

A Secure Mechanism for Address Block Allocation and Distribution

Damien Leroy* and Olivier Bonaventure

Universite catholique de Louvain (UCL) - Louvain-la-Neuve, Belgium
{Damien.Leroy,Olivier.Bonaventure}@uclouvain.be
<http://inl.info.ucl.ac.be>

Abstract. All equipments attached to the Internet are configured with one or several IP addresses. Most hosts are able to automatically request (e.g., from a DHCP server) or discover (e.g., by using stateless autoconfiguration) the IP address that they should use. This simplifies the configuration and the management of hosts. Unfortunately, these techniques do not apply on routers whose IP addresses and subnet prefixes for their directly attached LANs still need to be manually configured. This utilization of manual configuration is error-prone and a frequent source of errors. It is also one of the reasons why IP address renumbering is so difficult with both IPv4 and IPv6. In this paper, we propose a new address block allocation and distribution protocol that has been designed to be both secure and efficient. We first summarize the main requirements of an address block allocation mechanism. We then describe the operation of our proposed mechanism. Finally, we demonstrate the efficiency of our protocol by simulations.

1 Introduction

The growth of the BGP routing tables in the default-free zone is again a concern for many network operators [1]. A key reason for this is the way IP addresses are allocated and used. The pool of available IP addresses is managed by the regional registries (RIPE, APNIC, ...). These registries define two types of addresses : *Provider Aggregatable* (PA) and *Provider Independent* (PI). To limit the size of the BGP routing tables, only large ISPs should obtain PI addresses while customer networks should receive PA addresses from the PI block of their upstream provider. Unfortunately, if a corporate network uses a PA address block, it should renumber all its network when it changes from its upstream provider. For this reason, most corporate networks insist on obtaining PI addresses, even with IPv6 [2]. Combined with the growth of multihoming [3], this explains the growth of the BGP routing tables. This growth could be avoided if corporate networks and smaller ISP networks were able to more easily use PA addresses. Unfortunately, with the current Internet architecture, using PA addresses implies that each corporate network must be renumbered each time it changes from provider and several studies have shown this to be painful with both IPv4 and IPv6 [4].

* Supported by a grant from FRIA (Fonds pour la formation à la Recherche dans l'Industrie et dans l'Agriculture, rue d'Egmont 5 - 1000 Bruxelles, Belgium).

In the early days, IP addresses were allocated manually to both routers and endsystems. However, this manual allocation was a cause of errors and problems. As a consequence, most endsystems now obtain their IP address automatically either via DHCP or via auto-configuration. Despite the widespread use of automatic configuration of endsystems, the addresses used by the routers are still manually configured (except in small networks by using DHCP extensions) and several studies have shown that configuration errors are responsible for a large number of operational problems [5]. In this paper, we propose and evaluate a new distributed mechanism which is able to automatically allocate IPv6 prefixes to routers and LANs in an ISP, corporate or campus network.

Our address distribution mechanism is targeted for edge networks as well as ISP networks that need to provide address blocks to their subnets¹. The typical environment where it should be run is represented at Fig. 1. In this figure, arrows represent the direction of address block allocation. Our mechanism is composed of three main parts. One or several prefixes are obtained from upstream ISPs by border routers (A, B on Fig. 1). Subnets needing an address block ask for it at border routers (D, E, F on Fig. 1). The routers negotiate which parts of the obtained prefixes have to be allocated to subnets. The term “prefix” is used to indicate a prefix allocated by one of our providers (i.e. ISP α or ISP β in Fig. 1) while “address block” is used to indicate a group of addresses allocated by our protocol to subnets.




Fig. 1. The protocol does apply on hierarchical topologies

This paper is organized as follows. We first discuss in Sect. 2 the requirements for such an allocation and discuss a few existing proposals. Sect. 3 provides a detailed presentation of our solution by considering the security aspects, the protocol and the algorithm used to allocate address blocks. Finally, Sect. 4 contains a simulation-based evaluation of our solution. More details about our mechanism can be found in [6].

