

Network Working Group  
Request for Comments: 3041  
Category: Standards Track

T. Narten  
IBM  
R. Draves  
Microsoft Research  
January 2001

## Privacy Extensions for Stateless Address Autoconfiguration in IPv6

### Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

### Copyright Notice

Copyright (C) The Internet Society (2001). All Rights Reserved.

### Abstract

Nodes use IPv6 stateless address autoconfiguration to generate addresses without the necessity of a Dynamic Host Configuration Protocol (DHCP) server. Addresses are formed by combining network prefixes with an interface identifier. On interfaces that contain embedded IEEE Identifiers, the interface identifier is typically derived from it. On other interface types, the interface identifier is generated through other means, for example, via random number generation. This document describes an extension to IPv6 stateless address autoconfiguration for interfaces whose interface identifier is derived from an IEEE identifier. Use of the extension causes nodes to generate global-scope addresses from interface identifiers that change over time, even in cases where the interface contains an embedded IEEE identifier. Changing the interface identifier (and the global-scope addresses generated from it) over time makes it more difficult for eavesdroppers and other information collectors to identify when different addresses used in different transactions actually correspond to the same node.

## Table of Contents

|      |                                                      |    |
|------|------------------------------------------------------|----|
| 1.   | Introduction.....                                    | 2  |
| 2.   | Background.....                                      | 3  |
| 2.1. | Extended Use of the Same Identifier.....             | 3  |
| 2.2. | Address Usage in IPv4 Today.....                     | 4  |
| 2.3. | The Concern With IPv6 Addresses.....                 | 5  |
| 2.4. | Possible Approaches.....                             | 6  |
| 3.   | Protocol Description.....                            | 7  |
| 3.1. | Assumptions.....                                     | 8  |
| 3.2. | Generation Of Randomized Interface Identifiers.....  | 9  |
| 3.3. | Generating Temporary Addresses.....                  | 10 |
| 3.4. | Expiration of Temporary Addresses.....               | 11 |
| 3.5. | Regeneration of Randomized Interface Identifiers.... | 12 |
| 4.   | Implications of Changing Interface Identifiers.....  | 13 |
| 5.   | Defined Constants.....                               | 14 |
| 6.   | Future Work.....                                     | 14 |
| 7.   | Security Considerations.....                         | 15 |
| 8.   | Acknowledgments.....                                 | 15 |
| 9.   | References.....                                      | 15 |
| 10.  | Authors' Addresses.....                              | 16 |
| 11.  | Full Copyright Statement.....                        | 17 |

## 1. Introduction

Stateless address autoconfiguration [ADDRCONF] defines how an IPv6 node generates addresses without the need for a DHCP server. Some types of network interfaces come with an embedded IEEE Identifier (i.e., a link-layer MAC address), and in those cases stateless address autoconfiguration uses the IEEE identifier to generate a 64-bit interface identifier [ADDRARCH]. By design, the interface identifier is likely to be globally unique when generated in this fashion. The interface identifier is in turn appended to a prefix to form a 128-bit IPv6 address.

All nodes combine interface identifiers (whether derived from an IEEE identifier or generated through some other technique) with the reserved link-local prefix to generate link-local addresses for their attached interfaces. Additional addresses, including site-local and global-scope addresses, are then created by combining prefixes advertised in Router Advertisements via Neighbor Discovery [DISCOVERY] with the interface identifier.

Not all nodes and interfaces contain IEEE identifiers. In such cases, an interface identifier is generated through some other means (e.g., at random), and the resultant interface identifier is not globally unique and may also change over time. The focus of this document is on addresses derived from IEEE identifiers, as the

concern being addressed exists only in those cases where the interface identifier is globally unique and non-changing. The rest of this document assumes that IEEE identifiers are being used, but the techniques described may also apply to interfaces with other types of globally unique and/or persistent identifiers.

This document discusses concerns associated with the embedding of non-changing interface identifiers within IPv6 addresses and describes extensions to stateless address autoconfiguration that can help mitigate those concerns for individual users and in environments where such concerns are significant. Section 2 provides background information on the issue. Section 3 describes a procedure for generating alternate interface identifiers and global-scope addresses. Section 4 discusses implications of changing interface identifiers.

