A secure role-based address allocation and distribution mechanism

Damien Leroy[*] Olivier Bonaventure
Université catholique de Louvain (UCL), Belgium
{firstname.surname}@uclouvain.be

1. MOTIVATIONS

The growth of the BGP routing tables in the default-free zone is again a concern for many network operators. So as to tackle this problem, the IRTF has chartered the Routing Research Group to propose new solutions to improve the scalability of the Internet routing architecture. A key reason for the growth of the BGP routing tables 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 change the addresses of all its network when it changes from its upstream provider. For this reason, most corporate networks insist on obtaining PI addresses. Combined with the growth of multihoming, 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 the corporate networks must be renumbered each time it changes from provider and several studies have shown this to be painful with both IPv4 and IPv6.

During the last years, extensions to the Internet architecture have been proposed. Several of these solutions rely on the separation between the two distinct roles of IP addresses : **identifier** and **locator**. An identifier is an address used to identify (the applications running on) one endsystem. An endsystem usually has one identifier. It can be a cryptographic identifier (such as with HIP) or an IP address (such as with LISP) obtained from PI space. A locator indicates the attachment point of an endsystem or a router. A router and a multihomed host have multiple locators assigned to them. A mapping mechanism is used to derive one locator from an identifier. When an endsystem moves or its upstream provider changes, its locator(s) change(s) but its identifier remains the same.

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 autoconfiguration. 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.

In this paper, we propose a distributed mechanism that allows IP addresses used as locators to be automatically distributed and assigned to routers inside the network. The routers then are responsible for the suballocation of these locators to their locally connected endsystems and customers. What is called “subnet” in this paper is either a LAN in our own network or a customer network.

The zero-configuration protocol proposed in [1] tackles to similar problems. Nevertheless, their solution is limited to small networks and does not have any security consideration. The *autoconf* working group at IETF is also offering close solutions than ours [2]. However, their work is focussed on ad-hoc networks, thus they do not have the same hypothesis and objectives.

The next sections describe our solution in more details. First with a short description of the mechanism and next with the result of our simulations. This abstract ends with a description of what is our current and further work on this subject.

2. SECURE ROLE-BASED LOCATOR DISTRIBUTION

Our locator distribution mechanism is targetted for edge networks as well as ISP networks that need to provide loca-

[*] 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).

*CoNEXT’*07, December 10-13, 2007, New York, NY, U.S.A.
Copyright 2007 ACM 978-1-59593-770-4/ 07/ 0012 ...$5.00.

tors to their subnets. We encode locators as IPv6 addresses but there is no hindrance to apply the mechanism to another format. A full description of the protocol can be found in our technical paper[1]. The objective of the mechanism we propose is to dynamically assign and securely distribute locators in an entire network. We assume that each router is configured with an X.509 certificate indicating that it belongs to the network. In association with the corresponding public key, it will also be used to sign the locators attributed to client networks. Furthermore, we allow a network to divide its locator block in different roles (e.g. a sub-block reserved for servers, another one for customers, another one for loopback addresses, . . .). These roles are very important because they will allow the locators to be assigned in a way that simplifies the configuration of packet filters on the routers.

We consider that a locator is composed of three parts. The first part is the **prefix** given by the upstream ISP. These bits are common for all hosts/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 in our network and are called Subnet ID (**SID**). We will consider two parts in this SID: the attributed 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 delegated to the subnet.

Fig. 1 represents a typical environment where our mechanism could be used. In this figure, arrows represent the direction of locator allocation. Our mechanism is composed of three main parts: One or several prefixes are obtained from our upstream ISP by border routers; Our subnets needing a locator block ask for it at border routers; Core routers negotiate automatically which parts of the obtained address block have to be attributed to subnets.

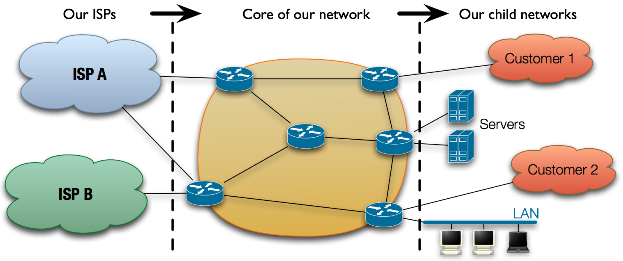

Figure 1: The protocol applies on hierarchical topology

Our locator distribution protocol behaves as follows. Routers communicate hop-by-hop, connections are only made with direct neighbors after authentication. The messages received by a router will be, if necessary, flooded to neighbors so that each router in the core receives the information. Global prefixes negotiated with ISPs are flooded to all core routers. These will append ASIDs to prefixes in order to obtain locators to deliver to subnets.

Each router in charge of one or several authenticated subnets starts off with choosing among free SIDs an address block where it can place all of them. Next, this address block is advertised through the network and the router waits a short while for a potential collision notification. If a collision appears, the procedure is restarted. Otherwise, all the routers retain this reservation and its initiator attributes locators to all of its subnets.

3. EVALUATION

In order to validate our assumptions, we have written a simulator that allows to evaluate the performances of the locator distribution protocol. We are also working on including the protocol in XORP. More evaluations are shown in our technical paper.

Simulations are run on a 110-routers-topology coming from a true ISP network. Child networks with random size are uniformally associated to routers to obtain 0.8 as an HD-Ratio[2]. So we have around 4,500 client networks. The tests consist of starting this configuration and applying regular modifications to the subnet topology. Modifications include client adding, removing, reducing and enlarging. We monitor 10,000 changes for each experiment.

Let us begin with the protocol overhead. First, we can consider the handshake messages between neighbors, i.e. the authenticated hello exchange and the keepalives. These messages never exceed 10 messages per minute with one hop. Next, let us consider the messages sent by the core routers when one of their subnets is modified. Since a router makes the reservation of an address block for all its clients, only one message is sent at startup by router having clients. Next modifications do not generate lots of messages since most changes can be done within the reserved blocks of the router. We have measured a mean rate of 0.03 messages received per router and per modification, i.e. 3 messages on each link for 100 modifications.

4. CONCLUSION

In this paper, we have proposed and briefly evaluated a secure role-based locator distribution mechanism. It will allow to automatically and securely distribute IP addresses used as locators to all routers in a large IP network. We are currently working on implementing our solution on the XORP platform to evaluate it in real networks.

5. REFERENCES

[1] G. Chelius, E. Fleury, and L. Toutain, “No Administration Protocol (NAP) for IPv6 router auto-configuration,” *Int. J. Internet Protocol Technology*, vol. 1, no. 2, 2005.

[2] E. Baccelli, K. Mase, S. Ruffino, and S. Singh, “Address autoconfiguration for MANET: Terminology and problem statement,” Internet Engineering Task Force, Internet Draft “draft-ietf-autoconf-statement-01”, Aug. 2007.

[1] Available at `http://inl.info.ucl.ac.be/protodist`

[2] HD-Ratio is defined in RFC 3194. 0.8 is the threshold defined by regional registries from which additional address allocation is justify (see ripe-388).