

Global Routing Operations  
Internet-Draft  
Expires: May 1, 2004

D. Plonka  
University of Wisconsin  
November 2003

**Embedding Globally Routable Internet Addresses Considered Harmful**  
**draft-ietf-grow-embed-addr-00**

Status of this Memo

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

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/1id-abstracts.txt>.

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

This Internet-Draft will expire on May 1, 2004.

Copyright Notice

Copyright (C) The Internet Society (2003). All Rights Reserved.

Abstract

Vendors of consumer electronics and network gear have produced and sold hundreds of thousands of Internet hosts with globally routable Internet Protocol addresses embedded within their products' firmware. These products are now in operation world-wide and primarily include, but are not necessarily limited to, low-cost routers and middleboxes for personal or residential use.

This "hard-coding" of globally routable IP addresses as identifiers within the host's firmware presents significant problems to the operation of the Internet and to the management of its address space.

This document means to clarify best current practices in the Internet community. It denounces the practice of embedding references to



unique, globally routable IP addresses in Internet hosts, describes some of the resulting problems, and considers selected alternatives. It is also intended to remind the Internet community of the ephemeral nature of unique, globally routable IP addresses and that the assignment and use of IP addresses as identifiers is temporary and therefore should not be used in fixed configurations.

## Revision History

The following is the revision history of this document since "-00":

\$Log: [draft-ietf-grow-embed-addr.xml](#),v \$

Revision 1.11 2003/12/02 22:28:04 plonka

renamed from [draft-plonka-embed-addr](#) to [draft-ietf-grow-embed-addr](#)

integrated suggestions from Paul Barford

reordered references to match the text

added quote from [RFC2101](#) re: use of IPv4 addresses as identifiers  
as mentioned by Brian Carpenter

Revision 1.10 2003/11/03 17:06:54 plonka

added background information in [appendix](#)

[Revision 1.9](#) 2003/11/03 16:39:30 plonka

various updates based on input from Mike O'Connor:

- indicated that DNS server(s) should be configurable
- clarified DNS round-robin behavior
- clarified "unsolicited traffic" by saying "IP traffic"

added revision history and [appendix A](#)

Figure 1



## **1. Introduction**

Internet hosts should not contain globally routable Internet Protocol addresses embedded within firmware or elsewhere as part of their default configuration influencing their run-time behavior.

Ostensibly, this practice arose as an attempt to simplify configuration of IP hosts by preloading them with IP addresses as service identifiers. Unfortunately, products that rely on such embedded IP addresses initially may appear convenient to both the product's designer and its operator or user, but this dubious benefit comes at the expense of others in the Internet community.



## 2. Problems

In a number cases, the embedding of IP addresses has caused Internet products to rely on a single central Internet service. This can result in a service outage when the aggregate workload overwhelms that service. When fixed addresses are embedded in an ever-increasing number of client IP hosts, this practice runs directly counter to the design intent of hierarchically deployed services that would otherwise be robust solutions.

The reliability, scalability, and performance of many Internet services require that the pool of users not directly access a service by IP address. Instead they rely on a level of indirection provided by the Domain Name System, [RFC 2219](#) [1]. DNS permits the service operator to reconfigure the resources for maintenance and load-balancing without the participation of the users. For instance, one common load-balancing technique employs multiple DNS records with the same name that are then rotated in a round-robin fashion in the set of answers returned by the Berkeley Internet Name Daemon (BIND) and other DNS server implementations. Upon receiving such a response to a query, resolvers typically use the first valid answer in the set, thus enabling the operator to distribute the user request load across a set of servers with discrete IP addresses that generally remain unknown to the user.