## 2. Background

This section discusses the problem in more detail, provides context for evaluating the significance of the concerns in specific environments and makes comparisons with existing practices.

### 2.1. Extended Use of the Same Identifier

The use of a non-changing interface identifier to form addresses is a specific instance of the more general case where a constant identifier is reused over an extended period of time and in multiple independent activities. Anytime the same identifier is used in multiple contexts, it becomes possible for that identifier to be used to correlate seemingly unrelated activity. For example, a network sniffer placed strategically on a link across which all traffic to/from a particular host crosses could keep track of which destinations a node communicated with and at what times. Such information can in some cases be used to infer things, such as what hours an employee was active, when someone is at home, etc.

One of the requirements for correlating seemingly unrelated activities is the use (and reuse) of an identifier that is recognizable over time within different contexts. IP addresses provide one obvious example, but there are more. Many nodes also have DNS names associated with their addresses, in which case the DNS name serves as a similar identifier. Although the DNS name associated with an address is more work to obtain (it may require a DNS query) the information is often readily available. In such cases, changing the address on a machine over time would do little to address the concerns raised in this document, unless the DNS name is changed as well (see Section 4).

Web browsers and servers typically exchange "cookies" with each other [COOKIES]. Cookies allow web servers to correlate a current activity with a previous activity. One common usage is to send back targeted advertising to a user by using the cookie supplied by the browser to identify what earlier queries had been made (e.g., for what type of information). Based on the earlier queries, advertisements can be targeted to match the (assumed) interests of the end-user.

The use of a constant identifier within an address is of special concern because addresses are a fundamental requirement of communication and cannot easily be hidden from eavesdroppers and other parties. Even when higher layers encrypt their payloads, addresses in packet headers appear in the clear. Consequently, if a mobile host (e.g., laptop) accessed the network from several different locations, an eavesdropper might be able to track the movement of that mobile host from place to place, even if the upper layer payloads were encrypted [SERIALNUM].

## 2.2. Address Usage in IPv4 Today

Addresses used in today's Internet are often non-changing in practice for extended periods of time, especially in non-home environments (e.g., corporations, campuses, etc.). In such sites, addresses are assigned statically and typically change infrequently. Over the last few years, sites have begun moving away from static allocation to dynamic allocation via DHCP [DHCP]. In theory, the address a client gets via DHCP can change over time, but in practice servers often return the same address to the same client (unless addresses are in such short supply that they are reused immediately by a different node when they become free). Thus, even within sites using DHCP, clients frequently end up using the same address for weeks to months at a time.

For home users accessing the Internet over dialup lines, the situation is generally different. Such users do not have permanent connections and are often assigned temporary addresses each time they connect to their ISP (e.g., AOL). Consequently, the addresses they use change frequently over time and are shared among a number of different users. Thus, an address does not reliably identify a particular device over time spans of more than a few minutes.

A more interesting case concerns always-on connections (e.g., cable modems, ISDN, DSL, etc.) that result in a home site using the same address for extended periods of time. This is a scenario that is just starting to become common in IPv4 and promises to become more of a concern as always-on internet connectivity becomes widely available. Although it might appear that changing an address regularly in such environments would be desirable to lessen privacy

concerns, it should be noted that the network prefix portion of an address also serves as a constant identifier. All nodes at (say) a home, would have the same network prefix, which identifies the topological location of those nodes. This has implications for privacy, though not at the same granularity as the concern that this document addresses. Specifically, all nodes within a home would be grouped together for the purposes of collecting information. This issue is difficult to address, because the routing prefix part of an address contains topology information and cannot contain arbitrary values.

Finally, it should be noted that nodes that need a (non-changing) DNS name generally have static addresses assigned to them to simplify the configuration of DNS servers. Although Dynamic DNS [DDNS] can be used to update the DNS dynamically, it is not yet widely deployed. In addition, changing an address but keeping the same DNS name does not really address the underlying concern, since the DNS name becomes a non-changing identifier. Servers generally require a DNS name (so clients can connect to them), and clients often do as well (e.g., some servers refuse to speak to a client whose address cannot be mapped into a DNS name that also maps back into the same address). Section 4 describes one approach to this issue.