2 Requirements

In this section, we summarize the main requirements that must be fulfilled by an address allocation mechanism as well as a state of the art of such mechanisms.

¹ In this paper, the term *subnet* is used to define a subnetwork (a LAN or a customer network) that needs to obtain an address block (e.g., customer 1 or LAN in Fig. 1).

Utilization of address space. A first goal is to support an Host-Density ratio (noted HD) equals to 0.8. The HD is a way to measure address space usage [7]. It is expressed as the ratio of $\log(\text{number of allocated address blocks})$ to $\log(\text{maximum number of allocatable address blocks})$. Several regional registries defined the value 0.8 as an acceptable IPv6 address utilization for justifying the allocation of additional address space [8].

Compatibility with existing protocols. Since the role of the mechanism is limited to address block allocation, it must be independent from the routing protocol.

Security. The Internet started as an open network mainly used by researchers. Nowadays, IP networks are used to support mission critical business services and security matters. We discuss the security threats in more details in Sect. 3.1.

Roles. In an ISP, enterprise or campus network, network operators usually group the hosts that have similar roles in contiguous address blocks. For example, all loopback addresses of routers usually belong to the same prefix, all servers are grouped in a few prefixes and the hosts used by students in a campus network are grouped in contiguous prefixes as well. ISPs also group their customers in prefixes depending on their type (e.g. home users, business customers, ...). This allows to define and deploy policies on routers (e.g., QoS, traffic engineering, ...) that are more scalable and easier to maintain. Common examples are the packet filters deployed on most routers and firewalls to filter unwanted packets. Chown et al. have shown that updating packet filters is a difficult problem when renumbering a network [9].

Prefix coloring. Another requirement is that the prefixes received from different providers are not necessarily equals. Due to routing policies, it can be necessary to classify the prefixes received from providers in different categories. For instance, if a National Research Network has been allocated a prefix from a backbone research network, it should only allocate address blocks from it to its customers that are research labs. K12 schools should not be allocated an address block from such a research prefix.

2.1 Existing Address Allocation Mechanisms

Several research groups have studied the problem of automatically configuring the addresses used by routers. First, studies within ad-hoc networks and notably within the *autoconf* working group of the IETF [10] tackle very different type of networks. The topologies they consider are changing frequently while we are looking for a stable configuration. The NAP Protocol (No Administration Protocol) by Chelius et al. [11] is targeted at small networks and unfortunately does not consider the problems faced by campus and enterprise networks such as the need for roles and the security issues. Other solutions proposed at IETF such as DHCP prefix delegation [12] or router renumbering [13] do not meet all requirements. They require manual configurations and cannot be adapted to provide minimum security features without difficulties. A detailed analysis of these solutions is available in [6].

3 Description of Our Solution

This section describes in more details our proposed mechanism. We first present the security risks and solutions. Next, we describe the address block distribution protocol. This section ends with the explanation of the address block allocation algorithm.

The distributed IPv6 addresses are composed of three parts as illustrated in Fig. 2. Here are some notations we use in this paper. The first part is the *prefix* given by the upstream ISP. These leftmost bits are common for all hosts and routers in our network. The last 64 bits of the address are used for the *Interface ID* (IID). The bits between these two parts uniquely identify a subnet and are called *Subnet ID* (SID). We will consider two parts in this SID: the *allocated SID* (ASID) and the *delegated SID* (DSID). The ASID is the part that our mechanism allocates and provides as a prefix to subnets. The DSID is the part of the SID that the subnet is free to allocate inside its network. In other words, a subnet obtains a $l(prefix_s + ASID_s)$ prefix².




Fig. 2. Different parts of an IPv6 address

The protocol for prefix delegation used with the providers of our network and with our customers is outside the scope of this paper. It could be designed as an extension to BGP, using DHCP with prefix delegation option [12] or a new protocol using the X.509 certificates being defined by the IETF *SIDR* working group. The only requirement we impose is security, as discussed below.

