

Best Current Operational Practice for operators: IPv6 prefix assignment for end-customers - persistent vs non-persistent, and what size to choose.

DRAFT v.4 (26th July 2017)

Authors: Jan Žorž <zorz@isoc.org>, Sander Steffann <sander@steffann.nl>, Primož Dražumerič <Primoz.Drazumeric@telekom.si>, Mark Townsley <townsley@cisco.com>, Andrew Alston <andrew.alston@liquidtelecom.com>, Gert Doering <gert@space.net>, Jordi Palet <jordi.palet@consulintel.es>, Jen Linkova <furry@google.com>, Luis Balbinot <lbalbinot@brfibra.com.br>, Kevin Meynell <meynell@isoc.org>, Lee Howard <lee.howard@retevia.net>

Working Groups: RIPE BCOP TF and RIPE IPv6 WG

Table of content:

[1. Executive Summary](#)

[2. What is a BCOP?](#)

[3. Introduction and incentives](#)

[4. Size of end-customer prefix assignment: /48, /56 or something else?](#)

[4.1. Numbering the WAN link \(interconnection between our network and the end-customer CPE\):](#)

[4.1.1. /64 prefix from a dedicated pool of IPv6 prefixes](#)

[4.1.2. Unnumbered](#)

[4.1.3. ULA](#)

[4.1.4. /64 prefix out of the IPv6 prefix assigned to the end-customer](#)

[4.1.5. Summary](#)

[4.2. Prefix assignment options](#)

[4.2.1. /48 for everybody](#)

[4.2.2. /48 for business customers and /56 for residential customers](#)

[4.2.3. Prefixes, longer than /56](#)

[4.2.4. Considerations for Cellular Operators](#)

[5. End-customer IPv6 prefix assignment: Persistent vs non-persistent](#)

[5.1. Why non-persistent assignments may be perceived as “easier” than static ones.](#)

[5.2. Why non-persistent assignments are considered harmful.](#)

[5.3. Why persistent prefix assignments are recommended.](#)

[6. Acknowledgements](#)

1. Executive Summary

This document discusses the main issues related to the operational practices for the assignment of IPv6 prefixes for end-customers.

Making wrong choices when designing your IPv6 network will sooner or later have negative implications on your deployment and require further effort such as renumbering when the network is already in operation. The temptation to take “easy” approaches for quicker deployment should therefore be resisted.

As a generic set of recommendations, you should consider the following:

- a) **IPv6 is not the same as IPv4.** In IPv6 you assign a short prefix to each end-customer site, so they are able to have as many subnets (/64s) as they need. You should not be concerned with exhausting the IPv6 addressing space, and you should think big when planning future requirements. If you need more space, you can go back to your Regional Internet Registry and providing your addressing plan justifies it, you can obtain more IPv6 addresses.
- b) It is **strongly discouraged** to assign prefixes longer than /56, so your choices are:
 - 1) If you want a simple addressing plan, /48 for each end-customer. This will work very well for customers coming from other ISPs, those that have their own ULA, or have been using transition mechanisms. This will also be easier when you have a mix of customers using the same infrastructure, whether they are residential customers, SMEs or even large corporates.
 - 2) Differentiate amongst types of customers, even if this will increase the complexity of your network and those of your customers, by offering /48 for business customers and /56 for residential customers.
 - 3) A trade-off amongst the previous two options by reserving a /48 for residential customers, but actually just assigning them the first /56.
There a specific case for cellular phones to be assigned a single /64 per each PDP context, but this is out of scope of this document.
- c) In order to facilitate troubleshooting and have a future proof network, you should consider **numbering the WAN links using GUAs** (Global Unicast Addresses), using a /64 prefix out of a dedicated pool of IPv6 prefixes. If you decide to use /127 for each point-to-point link, it is advisable to allocate a /64 for each link and just use one /127 out of it.
- d) Non-persistent prefixes are considered harmful in IPv6 as you can't avoid issues that may be caused by simple end-customer power outages, so assigning **persistent prefixes is a safer and simpler approach**. Furthermore, this avoids the need for expensive logging, increases your chances to offer new business to customers, and decreases your customer churn.
- e) In countries with **specific legal environments**, ISPs are encouraged to support both stable (persistent) and privacy-oriented (non-persistent) prefixes as options for customers. Persistent prefixes are recommended as the default, in the absence of legal requirements to the contrary.

