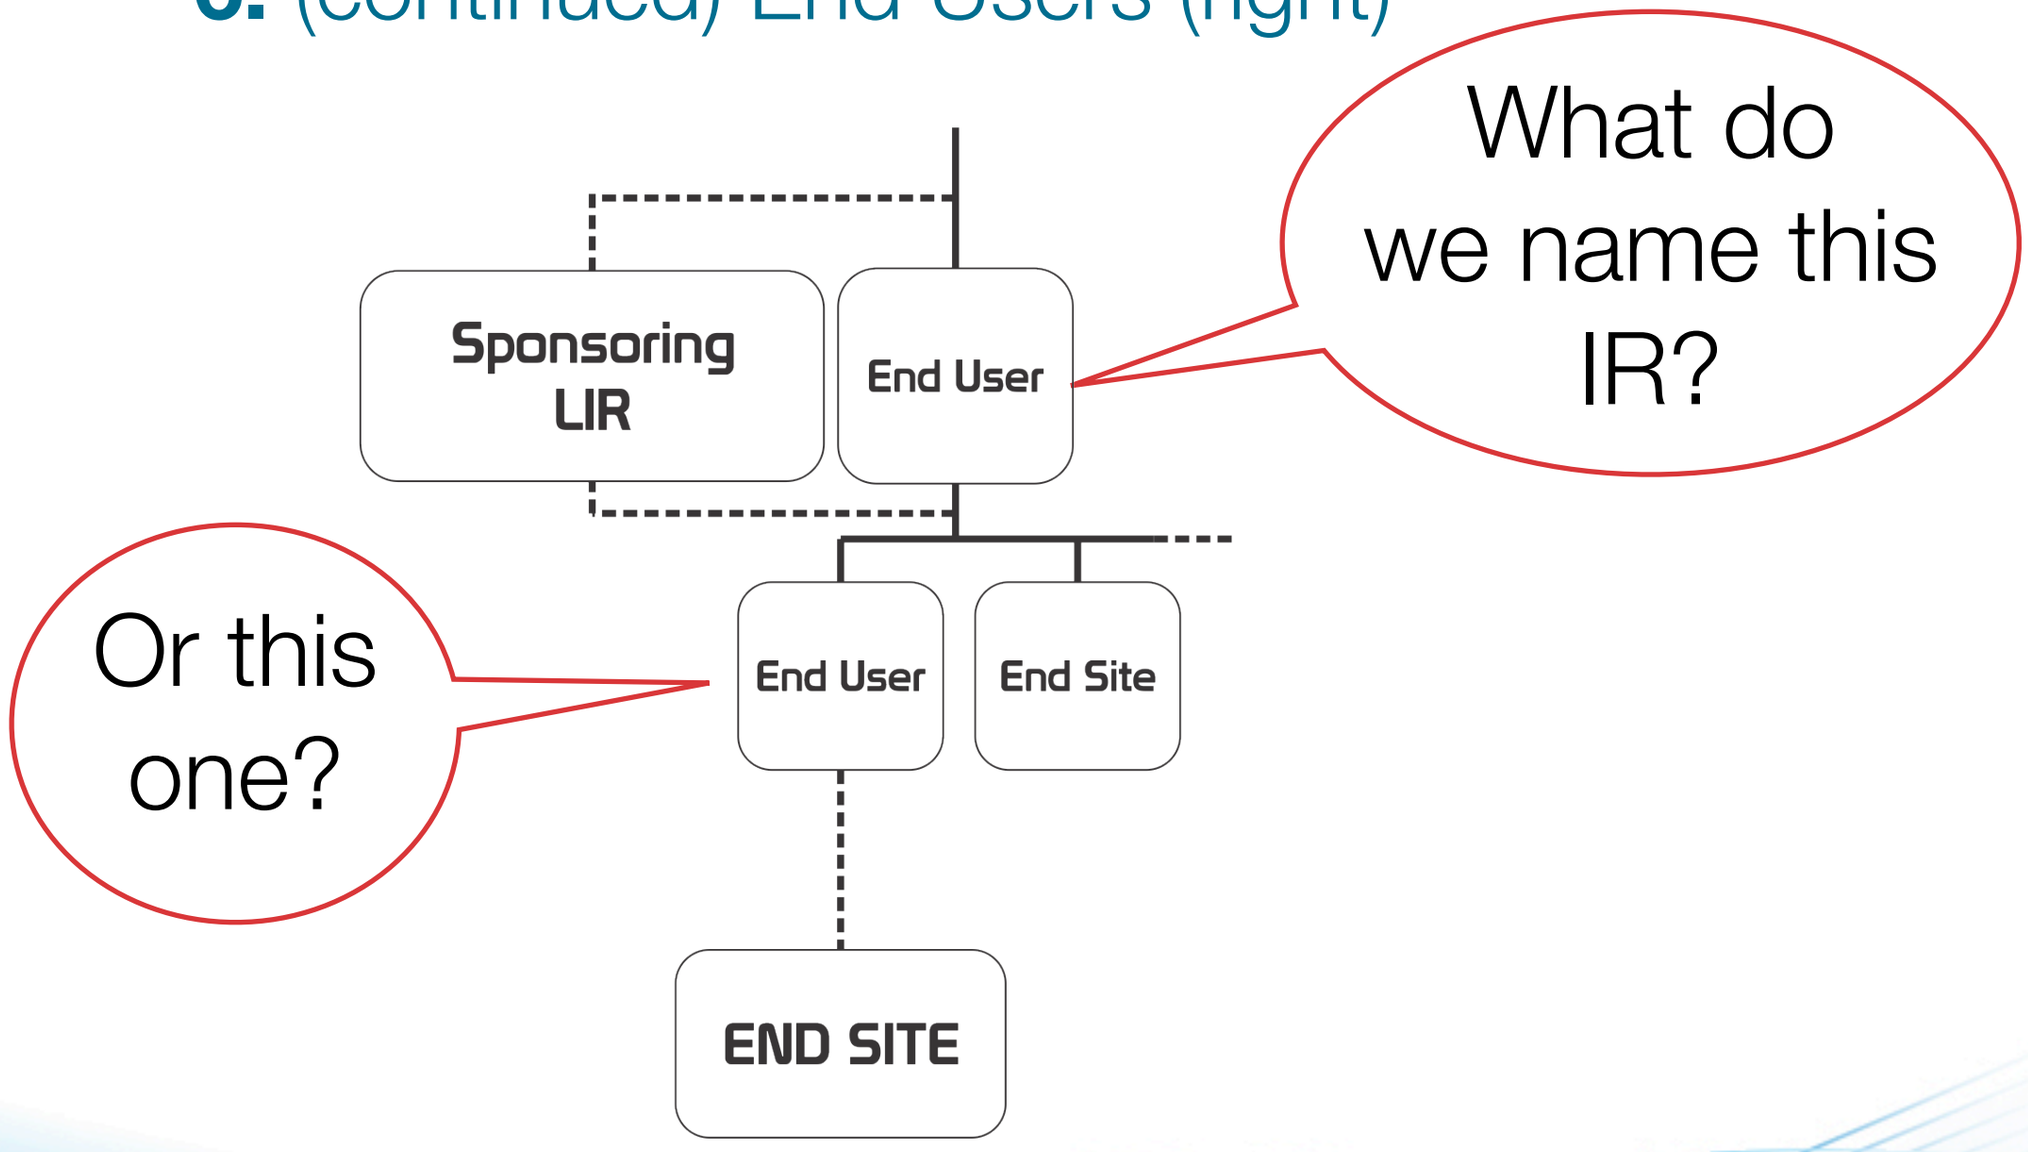

Questions raised on the mailing list

- **6.** (continued) End Users
  - we had the idea of PIR (Portable Internet Registry)
  - someone came with the idea of SIR (Sponsored Internet Registry)
  - other ideas were: Sub-LIR, Child LIR, associate member,
  - what about Sub-IR?
  - what do you think?

Questions raised on the mailing list - Billing

- How do you convince current PI users to pay more?
  - you don’t; we could keep the /48 at 50€
  - everyone pays the same
  - some other options

Questions raised on the mailing list - Billing

- Everyone pays the same
  - more than 20.000 organisations using independent resources
  - almost 10.000 LIRs
  - about 1.600 overlap
  - the payment is made per independent resource and not per organisation
  - the payment will be made per organisation and not per resource

Questions raised on the mailing list - Billing

- Everyone pays the same
  - Current users of independent resources will be paying at least 100€ once they start using IPv6 next to IPv4
  - the increase is from 100€ (if you use one v4 + one v6 independent resource) to 600-700€:
    - no matter how many resources the organisation will use
    - including all services, if you sign a membership contract
    - but only if it also applies to all current users of independent resources

Billing suggestions - Do you have any other?

