



Question(s): 3/20

Wuxi, 3-13 December 2018

TD
(Ref: [RIPE NCC - LS 3 -E](#))

Source: RIPE NCC

Title: LS/i/r on the results of the evaluation and review of draft Recommendation Y.IPv6RefModel (reply to SG20 LS71) [from RIPE NCC]

Purpose: Information

LIAISON STATEMENT

For action to: -

For comment to: -

For information to: SG20

Approval: -

Deadline: -

Contact: Marco Hogewoning
RIPE NCC

Contact: Chris Buckridge
RIPE NCC

This liaison answers [SG20 - LS 71](#).

A new liaison statement has been received from RIPE NCC.

This liaison statement follows and the original file can be downloaded from the ITU ftp server at <http://handle.itu.int/11.1002/ls/sp16-ripence-iLS-00003.docx>.

Keywords: IPv6; RIPE

Abstract: This Liaison Statement contains the results of the evaluation and review of draft Recommendation Y.IPv6RefModel by the RIPE community via the RIPE IPv6 Working Group mailing list and during the RIPE 76 meeting held from 13-18 May in Marseille, France.

Based on discussions in meetings of Study Group 20, and reflecting advice from a number of Study Group participants including the RIPE NCC, the Study Group decided earlier this year to seek feedback from the RIPE community on the draft Recommendation Y.IPv6RefModel.

RIPE (Réseaux IP Européens) is an open, regional forum that brings together network operators and relevant experts. It serves both as a venue for sharing experience and developing best practices, and as the authoritative body for IP address assignment policy in the RIPE NCC service region¹.

The evaluation of the draft Recommendation brought forward a number of key points that relate to both address policy and network operations, primarily:

- There is strong and consistent feedback from the RIPE community that this kind of standardisation work should not take place in an ITU venue, and that the technical problems in the draft Recommendation stem directly from its development in isolation from the experience and expertise of the network operator community. While participants in SG20 clearly have a legitimate interest in this topic, the development of such a Best Practice document should take place in one of the network operator community structures or the Internet Engineering Task Force.

Some of the specific technical problems identified in the draft Recommendation include:

- The connection between IPv4 and IPv6 in the numbering plan replicates IPv4 constraints in the IPv6 address space
- The proposed address plan cannot (easily) be implemented in network design or would cause severe impact on network operations and routing
- The Internet of Things does not require a dedicated addressing plan and the draft Recommendation fails to provide a rationale for such an approach.

RIPE Community Liaison Process

Following the Study Group's Liaison statement of 24 January 2018, the draft Recommendation was made available to RIPE's IPv6 Working Group, with an invitation for members of the working group to review and comment on its contents via the mailing list. In addition, the document's editor, Mr Ziegler, was invited to present at the RIPE 76 Meeting in Marseille and discuss the document with meeting attendees (an invitation he took up via a remote video link).

This meeting took place on Thursday, 17 May 2018 and a video recording of the proceedings, as well as live transcripts, are available on the RIPE 76 meeting website:

<https://ripe76.ripe.net/archives/#thu>

Archives of the RIPE IPv6 Working Group mailing list discussion can be accessed via the RIPE Forum: <https://www.ripe.net/participate/mail/forum/ipv6-wg>

¹ <https://www.ripe.net/about-us/what-we-do/ripe-ncc-service-region>

As of 20 June 2018, there was discussion of Y.IPv6RefModel in the following threads:

- **[ipv6-wg] Invitation to supply feedback on ITU draft Recommendation on IPv6 address planning for IoT** (9 mails): <https://www.ripe.net/participate/mail/forum/ipv6-wg/PDdmMGIyMzhiOTU0ODQ4NGM4Yjg5NjAxNzZmMjI1N2JmQGVsaXNhLmZpPg==>
- **[ipv6-wg] Y.IPv6RefModel is out of scope for the ITU** (8 mails): <https://www.ripe.net/participate/mail/forum/ipv6-wg/PDAzY2JiOTcxLTg4ZmYtMjhiNi1kODQzLWM3YjEyZDZjMDFjY0Bmb29iYXJub3JnPg==>
- **[ipv6-wg] comments on Y.Pv6RefModel** (11 mails): <https://www.ripe.net/participate/mail/forum/ipv6-wg/PEQzNjVBRDdELTAzNTQzNDY4Qj04MEQ3LURGQ0IxNzJCQzFGNEByZmMxMDM1LmNvbT4=>
- **[ipv6-wg] itu document on ipv6 addressing for iot** (5 mails): <https://www.ripe.net/participate/mail/forum/ipv6-wg/PDE2NjcxNzc2NDEuODQ4Nzc2MS4xNTI3MTcwNzUzMzgxQG1haWwueWFob28uY29tPg==>
- **[ipv6-wg] Still open for discussion: Invitation to supply feedback on ITU draft Recommendation on IPv6 address planning for IoT** (4 mails): <https://www.ripe.net/participate/mail/forum/ipv6-wg/PDczZGFIMTcwNTBjMjQyMzZiZmRhY2FmMmZmYmZmOTAyQGVsaXNhLmZpPg==>
- **Re: [ipv6-wg] Invitation to supply feedback on ITU draft Recommendation on IPv6 address planning for IoT** (1 mail): <https://www.ripe.net/participate/mail/forum/ipv6-wg/PEZDQTdBNkZCLUZBMjMtNDBDMi04REI1LUNGnjNGMTBENjlcNEBmcm9iYm10LnNIPg==>

Synopsis of RIPE Community Feedback

The feedback of the RIPE community should be considered in full. In the remainder of this document, however, we have attempted to distill some of the major points from that feedback, with relevant quotes from the mailing list and RIPE 76 meeting (URLs for the specific emails are included for reference at the end of the document).

A number of submissions stressed the point that any best practices regarding Internet addressing be developed within one of the Internet network operator communities or the IETF, as such work will necessarily emerge from the practices already being employed by operators:

“Whilst it's possible to define certain best practices that can apply, these are invariably based on operational experience and are not prescriptive (i.e. 'bottom up' rather than 'top down.')” [Jan Zorz]

“Internet standards were developed by organisations involved with its operation, and best practices evolved out of their direct operational experience. The network operator communities (e.g. RIPE) are therefore best placed to understand, define and implement IPv6 addressing plans for IoT devices.”

Other interested organisations including the ITU are of course encouraged to participate in this process, raise issues of concern to their communities, and contribute to the best current operational practice, but should not seek to prescribe solutions unilaterally.” [Jan Zorz]

“Solely based on this draft I would have to conclude that IPv6 (addressing) is probably not the ITUs area of expertise. Other venues/bodies, e.g. this WG, might be more suitable for giving guidance.” [Markus de Brun]

“If, and this is a big if, some coordination is to be created regarding best practices regarding address management of whatever is allocated to the end user, that obviously must be developed in the IETF, or possibly, the RIR community (or a mix between the two -- as normal). The IETF and RIR communities must in this specific case be clear and crisp in the response to the ITU-T that having ITU-T injecting themselves in matters like this is stepping too far into the area where other SDOs are responsible for.” [Patrik Fältström]

“There needs to be a crystal clear statement from the RIR community that this entire exercise is wildly out of scope for the ITU, that SG20 needs to cease and desist from reaching into areas outside their competence, and that the Y.IPv6RefModel work item needs to end, immediately and permanently.” [Nick Hilliard]

A number of commenters noted that the lack of demand identified for this kind of standardization, either by operators or in the draft Recommendation itself:

“The document suggests an “optional reference model of an IPv6 addressing plan for Internet of things (IoT) deployment by smart cities, public administrations and companies”. However, I am missing any identification of the special needs of these target audiences, any explanation of specific needs of “IoT” (with that being resolved per ITU-T Y.2060 as a “global infrastructure”), or the benefits of a unified addressing plan at all... In my view, the document clearly fails to motivate the need for a reference model for an addressing plan and any special role for “IoT” in any such plan (read: there is no problem statement).” [Peter Koch]

At several points the discussion focused on the disconnect from operator experiences in developing this draft, and noted that it is based on certain “outdated” concepts:

“The plans in section 11, finally, deal far more with the “non IoT” parts of the network (DMZ, servers, regular LAN) without justifying the recommendations based on operational experience. Also, while there is a postulated size for the address space for “IoT”, the size and need is not justified, nor is there any inner differentiation for systems or objects in the “IoT” category. Assigning addresses per building or floor is already outdated in today's networks.” [Peter Koch]