This approach is probably suited for most operator cases. However, we strongly encourage you to read through the rest of this document in order to understand the reasons for these recommendations. This will help you draw your own conclusions, and ensure that your solution matches your specific requirements.

2. What is a BCOP?

A Best Current Operational Practices (BCOP) document describes best current operational practice on a particular topic, as agreed by subject matter experts and periodically reviewed by the Internet community. This document is not intended to document what practices may be in future, but to reflect the best methods of implementing IPv6 at the time of publication. Updates to this document will be published to reflect changes in best current practices where there are developments in standards and implementations.

3. Introduction and incentives

The target audience of this document is Internet Service Providers (ISPs), which provide or intend to provide IPv6 services to end-customers, whether residential or business.

The task of ISPs is to provide their customers with Internet connectivity, and that includes providing those customers with IP addresses. Customers generally receive a single IPv4 address, but with IPv6 a customer requires a block of addresses, also known as a prefix. This document provides guidance as to what size such prefixes should be.

Network operators may find unexpected or unwanted results if they make a wrong decision when choosing the size of a prefix assignment, or the method of assigning it. This may be immediate in terms of other networks or content providers measuring the IPv6 “brokenness” who may stop serving you AAAA records, or it may happen in the future when, for example, the HOMENET standard results in more complex home networks being built and operated. In these cases, network operators will have to adapt their IPv6 implementations to new requirements, possibly including complete network renumbering, whilst their IPv6 network is already in production and has connected customers.

Up until now, there have been no clear recommendations as to what size of prefixes to choose and assign to customers, hence the decision to document some best current operational practices for network operators to consider when making their decisions.

4. Size of end-customer prefix assignment: /48, /56 or something else?

Network operators deploying IPv6 will sooner or later need to decide which size of IPv6 prefixes they will assign to their end-customers. At this point, the IPv4 scarcity needs to be

reconsidered because the abundance of IPv6 addresses enables numbering decisions to be taken differently.

Many operators might perceive that assigning large IPv6 prefixes to end customers is wasteful, but the reality is that decisions should be based on the IPv6 protocol architecture design. For example, Tony Hain calculated that assigning a /48 to every human on Earth, and never recovering those, will still mean that IPv6 would have a lifetime over the 480 years and we could repeat that several times. On that timescale, there will be other reasons, not just scarcity of IPv6 addresses, that will require the IETF to design a successor to IPv6.

4.1. Numbering the WAN link (interconnection between our network and the end-customer CPE):

Before going into details about the size of IPv6 prefix assignments, the choice for the WAN link needs to be understood. There are three options for addresses on the link between the operator network and the “end-customer” CPE (Customer Premises Equipment) WAN port. Note that CE (Customer Edge) is also commonly used for the CPE (RFC7084).

4.1.1. /64 prefix from a dedicated pool of IPv6 prefixes

Using a /64 prefix from a dedicated pool of IPv6 prefixes is the most common scenario and currently the best practice. A separate block of IPv6 space is allocated for the WAN links to the end customer CPEs, so that when CPE connects to the network and performs router discovery, a /64 prefix is used to number both ends of the connection.

In the case of an end-customer host (not router) connecting over PPPoE (for example), the setup process is concluded when the host is properly IPv6 numbered and can start sending and receiving traffic.

In the case of a router connecting to a network that is capable of issuing a DHCPv6 Prefix Delegation (PD) request, the IPv6 prefix delegation process is started and an IPv6 prefix is assigned to the CPE. This method is most probably the safer one, covering most of the scenarios regardless of underlying technologies, and is the one most used by operators.