- Suggestion1: **100€**/Non-LIR; **1750€**/LIR (2014 budget)
- Suggestion2: everyone pays the same: ~ **600-700€**
- Suggestion3: 80% of the budget paid by members (**1400€**) and 20% of the budget paid by Non-members (**200€**)
  - these options imply payment per organisation
- Suggestion4: **50€** per /48; **100-200€** per /32, around **1500€** per membership
  - this option implies payment per resource

(Un)foreseen consequences - 20 RIPE Documents affected

- **ripe-592** (IPv4 Policy) should be updated to reference this policy when approved
- **ripe-589** (IPv6 policy) will be updated
- **ripe-585, ripe-586, ripe-587** (Temporary assignments policy) need to be updated
- **ripe-573, ripe-574** (IPv6 PI request form and supporting notes) will be obsolete
- **ripe-575, ripe-576** (IXP request form and supporting notes) need to be updated
- **ripe-567, ripe-568** (Anycast request form and supporting notes) need to be updated
- **ripe-560, ripe-561, ripe-425, ripe-422** (IPv6 allocation requests) need to be updated
- **ripe-513** (value of status and assignment-size) will be obsolete
- **ripe-452** (Contractual requirements) should be updated to include IPv6 allocations
- **ripe-373, ripe-374** (End User Assignment request) need to be updated
- **ripe-233** (IPv6 for Internet Root Servers) will be obsolete

Questions ?

elvis@v4escrow.net
stolpe@resilans.se
Olaf.Sonderegger@abraxas.ch