### 2.3. The Concern With IPv6 Addresses

The division of IPv6 addresses into distinct topology and interface identifier portions raises an issue new to IPv6 in that a fixed portion of an IPv6 address (i.e., the interface identifier) can contain an identifier that remains constant even when the topology portion of an address changes (e.g., as the result of connecting to a different part of the Internet). In IPv4, when an address changes, the entire address (including the local part of the address) usually changes. It is this new issue that this document addresses.

If addresses are generated from an interface identifier, a home user's address could contain an interface identifier that remains the same from one dialup session to the next, even if the rest of the address changes. The way PPP is used today, however, PPP servers typically unilaterally inform the client what address they are to use (i.e., the client doesn't generate one on its own). This practice, if continued in IPv6, would avoid the concerns that are the focus of this document.

A more troubling case concerns mobile devices (e.g., laptops, PDAs, etc.) that move topologically within the Internet. Whenever they move (in the absence of technology such as mobile IP [MOBILEIP]), they form new addresses for their current topological point of attachment. This is typified today by the "road warrior" who has

Internet connectivity both at home and at the office. While the node's address changes as it moves, however, the interface identifier contained within the address remains the same (when derived from an IEEE Identifier). In such cases, the interface identifier can be used to track the movement and usage of a particular machine [SERIALNUM]. For example, a server that logs usage information together with a source addresses, is also recording the interface identifier since it is embedded within an address. Consequently, any data-mining technique that correlates activity based on addresses could easily be extended to do the same using the interface identifier. This is of particular concern with the expected proliferation of next-generation network-connected devices (e.g., PDAs, cell phones, etc.) in which large numbers of devices are in practice associated with individual users (i.e., not shared). Thus, the interface identifier embedded within an address could be used to track activities of an individual, even as they move topologically within the internet.

In summary, IPv6 addresses on a given interface generated via Stateless Autoconfiguration contain the same interface identifier, regardless of where within the Internet the device connects. This facilitates the tracking of individual devices (and thus potentially users). The purpose of this document is to define mechanisms that eliminate this issue, in those situations where it is a concern.

#### 2.4. Possible Approaches

One way to avoid some of the problems discussed above is to use DHCP for obtaining addresses. With DHCP, the DHCP server could arrange to hand out addresses that change over time.

Another approach, compatible with the stateless address autoconfiguration architecture, would be to change the interface id portion of an address over time and generate new addresses from the interface identifier for some address scopes. Changing the interface identifier can make it more difficult to look at the IP addresses in independent transactions and identify which ones actually correspond to the same node, both in the case where the routing prefix portion of an address changes and when it does not.

Many machines function as both clients and servers. In such cases, the machine would need a DNS name for its use as a server. Whether the address stays fixed or changes has little privacy implication since the DNS name remains constant and serves as a constant identifier. When acting as a client (e.g., initiating communication), however, such a machine may want to vary the addresses it uses. In such environments, one may need multiple addresses: a "public" (i.e., non-secret) server address, registered

in the DNS, that is used to accept incoming connection requests from other machines, and a "temporary" address used to shield the identity of the client when it initiates communication. These two cases are roughly analogous to telephone numbers and caller ID, where a user may list their telephone number in the public phone book, but disable the display of its number via caller ID when initiating calls.

To make it difficult to make educated guesses as to whether two different interface identifiers belong to the same node, the algorithm for generating alternate identifiers must include input that has an unpredictable component from the perspective of the outside entities that are collecting information. Picking identifiers from a pseudo-random sequence suffices, so long as the specific sequence cannot be determined by an outsider examining information that is readily available or easily determinable (e.g., by examining packet contents). This document proposes the generation of a pseudo-random sequence of interface identifiers via an MD5 hash. Periodically, the next interface identifier in the sequence is generated, a new set of temporary addresses is created, and the previous temporary addresses are deprecated to discourage their further use. The precise pseudo-random sequence depends on both a random component and the globally unique interface identifier (when available), to increase the likelihood that different nodes generate different sequences.