Some operators also use /126, /127, /112 or an alternative prefix assignment instead of /64 for managed links where addresses on the link are manually configured, and according to RFC6164, a /127 is suggested. Please note, that not all equipment currently supports this option, so /64 is still the safest approach and has the advantage of being future proof in the event that new standards make usage of the other 64 bits for other purposes or the link becomes point-to-multipoint, etc. Furthermore, some ISP hardware has limitations in using prefixes longer than /64, so the use of non-/64 prefixes will use more resources on these devices. If we decide to use /127 for each point to point link, then it is also advisable to allocate a /64 for each link and use just one /127 out of it to prevent the Neighbor Discovery exhaustion attack (RFC6583).

Finally, one additional benefit for a dedicated pool, is the possibility to apply security policies without harming the customers. This is only true if customers **always** have a CPE at their end of the WAN link.

4.1.2. Unnumbered

Some operators decide to leave the WAN link to the end-customer CPE unnumbered which means no GUA (IPv6 Global Unicast Addresses), so only link-local addresses are used at both ends of the link. While this might work for routers, it does not work for devices that cannot request a prefix delegation over DHCPv6 and are therefore left without any usable GUA to allow traffic in and out. In the case of a router on the end customer side, the route for the assigned prefix is pointed towards the link-local address on the CPE/router WAN port and the default route on the CPE/router is pointed towards the link-local address on the upstream network equipment port. This method may be seen as easier to implement, but it also brings some drawbacks such as difficulties with troubleshooting and monitoring that need to be considered as, for example, link local addresses do not appear in traceroute, the most common troubleshooting tool used, so you will not be able to easily locate the exact point of failure.

It is most useful in scenarios where it is known that only CPEs will be attached to the WAN link, and where a “pingable” address of the CPE is known (e.g. because the ISP-provided CPEs are always “pingable” on the first delegated address). It is less suitable in situations where unknown CPEs and/or non-CPE devices are attached to the WAN link.

In addition, non-routers (e.g. a older Windows or Linux PC or any other end-customer host) connecting to a network that have no ability to ask for prefix delegation over DHCPv6 might experience problems and will stay unnumbered upon connection if a /64 prefix is not used to number the CPE WAN link. This may be also the case for many CPEs, which will not be able to complete the DHCPv6-PD if they got an unnumbered link.

Finally, certain hardware in the ISP infrastructure may consume resources when using numbered links. This is a very specific situation that you may need to consider.

4.1.3. ULA

Some operators, use ULAs (IPv6 Unique Local Unicast Addresses) for numbering the WAN link to the end-customer CPE. This approach may cause numerous problems and is therefore **strongly discouraged**. For example, if the CPE needs to send an ICMPv6 message to a host outside an operator’s network (to the Internet), the packet with ULA source address will not get very far and Path MTU discovery (PMTUD) will break, which in turn will completely break that IPv6 connection when MTU is not the same for all the path.

4.1.4. /64 prefix out of the IPv6 prefix assigned to the end-customer

Some operators use the first or last /64 prefix to number the CPE WAN link, from the larger IPv6 prefix assigned to the end customer. This has been described by an old IETF Internet Draft (<https://tools.ietf.org/html/draft-palet-v6ops-point2point-01>).

This is sometimes a good idea, but in some cases may fail as some CPEs or scripts may try to use it for loopback and assign the next /64 to the LAN ports. Responsibility for the prefix is delegated to a CPE, but then one /64 is used for the WAN connection. Whilst there is a DHCPv6 option to let the CPE know if one of its prefixes have been “stolen” (RFC6603), not all CPEs support this and some issues will be encountered.

This is an option that simplifies routing and general provisioning provided you’re sure that the CPEs in your network don’t have the aforementioned issues.

4.1.5. Summary

Using a global /64 prefix for the WAN connection is the recommended choice.

