

INTERNET DRAFT
< [draft-ietf-ipv6-deprecate-site-local-00.txt](#) >
August 16, 2003
Expires February 16, 2003

C. Huitema
Microsoft
B. Carpenter
IBM

Deprecating Site Local Addresses

Status of this memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

This document is an Internet-Draft. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

Abstract

This document describes the issues surrounding the use of IPv6 site-local unicast addresses in their original form, and formally deprecates them. This deprecation does not prevent their continued use until a replacement has been standardized and implemented.

1 Introduction

For some time, the IPv6 working group has been debating a set of issues surrounding the use of "site local" addresses. In its meeting in March 2003, the group reached a measure of agreement that these issues were serious enough to warrant a replacement of site local addresses in their original form. Although the consensus was far from unanimous, the working group decided in its meeting in July 2003 to document these issues and the consequent decision to deprecate IPv6 site-local unicast addresses.

Site-local addresses are defined in the IPv6 addressing architecture [[RFC3513](#)], especially in [section 2.5.6](#).

The remainder of this document describes the adverse effects of site-local addresses according to the above definition, and formally deprecates them.

Carpenter, Huitema.

[Page 1]

INTERNET DRAFT Deprecating Site Local Addresses

August 16, 2003

Companion documents will describe the goals of a replacement solution [Hain/Templin] and specify a replacement solution [Hinden/Haberman]. However, the formal deprecation allows existing usage of site-local addresses to continue until the replacement is standardized and implemented.

[2](#) Adverse effects of site local addresses

Discussions in the IPv6 working group outlined several defects of the current site local addressing scope. These defects fall in two broad categories: ambiguity of addresses, and fuzzy definition of sites.

As currently defined, site local addresses are ambiguous: an address such as FEC0::1 can be present in multiple sites, and the address itself does not contain any indication of the site to which it belongs. This creates pain for developers of applications, for the designers of routers and for the network managers. This pain is compounded by the fuzzy nature of the site concept. We will develop the specific nature of this pain in the following section.

[2.1](#) Developer pain

Early feedback from developers indicates that site local addresses are hard to use correctly in an application. This is particularly true for multi-homed hosts, which can be simultaneously connected to multiple sites, and for mobile hosts, which can be successively connected to multiple sites.

Applications would learn or remember that the address of some correspondent was "FEC0::1234:5678:9ABC", they would try to feed the address in a socket address structure and issue a connect, and the call will fail because they did not fill up the "site identifier" variable, as in "FEC0::1234:5678:9ABC&1". The problem is compounded by the fact that the site identifier varies with the host instantiation, e.g. sometimes &1 and sometimes &2, and thus that the host identifier cannot be remembered in memory, or learned from a name server.

In short, the developer pain is caused by the ambiguity of site local addresses. Since site-local addresses are ambiguous,

application developers have to manage the "site identifiers" that qualify the addresses of the hosts. This management of identifiers has proven hard to understand by developers, and also hard to execute by those developers who understand the concept.

[2.2](#) Manager pain, leaks

The management of IPv6 site local addresses is in many ways similar to the management of [RFC 1918](#) [[RFC1918](#)] addresses in some IPv4 networks. In theory, the private addresses defined in [RFC 1918](#) should only be used locally, and should never appear in the

Carpenter, Huitema.

[Page 2]

INTERNET DRAFT Deprecating Site Local Addresses

August 16, 2003

Internet. In practice, these addresses "leak". The conjunction of leaks and ambiguity ends up causing management problems.

Names and literal addresses of "private" hosts leak in mail messages, web pages, or files. Private addresses end up being used as source or destination of TCP requests or UDP messages, for example in DNS or trace-route requests, causing the request to fail, or the response to arrive at unsuspecting hosts. Private addresses also end up being used as targets of reverse lookup requests in the DNS, uselessly overloading the DNS infrastructure.