### 3. Protocol Description

The goal of this section is to define procedures that:

- 1) Do not result in any changes to the basic behavior of addresses generated via stateless address autoconfiguration [ADDRCONF].
- 2) Create additional global-scope addresses based on a random interface identifier for use with global scope addresses. Such addresses would be used to initiate outgoing sessions. These "random" or temporary addresses would be used for a short period of time (hours to days) and would then be deprecated. Deprecated address can continue to be used for already established connections, but are not used to initiate new connections. New temporary addresses are generated periodically to replace temporary addresses that expire, with the exact time between address generation a matter of local policy.
- 3) Produce a sequence of temporary global-scope addresses from a sequence of interface identifiers that appear to be random in the sense that it is difficult for an outside observer to predict a

future address (or identifier) based on a current one and it is difficult to determine previous addresses (or identifiers) knowing only the present one.

- 4) Generate a set of addresses from the same (randomized) interface identifier, one address for each prefix for which a global address has been generated via stateless address autoconfiguration. Using the same interface identifier to generate a set of temporary addresses reduces the number of IP multicast groups a host must join. Nodes join the solicited-node multicast address for each unicast address they support, and solicited-node addresses are dependent only on the low-order bits of the corresponding address. This decision was made to address the concern that a node that joins a large number of multicast groups may be required to put its interface into promiscuous mode, resulting in possible reduced performance.

### 3.1. Assumptions

The following algorithm assumes that each interface maintains an associated randomized interface identifier. When temporary addresses are generated, the current value of the associated randomized interface identifier is used. The actual value of the identifier changes over time as described below, but the same identifier can be used to generate more than one temporary address.

The algorithm also assumes that for a given temporary address, an implementation can determine the corresponding public address from which it was generated. When a temporary address is deprecated, a new temporary address is generated. The specific valid and preferred lifetimes for the new address are dependent on the corresponding lifetime values in the public address.

Finally, this document assumes that when a node initiates outgoing communication, temporary addresses can be given preference over public addresses. This can mean that all connections initiated by the node use temporary addresses by default, or that applications individually indicate whether they prefer to use temporary or public addresses. Giving preference to temporary address is consistent with on-going work that addresses the topic of source-address selection in the more general case [ADDR\_SELECT]. An implementation may make it a policy that it does not select a public address in the event that no temporary address is available (e.g., if generation of a useable temporary address fails).



### 3.2. Generation Of Randomized Interface Identifiers.

We describe two approaches for the maintenance of the randomized interface identifier. The first assumes the presence of stable storage that can be used to record state history for use as input into the next iteration of the algorithm across system restarts. A second approach addresses the case where stable storage is unavailable and there is a need to generate randomized interface identifiers without previous state.

#### 3.2.1. When Stable Storage Is Present

The following algorithm assumes the presence of a 64-bit "history value" that is used as input in generating a randomized interface identifier. The very first time the system boots (i.e., out-of-the-box), a random value should be generated using techniques that help ensure the initial value is hard to guess [RANDOM]. Whenever a new interface identifier is generated, a value generated by the computation is saved in the history value for the next iteration of the algorithm.

A randomized interface identifier is created as follows:

- 1) Take the history value from the previous iteration of this algorithm (or a random value if there is no previous value) and append to it the interface identifier generated as described in [ADDRARCH].
- 2) Compute the MD5 message digest [MD5] over the quantity created in the previous step.
- 3) Take the left-most 64-bits of the MD5 digest and set bit 6 (the left-most bit is numbered 0) to zero. This creates an interface identifier with the universal/local bit indicating local significance only. Save the generated identifier as the associated randomized interface identifier.
- 4) Take the rightmost 64-bits of the MD5 digest computed in step 2) and save them in stable storage as the history value to be used in the next iteration of the algorithm.

MD5 was chosen for convenience, and because its particular properties were adequate to produce the desired level of randomization. IPv6 nodes are already required to implement MD5 as part of IPsec [IPSEC], thus the code will already be present on IPv6 machines.