Besides being a safe choice, using a /64 is sometimes required when there are more devices than the two endpoints on a WAN link (e.g. intermediary bridges or repeaters) that require management, or if there is the need for redundancy (e.g. VRRPv3 or multiple routers at the customer premises).

When choosing which /64 to use, the recommended option is to dedicate a separate pool of prefixes for the WAN links. While the addressing plan and administration might be easier when selecting the /64 from the prefix delegated to the customer, this is technically “stealing” because the customer’s CPE has been informed that whole prefix will be delegated to it, so it should not also be used on the WAN link unless it is known that all CPEs will support RFC6603 to negotiate this.

When providing address space to a smaller ISP you should give them a sub-allocation (not an assignment) that allows them to make properly sized (see the rest of this BCOP) assignments to their customers.

Using ULA addresses on the WAN link is very strongly discouraged.

4.2. Prefix assignment options

To keep addressing plans usable and understandable, and to align with DNS reverse zone delegations, the size of the delegated prefix should align with a nibble boundary. Each hexadecimal character in an IPv6 prefix represents one nibble, which is 4 bits. The length of a delegated prefix should therefore always be a multiple of 4.

A single network at a customer site will be a /64. RIR policies permit assignment of a /48 per site, so the possible options when choosing a prefix size to delegate are /48, /52, /56, /60 and /64. However, /64 is not sustainable, doesn’t allow customer subnetting, and it doesn’t follow IETF recommendations of “at least” multiple /64s per customer. Moreover, new IETF work allows the assignment of a complete /64 to a single interface.

The following sections explain why /48 and /56 are the recommended prefix assignment sizes for end customers.

4.2.1. /48 for everybody

This is probably the most practical and pragmatic way to assign IPv6 prefixes to end customer CPE devices. In this case everyone has a /48 prefix and advanced end-customers are less likely to make mistakes when addressing their networks and devices, resulting in much less call-centre time to sort out problems. It also has the advantage of sharing the same prefix size as ULAs and some transition mechanisms, so this facilitates a direct mapping of existing customer addressing plans to the delegated prefix.

Example of assigned prefixes out of 2001:db8::/40 :

```
2001:db8:aaaa::/48
2001:db8:aaab::/48
2001:db8:aaac::/48
2001:db8:aaad::/48
```

For the first prefix, **2001:db8:aaaa:XXXX::/48**, the XXXX is the part that the end-customer can use to make their own addressing plan. This is also the structure used in most addressing plan manuals and guides.

4.2.2. /48 for business customers and /56 for residential customers

Some operators decide to give a /48 prefix to their business customers and a /56 to their residential customers. This rationale is understood to be mainly coming from sales and marketing departments where they wish to create some distinction in services between different types of customer. This method can be considered as pragmatic, future-proof and has nearly no downsides, the same as “/48 for everyone approach”.

For more advanced home users who manually configure their servers and routers, local addressing is a bit harder for /56 prefixes, as you are unable to use all 4 digits between the double colons in the address notation. This means that there is a higher risk of making mistakes whilst making the addressing plans and sub-assign address to the devices outside the assigned prefix. This should not be a problem in simple cases where a CPE assigns all prefixes, but it might cause annoyances for those advanced users that want to offer content and services from their home network.

It should be remembered that some protocols/mechanisms use a default /48 prefix size (tunnel brokers, ULA, 6to4, among others), so if a customer already has an addressing plan, they will be forced to redo it and renumber in case of a different prefix size. Instead, using the same prefix size, allow them to keep the addressing plan and just renumber some of the prefix bits rather than have a complete numbering plan change.

Example of assigned prefixes out of 2001:db8::/40 :

```
2001:db8:aaaa:1a00::/56
2001:db8:aaaa:1b00::/56
2001:db8:aaaa:1c00::/56
```

2001:db8:aaaa:1d00::/56

For the first prefix, **2001:db8:aaaa:1aXX::/56**, the XX is the part that advanced customers can use to make their own addressing plan.