3.1 Security Threat Model and Solutions

It is easy to be convinced that an address block allocation mechanism is a key component of an IP network and could be the target of attacks. Indeed, this protocol, in the network configuration process, runs prior to the other IP protocols. Therefore, if this protocol were compromised, the entire network or a large part of it could be compromised.

The first risk comes from a malicious router that could be interpolated in the network. In addition to the security threats on all the routing and IP protocols, this router could send several kinds of messages in order to confuse the address block distribution protocol. It should be possible for a router to detect whether its neighbor is a legitimate router. For instance, if router D on Fig. 1 were unauthorized, router A and E should not trust it and thus not forward its messages.

The second risk could be another possibility of obtaining similar rights but without inserting a router. It consists in intercepting the packets exchanged between two routers and performing a man-in-the-middle attack.

The last two risks are based on abusing a link between border routers and their providers or customers (e.g. between ISP α and router A or between router D and customer 1 on Fig. 1). Invalid prefix announcements can cause a global unavailability or traffic hijacking. Abuse can also appear with address block requests if there are automated and not secure.

² s subscripted means “size” in this paper. For instance, SID_s for the size of the SID .

Solutions. The first idea is that IPv6 prefix delegations and certificate delegations go together. In other words, each network that possesses a prefix has obtained, from its ISP, a certificate proving that it is allowed to use and delegate it. In our mechanism, these are mainly used to authenticate our ISP's routers. Actually, we rely on the X.509 extensions for IP addresses [14].

To allow a provider to authenticate its customers, our network has its own certification authority and each router is configured with a certificate (as suggested by Greenberg et al. [15]) to confirm that it does belong to the network. An X.509 certificate signed by our local certification authority (CA) can be used by its owner to provide an authorization evidence to nodes of our network. A certificate itself identifies its owners and we associate them with attribute certificates [16] to provide authorization. These attributes are used by our subnets to specify their role, their provided prefix colors and their address block size. These certificates could be given offline to customers when the contract is signed and are used when the customer connects to one of our border routers. Secondly, attribute certificates are also used to authenticate connections between routers : a neighbor is considered as a legitimate router if and only if it can be authenticated. The following packets exchanged between routers are secured with IPSec.

3.2 Address Block Distribution Protocol

Our protocol is distributed, each router is responsible for its address blocks and has to choose and advertise them. In practice, each router chooses one or several address blocks and advertises them through the network. A distributed protocol avoids the single point of failure problem. Chelius et al. also showed that it is more efficient for address block distribution, especially if the network is multihomed [11].

Each router is configured with one 64-bits router id (derived from MAC address or crypto based), its X.509 and attribute certificates and information about the subnets attached to itself. The loopback address of each router can be derived from the router id using a fixed SID for all routers.

Routers use *discover* and *hello messages* for, respectively, discovering their direct neighbors and initializing connections with them. Since this is very similar to other protocols such as LDP, we will not enter into detail about these. The *prefix* and *address block advertisement messages* are flooded to all routers.

Prefix advertisement message. When a border router learns that a new global prefix has been allocated to the network, it floods immediately the information. Each prefix is associated with a preferred and a valid lifetime and so must be regularly renewed as long as it remains valid. A *prefix advertisement message* contains : the size of the prefix, the prefix itself, the deprecated and validity lifetime, the prefix color, a sequence number and security parameters. The sequence number is incremented each time a new message about a specific prefix is generated. Until the lifetime expires, an entry per prefix is stored in a prefix table in each router.