In theory, generating successive randomized interface identifiers using a history scheme as above has no advantages over generating them at random. In practice, however, generating truly random numbers can be tricky. Use of a history value is intended to avoid the particular scenario where two nodes generate the same randomized

interface identifier, both detect the situation via DAD, but then proceed to generate identical randomized interface identifiers via the same (flawed) random number generation algorithm. The above algorithm avoids this problem by having the interface identifier (which will often be globally unique) used in the calculation that generates subsequent randomized interface identifiers. Thus, if two nodes happen to generate the same randomized interface identifier, they should generate different ones on the followup attempt.

### 3.2.2. In The Absence of Stable Storage

In the absence of stable storage, no history value will be available across system restarts to generate a pseudo-random sequence of interface identifiers. Consequently, the initial history value used above will need to be generated at random. A number of techniques might be appropriate. Consult [RANDOM] for suggestions on good sources for obtaining random numbers. Note that even though machines may not have stable storage for storing a history value, they will in many cases have configuration information that differs from one machine to another (e.g., user identity, security keys, serial numbers, etc.). One approach to generating a random initial history value in such cases is to use the configuration information to generate some data bits (which may remain constant for the life of the machine, but will vary from one machine to another), append some random data and compute the MD5 digest as before.

### 3.3. Generating Temporary Addresses

[ADDRCONF] describes the steps for generating a link-local address when an interface becomes enabled as well as the steps for generating addresses for other scopes. This document extends [ADDRCONF] as follows. When processing a Router Advertisement with a Prefix Information option carrying a global-scope prefix for the purposes of address autoconfiguration (i.e., the A bit is set), perform the following steps:

- 1) Process the Prefix Information Option as defined in [ADDRCONF], either creating a public address or adjusting the lifetimes of existing addresses, both public and temporary. When adjusting the lifetimes of an existing temporary address, only lower the lifetimes. Implementations must not increase the lifetimes of an existing temporary address when processing a Prefix Information Option.
- 2) When a new public address is created as described in [ADDRCONF] (because the prefix advertised does not match the prefix of any address already assigned to the interface, and the Valid Lifetime in the option is not zero), also create a new temporary address.

- 3) When creating a temporary address, the lifetime values are derived from the corresponding public address as follows:
  - Its Valid Lifetime is the lower of the Valid Lifetime of the public address or TEMP\_VALID\_LIFETIME.
  - Its Preferred Lifetime is the lower of the Preferred Lifetime of the public address or TEMP\_PREFERRED\_LIFETIME - DESYNC\_FACTOR.

A temporary address is created only if this calculated Preferred Lifetime is greater than REGEN\_ADVANCE time units. In particular, an implementation must not create a temporary address with a zero Preferred Lifetime.

- 4) New temporary addresses are created by appending the interface's current randomized interface identifier to the prefix that was used to generate the corresponding public address. If by chance the new temporary address is the same as an address already assigned to the interface, generate a new randomized interface identifier and repeat this step.
- 5) Perform duplicate address detection (DAD) on the generated temporary address. If DAD indicates the address is already in use, generate a new randomized interface identifier as described in Section 3.2 above, and repeat the previous steps as appropriate up to 5 times. If after 5 consecutive attempts no non-unique address was generated, log a system error and give up attempting to generate temporary addresses for that interface.

Note: because multiple temporary addresses are generated from the same associated randomized interface identifier, there is little benefit in running DAD on every temporary address. This document recommends that DAD be run on the first address generated from a given randomized identifier, but that DAD be skipped on all subsequent addresses generated from the same randomized interface identifier.

#### 3.4. Expiration of Temporary Addresses

When a temporary address becomes deprecated, a new one should be generated. This is done by repeating the actions described in Section 3.3, starting at step 3). Note that, except for the transient period when a temporary address is being regenerated, in normal operation at most one temporary address corresponding to a public address should be in a non-deprecated state at any given time. Note that if a temporary address becomes deprecated as result of processing a Prefix Information Option with a zero Preferred Lifetime, then a new temporary address must not be generated. The Prefix Information Option will also deprecate the corresponding public address.

To insure that a preferred temporary address is always available, a new temporary address should be regenerated slightly before its predecessor is deprecated. This is to allow sufficient time to avoid race conditions in the case where generating a new temporary address is not instantaneous, such as when duplicate address detection must be run. It is recommended that an implementation start the address regeneration process `REGEN_ADVANCE` time units before a temporary address would actually be deprecated.