An alternative is to reserve a /48 for residential customers, but actually assign them just the first /56. If subsequently required, they can then be upgraded to the required prefix size without the need to renumber, or the spare prefixes can be used for new customers if it is not possible to obtain a new allocation from your RIR (which should not happen according to current IPv6 policies).

Finally, take in consideration as well, that SMEs often use residential services, which means that you may be actually assigning /56 to business customers.

4.2.3. Prefixes, longer than /56

It is **strongly discouraged** to assign prefixes longer than /56 unless there are very strong and unsolvable technical reasons for doing this.

There are enough IPv6 addresses to delegate end-customer's a /48, so a /56 already represents a sizeable restriction. There is no need to delegate fewer addresses than that, so if your IPv6 allocation is insufficient to provide a /56 to each end customer (remember there are 134 million /56s in a /29 and 16,7 million in a /32), explain to your RIR that your initial allocation was too small and that you require a larger allocation based on your IPv6 implementation plan. Offering less than a /56 can be feasible if just a /29 is allocated and 6RD is being used, as an IPv4 address is embedded in the IPv6 prefix. However, even in this case, a larger allocation from your RIR can be justified and the transition protocol allows other ways to sort it out.

Assigning a /64 or longer prefix is really dangerous and **highly discouraged**, does not conform to IPv6 standards, and will break functionality in customer LANs. With a single /64 the end customer CPE will have just one possible network on the LAN side and it will not be possible to subnet, assign VLANs, alternative SSIDs, or have several chained routers in the same customer network, etc...

Some CPEs use a /64 for the loopback interface and some may have multiple LAN segments predefined (for example Guest WiFi network and LAN), so as soon as there is more than one LAN segment behind the CPE, exceptions will have to be added to the ISP provisioning system that will greatly complicate management. It is not possible to assign less than a /64 to each LAN port/segment/subnet/VLAN due to IPv6 protocol requirements (SLAAC, ND, etc..) that reserve the last 64 bits of an IPv6 address for the host.

Moreover, there is work in progress at IETF showing that it can be convenient in some cases to assign a /64 to each host (e.g. for security, isolation of customers, or having many virtual machines on a single host), again explaining the need for delegating many /64s to each customer.

A growing number of operators are also using prefix colouring to deploy products with distinct Service Level Agreements (e.g. voice, data, video) and this requires at least a unique /64 for each product or service. If combined with home networking technologies under development, the number of prefixes increases quite quickly and delegation of multiple IPv6 prefixes at the same time will occur, so assignments of less than /56 will probably require renumbering in the near future.

The Broadband Forum also recommends a similar approach in their TR-177 document (2010): <https://www.broadband-forum.org/technical/download/TR-177.pdf>

"A delegated prefix for use within the home network (mandatory). The Broadband Forum suggests a size for the delegated prefix of at least a /60 for home network or SOHO environments, with a recommended prefix length of /56. The delegated prefix may be extended to a /48 for larger organizations."

4.2.4. Considerations for Cellular Operators

There is a clear exception to the rule described above when assigning prefixes in a cellular network. In this case, a /64 will need to be provided for each PDP context for cellular phones, whereas for LTE modems/routers, i.e. in the case of broadband by means of cellular access, it will still be necessary to choose a /48 or /56 in accordance with the aforementioned considerations.

5. End-customer IPv6 prefix assignment: Persistent vs non-persistent

After determining the size of the IPv6 prefixes to be assigned to end customers, it is necessary to establish how to actually assign them. There are two methods: Persistent and non-persistent assignment.

Persistent assignment means that a prefix is assigned to a customer (typically by means of an AAA - Authentication, Authorization and Accounting, which may be dependent on DHCPv6 and/or other provisioning systems) and that prefix will always remain the same regardless of how many times the customer (re)connects or renews the lease. It is "persistent" for the same customer in the same link/location regardless of the provisioning system being used.