Address block advertisement message. When a router chooses new address blocks, it floods an *address block advertisement message*. This message contains a list of address block entries, one for each entry it chooses. When such a message is received, the proposed address blocks are checked and new address blocks are stored with the already

allocated ones. The way address blocks are chosen and stored is explained in Sect. 3.3. An *address block advertisement message* contains the router id of its issuer and a list of address block entries. An entry contains a role, P_s , $ASID_s$, $ASID$, a timestamp, the preferred and valid lifetime and a sequence number. As for prefixes, the sequence number is incremented each time an address block entry is changed (ASID changed, preferred time redefined, ...). As the mechanism is distributed, there is a non-zero probability that two routers choose conflicting blocks nearly at the same time; we call this a collision. The issue is solved by using logical clocks and randomness for tie-breaking.

3.3 Address Block Allocation Algorithm

This section explains the way subnet address blocks are chosen in the address space. The following rules apply to any prefix for which the subnets should obtain an address block. The only part of the address block that our mechanism has to determine is the ASID as shown in Fig. 2. For instance, consider that the network has obtained two /48 prefixes, 2001:4bc7:9377::/48 and 2001:3def:73cb::/48, from its ISPs. If the ASID “1ed6” is allocated a /64 subnet, this subnet will obtain as prefixes 2001:4bc7:9377:1ed6::/64 and 2001:3def:73cb:1ed6::/64.

In our mechanism, each router chooses a large address block that can contain all the subnets attached to itself, we will call it the “router block”. Inside this block, each router allocates address blocks for its subnets, we will call these the “subnet blocks”. In the topology depicted in Fig. 1, it means that router F allocates two address blocks inside its own router block : one for *Customer 2* and one for *LAN*. The routers blocks are chosen by routers using heuristics such as role aggregation and optimization of the address space. Due to space limitation, these cannot be explained in this paper. The addresses of the router blocks are flooded through the network and are therefore known by all the routers. On the other hand, the subnet block allocations are only known by the router in charge of them. More details can be found in [6].

4 Evaluation

In order to validate our assumptions, we have written a simulator [6] that allows us to evaluate the performance of our mechanism. Some results of our simulations are shown in this section.

Simulations are run on a 110-routers-topology coming from a real ISP network. Subnets with random $DSID_s$ are uniformly associated to routers to obtain the HD-Ratio we want to observe. Since we want to evaluate the protocol on a large number of subnets, the size of blocks requested is quite small, DSID sizes are comprised between 0 and 2. We consider that we obtain a /48 from our ISPs, so SID size equals to 16. The test set consists in starting this configuration and applying every minute modifications to the subnet topology. Modifications include adding new subnets, removing subnets, reducing and enlarging the address block they need. These modifications are uniformly performed unless we obtain a larger HD-Ratio than the one we want to maintain. We consider 10,000 changes for each experiment.

The set of address blocks globally reserved is one of the information that is stored on all routers. Therefore, the number of such allocations should be as low as possible

and can be viewed as a way to measure the performance of the allocation algorithm. If there is at least one subnet managed by each router, the minimum value of this is the number of routers. In our test bed, this value should be as close as possible to 110. Figure 3 shows the evolution of this value when the HD-Ratio is changing. The upper part shows number of global address blocks for each experiment. The lower part represents the mean number of subnets per router. The first vertical line shows the 0.8 value of HD-Ratio, i.e. the ratio we want to reach without any problem. The second vertical line shows the HD value (0.95) from which our mechanism is not able to place all the subnets. The dotted horizontal line represent the number of routers (110), i.e. the optimum value we could obtain. Note that for the results obtained in these experiments, we did not have to move any already allocated address block.




Fig. 3. Evolution of the number of global address blocks reserved when HD is increased

We can see at the left of Fig. 3 some points below the “110 line”. They correspond to experiments in which some routers have no subnet to place. Next, we see an increase followed by a drop. In these experiments, there is a small number of customers allocated in each router address block. Therefore, when modifications are performed (e.g. addition of new subnet) the router address block can be full and a new address block requested. When the number of customers increases, these modifications do not fill entirely the router address blocks anymore. That is why the number of address blocks decreases and stays equals or nearly equals to 110. From 0.95 as an HD-ratio, some routers do not succeed in allocating address block for all their subnets.