“As IoT devices are connected to the Internet like any Internet Connected Device there is no difference between IoT deployment and non-IoT deployment. The contrary. We will have IoT deployment everywhere. In every house, home, car, oven, airplane, ship, watch we will have some IoT devices or IoT functionality. Having specific IoT reference models and then non-IoT allocation models as well will absolutely create a digital divide and fragmentation of the Internet. If any reference model is to be created (or Best Practice) it must be created from the general use of IPv6 in the world, and not for IoT specifically.” [Patrik Fältström]

“In Section 11 there is a requirement that the geographical location of a device is to be visible as a specific set of bits (the first four). This implies a device is changing IPv6 address when being moved from one building to another, and that one can not have a subnet covering more than 16 buildings. I do not give more comments on this.” [Patrik Fältström]

“This plan however seems too far removed from operational reality to be really usable.” [Sander Steffann]

There were a number of comments related to use of IPv4 as a basis for the IPv6 number plan:

“I understand the incentive to map IPv4 prefixes to IPv6 prefixes in the short term. The problem with that strategy is that the given approach results in a very bad IPv6 addressing plan. And it won't be applicable unless the IPv4 addressing plan already conforms to a whole set of conditions. If the IPv4 plan deviates even slightly from the proposed plan then it can't be used.” [Sander Steffann]

"...the proposed plans are heavily encumbered by an unmotivated "consistency" with IPv4 addressing." [Peter Koch]

"Attempting to map IPv4 and IPv6 in a 1:1 fashion carries all the legacy issues we have with IPv4, like broadcast-based applications leading to undesirable network topologies, over into the IPv6 world." [Benedikt Stokebrand]

"The whole document includes in a number of cases statements that claim transmission from IPv4 to IPv6 is easier if there is a mechanical mapping between the addresses in the two protocols. This is wrong in so many ways. First of all that there is a transition. What is happening is that IPv6 is deployed. Then one day maybe IPv4 is depleted, or rather, made less important. Sure, various applications will move from using IPv4 to using IPv6, but having application layer transition their default protocol has not much or anything to do with the address plan. Further, due to the limited number of IPv4 addresses, the subletting (if any) in IPv4 space is constrained in a way that IPv6 do not have to be. Limiting IPv6 address plans to the same scheme we have been squeezed into for IPv4 is a bad idea from the beginning. Its like adding ICT to existing processes instead of reviewing and optimizing processes the way possible thanks to ICT."
[Patrik Fältström]

In relation to the connection between the addressing plan and network design, some commenters noted that:

"IP addresses by their very design are supposed to hold all the information needed to route IP packets from their source to their destination. As such IP addresses must be assigned in a way that matches the chosen network topology, and nothing else."

The proposed model however effectively re-purposes two octets of data for purposes unrelated to routing." [Benedikt Stokebrand]

"A general reality is that network operators determine their IPv6 addressing plans depending on the type of network they are building and the purpose. Whilst it's possible to define certain best practices that can apply, these are invariably based on operational experience and are not prescriptive (i.e. 'bottom up' rather than 'top down'). IoT devices will be utilised in many different ways and in different types of networks, and whilst there may be merit in starting to think about IPv6 addressing plans for IoT, these should be formulated based on the collective experience of multiple diverse network operators." [Jan Zorz]

Some people also noted that the current draft has security-related implications:

"In enterprise and other data center environments, microsegmentation and hierarchical security zones as well as a number of other, more specific designs, are used to reach a sufficient level of security."

All of these measures however require a network design that can't be implemented within the constraints of the proposed model except through an excessive bloat of the routing tables, firewall configurations and application based access control lists involved." [Benedikt Stokebrand]

"Section 11 talks about DMZ in singular. This is built upon the thinking used in the 20th century that one have a firewall (i.e. one firewall) that separates the "inside" from the "outside" and that everything on the "inside" of the firewall was

secure, safe and trusted. And further, that devices that where to be reached from the outside had to be in a so called DMZ. This is a model that no longer is valid. All devices and services should not be reachable, and this can not be based on some firewall. Every active network element that can filter traffic must do so, and the devices themselves must filter explicitly what services they expose and make available. And this independently of whether it is on the "inside" or "outside" of whatever layer in the layered protection designed in the network.” [Patrik Fältström]