DHCPv6-PD is commonly used for non-persistent assignments, where a bigger pool of IPv6 space is assigned to each termination point, and so as customers connect, they are randomly assigned prefixes from this when they (re-)connect or the lease time expires. They may get the same prefix, but typically it will be different (non-persistent) unless the lease time is so high that it effectively become "persistent". In this case, a customer will get the same prefix even if they switch off their router for several days, weeks or months, depending on the lease time.

5.1. Why non-persistent assignments may be perceived as “easier” than static ones.

When faced with the task of how to deploy IPv6, it is easy to go for an option that initially requires less effort, time and energy. However, this will usually create more problems later on.

The easiest method is to assign prefixes from a pool of IPv6 prefixes to termination points (BNGs or other equipment, depending on the type of access network) and let the provisioning system decide what customers will get when they connect and ask for a prefix delegation over DHCPv6. This practice is carried over from the IPv4 world where addresses are commonly assigned dynamically and where CPEs perform NAT to conserve IPv4 space. Bear in mind that end customers with an IPv4 subnet behind their CPE never got “non-persistent” assigned IPv4 prefixes as this would require reconfiguration of all hosts on their network every few hours or days. By contrast, every IPv6 customer gets an IPv6 subnet so it is unnecessary to apply this “IPv4 model” to IPv6.

Non-persistent prefix assignment also appears initially easier as it facilitates aggregation of internal routing tables according to end customer connection termination points. Every termination point has its own pool of IPv6 prefixes that are nicely aggregable, whilst with persistent IPv6 prefix assignments it is necessary to discover which customer is terminated at which termination point, group them into larger IPv6 pools, and then update our database accordingly. This is unless your provisioning system is doing it in the other way around from an AAA database and is typically tied to an IPAM for the initial assignment to each new customer.

Some operators may ask what happens if a customer moves to another city or neighbourhood, but if they move to different location then all the equipment is switched off whilst transporting it. Since it will be necessary to provision the new connection elsewhere, then it is perfectly fine to also change the prefix of the customer to match the criteria for that aggregation point unless it is in the same “aggregation region/point” or the customer specifically requests that the same IPv6 prefix moves with them. In this case, the internal routing table will be one entry bigger and the decision may be to offer that service at an extra cost, but that should be an uncommon scenario.

5.2. Why non-persistent assignments are considered harmful.

Taking a scenario where there is a connected CPE with a dynamically assigned prefix (e.g. 2001:db8:aaaa::/48). The CPE may sub-assign local loopback addresses out of the first /64 segment (2001:db8:aaaa:0::/64) and sub-assign 2001:db8:aaaa:1::/64 to the first LAN interface (and possibly 2001:db8:aaaa:2::/64 to the guest WiFi interface, and so on). So, the hosts on the network autoconfigure IPv6 addresses out of 2001:db8:aaaa:1::/64 and/or 2001:db8:aaaa:2::/64 segments and start communicating with the rest of the Internet.

Now all of a sudden there is a power outage (which is very common in many regions/countries in the world, even “highly-developed” countries) or the CPE freezes and reboots and the connection has to be established again, but this time with a new IPv6 assigned prefix of 2001:db8:bbbb::/48. If the CPE knows that the delegated prefix has changed it should send out RA (Router Advertisement) packets with a prefix valid lifetime of 0 to tell all devices that the old addresses are no longer valid. However, the CPE rarely knows that before the reboot there was a different prefix on the network, and the packets to revoke the old IPv6 addresses do not get sent. In this case, multiple IPv6 addresses from completely different assigned prefixes end up on the same network interface, some of which will no longer work and may imply increase the number of claims to the call-centre.

Different OS vendors treat this scenario differently, but there often ends up being a wrong source IPv6 address for sending packets through the CPE out to the Internet. As the network operator’s equipment deleted the previous delegated prefix route back to the CPE, any return packets never reach the originating device and IPv6 connectivity will be broken until the old IPv6 addresses time-out and are automatically removed from the interfaces. If another reconnection to the ISP is required in the interim, there will be a third set of IPv6 addresses from yet another assigned prefix, and this will cause even more confusion.