The leakage issue is largely unavoidable. While some applications are intrinsically scoped (eg. RA, ND), most applications have no concept of scope, and no way of expressing scope. As a result, "stuff leaks across the borders". Since the addresses are ambiguous, the network managers cannot easily find out "who did it". Leaks are thus hard to fix, resulting in a lot of frustration.

[2.3](#) Router pain, routing tables

The ambiguity of site local addresses also creates problems for the routers. In theory, site local addresses are only used within a contiguous site, and all routers in that site can treat them as if they were not ambiguous. In practice, problem occurs when sites are disjoint, or when routers have to handle several sites.

In theory, sites should never be disjoint. In practice, if site local addressing is used throughout a large network, some elements of the site will not be directly connected. This will create a demand to route the site-local packets across some intermediate network. In practice, this leads to an extensive use of tunneling techniques, or the use of multi-sited routers, or both.

Ambiguous addresses have fairly obvious consequences on multi-sited

routers. In classic router architecture, the exit interface is a direct function of the destination address, as specified by a single routing table. However, if a router is connected to multiple sites, the routing of site local packets depends on the interface on which the packet arrived. Interfaces have to be associated to sites, and the routing entries for the site local addresses are site-dependent. The route management software and the routing protocols have to account for the site boundaries.

In multi-homed routers, such as for example site border routers, the routing process should be complemented by a filtering process, to guarantee that packets sourced with a site local address never leave the site. This filtering process will in turn interact with the routing of packets, as it may cause the drop of packets sent to a global address, even if that global address happen to belong to the target site.

In summary, the ambiguity of site local addresses makes them hard to

Carpenter, Huitema.

[Page 3]

INTERNET DRAFT Deprecating Site Local Addresses

August 16, 2003

manage in multi-sited routers, while the requirement to support disjoint sites creates a demand for such routers.

[2.4](#) Site is an ill-defined concept

The current definition of scopes follows an idealized "concentric scope" model. Hosts are supposed to be attached to a link, which belongs to a site, which belongs to the Internet. Packets could be sent to the same link, the same site, or outside that site. However, experts have been arguing about the definition of sites for years and have reached no sort of consensus. That suggests that there is in fact no consensus to be reached.

Apart from link-local, scope boundaries are ill-defined. What is a site? Is the whole of a corporate network a site, or are sites limited to single geographic locations? Many networks today are split between an internal area and an outside facing "DMZ", separated by a firewall. Servers in the DMZ are supposedly accessible by both the internal hosts and external hosts on the Internet. Does the DMZ belong to the same site as the internal host?

Depending on whom we ask, the definition of the site scope varies. It may map security boundaries, reachability boundaries, routing boundaries, QOS boundaries, administrative boundaries, funding boundaries, some other kinds of boundaries, or a combination. It is very unclear that a single scope could satisfy all these requirements.

There are some well known and important scope-breaking phenomena, such as intermittently connected networks, mobile nodes, mobile networks, inter-domain VPNs, hosted networks, network merges and splits, etc. Specifically, this means that scope *cannot* be mapped into concentric circles such as a naive link/local/global model. Scopes overlap and extend into one another. The scope relationship between two hosts may even be different for different protocols.

In summary, the current concept of site is naive, and does not map operational requirements.

3 Development of a better alternative

The previous section reviewed the arguments against site-local addresses. Obviously, site locals also have some benefits, without which they would have been removed from the specification long ago. The perceived benefits of site local are that they are simple, stable, and private [Hain/Templin]. However, it appears that these benefits can be also obtained with an alternative architecture, for example [Hinden/Haberman], in which addresses are not ambiguous and do not have a simple explicit scope.

Having non ambiguous address solves a large part of the developers' pain, as it removes the need to manage site identifiers. The

Carpenter, Huitema.