As an optional optimization, an implementation may wish to remove a deprecated temporary address that is not in use by applications or upper-layers. For TCP connections, such information is available in control blocks. For UDP-based applications, it may be the case that only the applications have knowledge about what addresses are actually in use. Consequently, one may need to use heuristics in deciding when an address is no longer in use (e.g., the default `TEMP_VALID_LIFETIME` suggested above).

### 3.5. Regeneration of Randomized Interface Identifiers

The frequency at which temporary addresses should change depends on how a device is being used (e.g., how frequently it initiates new communication) and the concerns of the end user. The most egregious privacy concerns appear to involve addresses used for long periods of time (weeks to months to years). The more frequently an address changes, the less feasible collecting or coordinating information keyed on interface identifiers becomes. Moreover, the cost of collecting information and attempting to correlate it based on interface identifiers will only be justified if enough addresses contain non-changing identifiers to make it worthwhile. Thus, having large numbers of clients change their address on a daily or weekly basis is likely to be sufficient to alleviate most privacy concerns.

There are also client costs associated with having a large number of addresses associated with a node (e.g., in doing address lookups, the need to join many multicast groups, etc.). Thus, changing addresses frequently (e.g., every few minutes) may have performance implications.

This document recommends that implementations generate new temporary addresses on a periodic basis. This can be achieved automatically by generating a new randomized interface identifier at least once every  $(\text{TEMP\_PREFERRED\_LIFETIME} - \text{REGEN\_ADVANCE} - \text{DESYNC\_FACTOR})$  time units. As described above, generating a new temporary address `REGEN_ADVANCE` time units before a temporary address becomes deprecated produces addresses with a preferred lifetime no larger than `TEMP_PREFERRED_LIFETIME`. The value `DESYNC_FACTOR` is a random value (different for each client) that ensures that clients don't

synchronize with each other and generate new addresses at exactly the same time. When the preferred lifetime expires, a new temporary address is generated using the new randomized interface identifier.

Because the precise frequency at which it is appropriate to generate new addresses varies from one environment to another, implementations should provide end users with the ability to change the frequency at which addresses are regenerated. The default value is given in TEMP\_PREFERRED\_LIFETIME and is one day. In addition, the exact time at which to invalidate a temporary address depends on how applications are used by end users. Thus the default value given of one week (TEMP\_VALID\_LIFETIME) may not be appropriate in all environments. Implementations should provide end users with the ability to override both of these default values.

Finally, when an interface connects to a new link, a new randomized interface identifier should be generated immediately together with a new set of temporary addresses. If a device moves from one ethernet to another, generating a new set of temporary addresses from a different randomized interface identifier ensures that the device uses different randomized interface identifiers for the temporary addresses associated with the two links, making it more difficult to correlate addresses from the two different links as being from the same node.

#### 4. Implications of Changing Interface Identifiers

The IPv6 addressing architecture goes to some lengths to ensure that interface identifiers are likely to be globally unique where easy to do so. During the IPng discussions of the GSE proposal [GSE], it was felt that keeping interface identifiers globally unique in practice might prove useful to future transport protocols. Usage of the algorithms in this document may complicate providing such a future flexibility.

The desires of protecting individual privacy vs. the desire to effectively maintain and debug a network can conflict with each other. Having clients use addresses that change over time will make it more difficult to track down and isolate operational problems. For example, when looking at packet traces, it could become more difficult to determine whether one is seeing behavior caused by a single errant machine, or by a number of them.

Some servers refuse to grant access to clients for which no DNS name exists. That is, they perform a DNS PTR query to determine the DNS name, and may then also perform an A query on the returned name to verify that the returned DNS name maps back into the address being used. Consequently, clients not properly registered in the DNS may

be unable to access some services. As noted earlier, however, a node's DNS name (if non-changing) serves as a constant identifier. The wide deployment of the extension described in this document could challenge the practice of inverse-DNS-based "authentication," which has little validity, though it is widely implemented. In order to meet server challenges, nodes could register temporary addresses in the DNS using random names (for example a string version of the random address itself).