A number of these issues were also addressed during the RIPE 76 meeting of the working group, as can be seen in the video and transcripts on the RIPE 76 website, and in the meeting minutes. In particular however, we would like to highlight the comments from Tahar Schaa, on behalf of the German Ministry of the Interior, following an exchange about the status of an ITU-T Recommendation:

“Tahar said that the German government is enrolling in the public administration IPv6 already, with success. The schemas that they have used have nothing to do with what Sébastien presented. So if this recommendation becomes local legislation, this will stop Germany from enrolling IPv6 in the public administration.”[RIPE 76 IPv6 Working Group meeting minutes]

In line with his written comments, IPv6 Working Group Co-Chair Benedikt Stokebrand highlighted that the proposed addressing plan would impact the lifespan of IPv6 if it were implemented, as the Recommendation limits available addresses:

“if you are repurposing bits of the address fields of the headers for things other than what addresses are meant for, which means routing information, what you do is you artificially shorten the addresses and with an internet growing at about 30% a year, that basically means you will shorten the lifetime of IPv6 by at least 20 years.”

Conclusion

The RIPE NCC, as a Sector Member of the ITU-T and a participant in Study Group 20, would strongly recommend that the Study Group review the comments of the RIPE community and evaluate whether, based on that feedback, the Study Group should continue with this Work Item.

References

[Jan Zorz] <https://www.ripe.net/participate/mail/forum/ipv6-wg/PDdmMGIyMzhiOTU0ODQ4NGM4Yjg5NjAxNzZmMjI1N2JmQGVsaXNhLmZpPg==#PGY3OTgyOTc3LWE5NDEtMjQwYy03OTI5LTU3NGU4MGJiNjIwZkZkbnZyuc2k+>

[Markus de Brun] <https://www.ripe.net/participate/mail/forum/ipv6-wg/PDdmMGIyMzhiOTU0ODQ4NGM4Yjg5NjAxNzZmMjI1N2JmQGVsaXNhLmZpPg==#PDIwMTgwNTA5MTM1My4wMzE4My5tYXJrdXMuZGVicnVlbnZkbnZyuc2k+>

[Patrik Fältström] <https://www.ripe.net/participate/mail/forum/ipv6-wg/PEZDQTDtBNkZCLUZBMjMtNDBDmi04REI1LUNGNjNGMTBENjICNEBmcm9iYml0LnNIPg==#PEZDQTDtBNkZCLUZBMjMtNDBDmi04REI1LUNGNjNGMTBENjICNEBmcm9iYml0LnNIPg==>

[Nick Hilliard] <https://www.ripe.net/participate/mail/forum/ipv6-wg/PDAzY2JiOTcxLTg4ZmYtMjhiNi1kODQzLWM3YjEYyZDJjMDFjY0Bmb29iYXJub3JnPg==#PDAzY2JiOTcxLTg4ZmYtMjhiNi1kODQzLWM3YjEYyZDJjMDFjY0Bmb29iYXJub3JnPg==>

[Sander Steffann] <https://www.ripe.net/participate/mail/forum/ipv6-wg/PDdmMGIyMzhiOTU0ODQ4NGM4Yjg5NjAxNzZmMjI1N2JmQGVsaXNhLmZpPg==#PEQ2MjcxNDhGLTFCMzctNDU4NC05NThBLUY5MDg3NzVGOzYyNkZkdGVmZmFubi5ubD4=>

[Peter Koch] <https://www.ripe.net/participate/mail/forum/ipv6-wg/PDdmMGIyMzhiOTU0ODQ4NGM4Yjg5NjAxNzZmMjI1N2JmQGVsaXNhLmZpPg==#PDIwMTgwNTI0MDkzNjA1LkdDNDU4NjNAeDExNjcyLmFkbS5kZW5pYy5kZT4=>

[Benedikt Stokebrand] <https://www.ripe.net/participate/mail/forum/ipv6-wg/PDdmMGIyMzhiOTU0ODQ4NGM4Yjg5NjAxNzZmMjI1N2JmQGVsaXNhLmZpPg==#PDg3bGdjOXI0Y3QuZnNmQG1vYS5zdGVwbGFkZGVyLWl0LmNvbT4=>

[RIPE 76 IPv6 Working Group meeting minutes] <https://www.ripe.net/participate/ripe/wg/ipv6/minutes/ipv6-working-group-minutes-ripe-76>