[Page 4]

INTERNET DRAFT Deprecating Site Local Addresses

August 16, 2003

application can use the addresses as if they were regular global addresses, and the stack will be able to use standard techniques to discover which interface should be used. Some level of pain will remain, as these addresses will not always be reachable; however, applications can deal with the un-reachability issues by trying connections at a different time, or with a different address. Speculatively, a more sophisticated scope mechanism might be introduced at a later date.

Having non ambiguous addresses will not eliminate the leaks that cause management pain. However, since the addresses are not ambiguous, debugging these leaks will be much simpler.

Having non ambiguous addresses will solve a large part of the router issues: since addresses are not ambiguous, routers will be able to use standard routing techniques, and will not need different routing tables for each interface. Some of the pain will remain at border routers, which will need to filter packets from some ranges of source addresses; this is however a fairly common function.

Avoiding the explicit declaration of scope will remove the issues linked to the ambiguity of the site concept. Non-reachability can be obtained by using "firewalls" where appropriate. The firewall rules can explicitly accommodate various network configurations, by accepting or refusing traffic to and from ranges of the new non-ambiguous addresses.

One question remains, anycast addressing. Anycast addresses are ambiguous by construction, since they refer by definition to any host that has been assigned a given anycast address. Link-local or global anycast addresses can be "baked in the code". Further study is required on the need for anycast addresses with scope between link-local and global.

[4](#) Deprecation

This document formally deprecates the IPv6 link-local unicast prefix defined in [[RFC3513](#)], i.e. 1111111011 binary or FEC0::/10. The special behavior of this prefix MUST no longer be supported in new implementations. The prefix MUST NOT be reassigned for other use except by a future IETF standards action. Future versions of the addressing architecture [[RFC3513](#)] will include this information.

However, router implementations SHOULD be configured to prevent routing of this prefix by default.

Existing implementations and deployments MAY continue to use this prefix.

[5](#) Security Considerations

The link-local unicast prefix allows for some blocking action in

Carpenter, Huitema.

[Page 5]

INTERNET DRAFT Deprecating Site Local Addresses August 16, 2003

firewall rules and address selection rules, which are commonly viewed as a security feature since they prevent packets crossing administrative boundaries. However, such blocking rules can be configured for any prefix, including the expected future replacement for the site-local prefix. Thus the deprecation of the site-local prefix does not endanger security.

[6](#) IANA Considerations

IANA is specifically requested not to reassign the 1111111011 binary or FEC0::/10 prefix unless requested to do so by a future IETF standards action.

7 Copyright

The following copyright notice is copied from [RFC 2026](#) [Bradner, 1996], [Section 10.4](#), and describes the applicable copyright for this document.

Copyright (C) The Internet Society August 13, 2003. 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 assignees.

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.

8 Intellectual Property

The following notice is copied from [RFC 2026](#) [Bradner, 1996], [Section 10.4](#), and describes the position of the IETF concerning intellectual property claims made against this document.

Carpenter, Huitema.

[Page 6]

INTERNET DRAFT Deprecating Site Local Addresses

August 16, 2003

The IETF takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use other technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Information on the

IETF's procedures with respect to rights in standards-track and standards-related documentation can be found in [BCP-11](#). Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF Secretariat.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this standard. Please address the information to the IETF Executive Director.

[9](#) Acknowledgements

[10](#) References

Normative References

[RFC3513] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", [RFC 3513](#), April 2003

Informative references

[RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G. J. and E. Lear, "Address Allocation for Private Internets", [RFC 1918](#), February 1996

[Hain/Templin] Hain, T. and F. Templin, "Addressing Requirements for Local Communications within Sites", work in progress.

[Hinden/Haberman] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", work in progress.

[11](#) Authors' Addresses

Christian Huitema
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052-6399
USA
Email: huitema@microsoft.com

Brian Carpenter
IBM Corporation

Carpenter, Huitema.

[Page 7]

Sauemerstrasse 4
8803 Rueschlikon
Switzerland
Email: brc@zurich.ibm.com