Use of the extensions defined in this document may complicate debugging and other operational troubleshooting activities. Consequently, it may be site policy that temporary addresses should not be used. Implementations may provide a method for a trusted administrator to override the use of temporary addresses.

## 5. Defined Constants

Constants defined in this document include:

TEMP\_VALID\_LIFETIME -- Default value: 1 week. Users should be able to override the default value.  
TEMP\_PREFERRED\_LIFETIME -- Default value: 1 day. Users should be able to override the default value.  
REGEN\_ADVANCE -- 5 seconds  
MAX\_DESYNC\_FACTOR -- 10 minutes. Upper bound on DESYNC\_FACTOR.  
DESYNC\_FACTOR -- A random value within the range 0 - MAX\_DESYNC\_FACTOR. It is computed once at system start (rather than each time it is used) and must never be greater than (TEMP\_VALID\_LIFETIME - REGEN\_ADVANCE).

## 6. Future Work

An implementation might want to keep track of which addresses are being used by upper layers so as to be able to remove a deprecated temporary address from internal data structures once no upper layer protocols are using it (but not before). This is in contrast to current approaches where addresses are removed from an interface when they become invalid [ADDRCONF], independent of whether or not upper layer protocols are still using them. For TCP connections, such information is available in control blocks. For UDP-based applications, it may be the case that only the applications have knowledge about what addresses are actually in use. Consequently, an implementation generally will need to use heuristics in deciding when an address is no longer in use (e.g., as is suggested in Section 3.4).

The determination as to whether to use public vs. temporary addresses can in some cases only be made by an application. For example, some applications may always want to use temporary addresses, while others may want to use them only in some circumstances or not at all. Suitable API extensions will likely need to be developed to enable individual applications to indicate with sufficient granularity their needs with regards to the use of temporary addresses.

## 7. Security Considerations

The motivation for this document stems from privacy concerns for individuals. This document does not appear to add any security issues beyond those already associated with stateless address autoconfiguration [ADDRCONF].

## 8. Acknowledgments

The authors would like to acknowledge the contributions of the IPNGWG working group and, in particular, Matt Crawford, Steve Deering and Allison Mankin for their detailed comments.

## 9. References

- [ADDRARCH] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 2373, July 1998.
- [ADDRCONF] Thomson, S. and T. Narten, "IPv6 Address Autoconfiguration", RFC 2462, December 1998.
- [ADDR\_SELECT] Draves, R. "Default Address Selection for IPv6", Work in Progress.
- [COOKIES] Kristol, D. and L. Montulli, "HTTP State Management Mechanism", RFC 2965, October 2000.
- [DHCP] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.
- [DDNS] Vixie, R., Thomson, S., Rekhter, Y. and J. Bound, "Dynamic Updates in the Domain Name System (DNS UPDATE)", RFC 2136, April 1997.
- [DISCOVERY] Narten, T., Nordmark, E. and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", RFC 2461, December 1998.

- [GSE] Crawford, et al., "Separating Identifiers and Locators in Addresses: An Analysis of the GSE Proposal for IPv6", Work in Progress.
- [IPSEC] Kent, S., Atkinson, R., "Security Architecture for the Internet Protocol", RFC 2401, November 1998.
- [MD5] Rivest, R., "The MD5 Message-Digest Algorithm", RFC 1321, April 1992.
- [MOBILEIP] Perkins, C., "IP Mobility Support", RFC 2002, October 1996.
- [RANDOM] Eastlake 3rd, D., Crocker S. and J. Schiller, "Randomness Recommendations for Security", RFC 1750, December 1994.
- [SERIALNUM] Moore, K., "Privacy Considerations for the Use of Hardware Serial Numbers in End-to-End Network Protocols", Work in Progress.

#### 10. Authors' Addresses

Thomas Narten  
IBM Corporation  
P.O. Box 12195  
Research Triangle Park, NC 27709-2195  
USA

Phone: +1 919 254 7798  
EMail: narten@raleigh.ibm.com

Richard Draves  
Microsoft Research  
One Microsoft Way  
Redmond, WA 98052

Phone: +1 425 936 2268  
EMail: richdr@microsoft.com



11. Full Copyright Statement

Copyright (C) The Internet Society (2001). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.