5 Conclusion

In this paper, we have proposed a distributed mechanism to allocate and distribute address blocks in ISP, campus and enterprise networks. We have first listed the requirements for such a protocol and discussed the security threats and how to handle them.

The allocation based on roles permits to aggregate subnets to make rules, such as fire-wall ones, simpler and shorter. When subnets have obtained address blocks, renumbering and topology changes can be performed without introducing any important perturbation in the rest of the network. Our simulations shows that our protocol holds 0.8 as HD-Ratio without any problem. However, lots of parameters of the mechanism can still be discussed, most of all concerning allocation, even though some of them could be found in our technical report. We are currently implementing the proposed protocol in the XORP platform.

References

1. Meyer, D., Zhang, L., Fall, K.: Report from the IAB Workshop on Routing and Addressing. RFC 4984, Internet Engineering Task Force (2007)
2. Palet, J.: Provider independent (PI) IPv6 assignments for end user organisations. (RIPE Policy Proposal 2006-01)
<http://www.ripe.net/ripe/policies/proposals/2006-01.html>
3. Huston, G.: Analyzing the Internet's BGP routing table. The Internet Protocol J.14 (2001)
4. Baker, F., Lear, E., Droms, R.: Procedures for renumbering an IPv6 network without a flag day. RFC 4192, Internet Engineering Task Force (2005)
5. Mahajan, R., Wetherall, D., Anderson, T.: Understanding BGP misconfiguration. In: SIGCOMM 2002: Proceedings of the 2002 conference on Applications, technologies, architectures, and protocols for computer communications, pp. 3–16. ACM, New York (2002)
6. Leroy, D., Bonaventure, O.: A secure mechanism for address block allocation and distribution. Technical report, Universite catholique de Louvain (UCL), Belgium (work in progress, 2008), <http://inl.info.ucl.ac.be/addr-alloc-distrib>
7. Durand, A., Huitema, C.: The Host-Density Ratio for Address Assignment Efficiency: An update on the H ratio. RFC 3194, Internet Engineering Task Force (2001)
8. APNIC, ARIN, RIPE NCC: IPv6 address allocation and assignment policy. ripe-412 (2007), <http://www.ripe.net/ripe/docs/ipv6policy.html>
9. Chown, T., Ford, A., Venaas, S.: Things to think about when renumbering an IPv6 network. Internet Draft, Internet Engineering Task Force “draft-chown-v6ops-renumber-thinkabout-05” (work in progress, 2006)
10. Baccelli, E., Mase, K., Ruffino, S., Singh, S.: Address autoconfiguration for MANET: Terminology and problem statement. Internet Draft (work in progress 2007), “draft-ietf-autoconf-statement-02”, Internet Engineering Task Force
11. Chelius, G., Fleury, E., Toutain, L.: No Administration Protocol (NAP) for IPv6 router auto-configuration. Int. J. Internet Protocol Technology 1 (2005)
12. Troan, O., Droms, R.: IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6. RFC 3633, Internet Engineering Task Force (2003)
13. Crawford, M.: Router renumbering for IPv6. RFC 2894, Internet Engineering Task Force (2000)
14. Lynn, C., Kent, S., Seo, K.: X.509 Extensions for IP Addresses and AS Identifiers. RFC 3779, Internet Engineering Task Force (2004)
15. Greenberg, A., Hjalmtysson, G., Maltz, D.A., Myers, A., Rexford, J., Xie, G., Yan, H., Zhan, J., Zhang, H.: Refactoring Network Control and Management: A Case for the 4D Architecture. Technical Report CMU-CS-05-117, CMU CS (2005)
16. Farrell, S., Housley, R.: An Internet Attribute Certificate Profile for Authorization. RFC 3281, Internet Engineering Task Force (2002)