Some big content providers are measuring “IPv6 brokenness” in operator networks by matching the SYN, SYN ACK and ACK packets received/sent to/from a single source. If they see a SYN and send back SYN ACK, but never receive ACK in a timely manner, they slowly stop serving AAAA records for the AS number where they see this. So, if you decide to assign IPv6 non-persistent prefixes, don’t be surprised if content providers start ignoring your IPv6 traffic quite quickly and force your customers back to IPv4. And these are usually destinations (big content providers) where you will be sending a significant amount of traffic.

Using non-persistent prefixes also means that it is necessary to have a logging system that registers which WAN and LAN prefixes are being used at each moment by each customer site in order to comply with legal/regulatory requirements.

Finally, if users have services such as web, email or VPNs, they typically need to manually configure the addresses of those, so using non-persistent prefixes is not an option in this case. With the growth of broadband capacities (such as FTTH), it is becoming more and more common that end-customers run servers or services on their LANs, including the likes of IP cameras or security systems in a home. Persistent prefixes will allow DNS names to be assigned to individual addresses of those prefixes.

If you wish to do non-persistent prefix delegation, you must verify that the CPEs used by the majority of your users will not have the problems described above. If this can’t be ensured, then the recommendation is to avoid using non-persistent prefixes and revisit the CPE and end user device implementations periodically to see whether this has become feasible if you still need that.

It also needs to be realised that non-persistent prefixes have other implications, because it is not only the CPEs that need to be numbered, but all the devices behind them in the customer LANs as well. Different implementations can have very different behaviors that may affect the number of support calls to your call-centre every time a prefix changes.

5.3. Why persistent prefix assignments are recommended.

As already mentioned, the best practice in IPv4 was to assign persistent IPv4 prefixes whilst giving an end-customer a routed prefix, and same is true for IPv6 where persistent prefix assignment is strongly recommended. If required for provisioning, the connection between a network and a customer CPE can be non-persistent using /64 (or even /127 if the equipment supports it), but choosing persistent IPv6 assigned prefixes for end-customer LANs will avoid a lot of difficulties experienced by some earlier IPv6 network operators.

This is especially the case where there are non-residential customers as well, as this means you can have a single provisioning system and it is unnecessary to maintain a logging system. This is because each user/site always have the same prefix(es) and can be identified in the AAA or provisioning system, reducing complexity and making things cheaper. A bit of initial thought, planning and consideration for future needs can therefore save a great deal of time and energy when IPv6 has been deployed.

From the economic/business perspective, another important consideration with persistent prefixes is that it is possible to 'name' value added services using DNS (e.g. camera1.username.ispname.com) that can result in considerable new income streams. Trying to deploy new services or applications with non-persistent prefixes is always more difficult and costly, and will increase time spent on troubleshooting. This is just one example to make clear that persistent prefixes enable innovation, new services and applications, which in turn means new business possibilities.

Further to the above, it should be noted that persistent assignments of prefixes to customers may increase the incentive for customers to remain with that ISP, and result in decreased customer churn.

Finally, if you're unconvinced about the use of persistent prefixes, you should at least consider using non-persistent but long-lived assignments so the lease time is as long as your provisioning system allows.

6. Acknowledgements

The authors would like to thank Nathalie Kunneke-Trenaman, Mikael Abrahamsson, Jason Fesler, Martin Levy, Ian Dickinson, Philip Homburg, Ivan Pepelnjak, Matthias Kluth, Ondřej Caletka, Nick Hilliard, Paul Hoffman, Tim Chown and Roman Nurul Islam for their comments and suggestions as to how to improve this document. Special thanks go to RIPE IPv6 Working Group community and Chairs for accepting this document for technical review, and

also the RIPE BCOP Task Force community and Chairs for ensuring it does conform with actual best operational practice.