Embedding globally unique IP addresses taints the IP address blocks in which they reside, lessening the usefulness and portability of those IP address blocks and increasing the cost of operation. Unsolicited traffic may continue to be delivered to the embedded addresses, even after the IP address or block has been reassigned and no longer hosts the service for which that traffic was meant. Circa 1997, the authors of [RFC 2101](#) [3] made this observation:

Due to dynamic address allocation and increasingly frequent network renumbering, temporal uniqueness of IPv4 addresses is no longer globally guaranteed, which puts their use as identifiers into severe question.

In this way, IP address blocks containing addresses that have been embedded into the configuration of many Internet hosts become encumbered by their historical use. This may interfere with the ability of the Internet Assigned Numbers Authority (IANA) and the Internet Registry (IR) hierarchy to usefully reallocate IP address blocks. This is of particular concern as the IPv4 address space nears exhaustion. Note that, to facilitate IP address reuse, [RFC 2050](#) [2], encourages Internet Service Providers (ISPs) to treat address assignments as "loans".





Because consumers are not necessarily capable, experienced operators of Internet hosts, they are not able to be relied upon to implement a fix if and when problems arise. As such, a significant responsibility lies with the manufacturer or vendor of the Internet host to avoid embedding IP addresses.

### **3. Recommendations**

Internet host and router designers, including network product manufacturers, should not assume that their products will only be deployed on a single global Internet, that they happen to observe today. A myriad of private internets in which these products will be used will often not allow these hosts to establish end-to-end communications with arbitrary hosts on the global Internet.

Vendors should, by default, disable unnecessary features in their products. This is especially true of features that generate unsolicited IP traffic. In this way these hosts will be conservative regarding the unsolicited Internet traffic they produce. For instance, one of the most common uses of embedded IP addresses has been the hard-coding of addresses of well know public Simple Network Time Protocol (SNTP [RFC 2030](#) [4]) servers, even though only a small fraction of the users benefits from these products even having some notion of the current date and time.

Vendors should provide an operator interface for every feature that generates unsolicited IP traffic. A prime example of this that the Domain Name System resolver should have an interface enabling the operator to either explicitl set the servers of his choosing or to enable the use of a standard automated configuration protocol such as DHCP, defined by [RFC 2132](#) [5]. Within the operator interface, these features should be disabled by default so that one consequence of enabling these features is that the operator becomes aware that the feature exists. This will mean that it is more likely that the product's owner or operator can participate in problem determination and mitigation when problems arise.

Internet hosts should use the Domain Name System to determine the routable IP addresses associated with the Internet services they require. However, note that simply hard-coding DNS names rather than IP addresses is not a panacea. Entries in the domain name space are also ephemeral and can change owners for various reasons including such as acquisitions and litigation. A given vendor ought not assume that it will retain control of a given zone indefinitely.

Whenever possible, default configurations, documentation, and example configurations for Internet hosts should use Private Internet Addresses, as defined by [RFC 1918](#) [6], rather than unique, globally routable IP addresses.

Service providers and enterprise network operators should advertise the identities of suitable local services. For instance, the DHCP protocol, as defined by [RFC 2132](#) [5], enables one to configure a server to answer queries regarding available servers to clients that



ask for them. Unless the advertisement of local services is ubiquitous, designers may resort to ad hoc mechanisms that rely on central services.

Operators that provide public services on the global Internet, such as the NTP community, should deprecate the explicit advertisement of IP addresses of public services. These addresses are ephemeral. As such, their widespread citations in public service indexes interferes with the ability to reconfigure the service as necessary to address unexpected, increased workloads.

#### **4. Security Considerations**

Embedding or "hard-coding" IP addresses within a host's configuration almost always means that some sort of host-based trust model is being employed, and that the Internet host with the given address is trusted in some way. Due to the ephemeral roles of routable IP addresses, the practice of embedding them within products' firmware or default configurations presents a security risk.

An Internet host designer may be tempted to implement some sort of remote control mechanism within a product, by which its Internet host configuration can be changed without reliance on, interaction with, or even the knowledge of its operator or user. This raises security issues of its own. If such a scheme is implemented, this should be fully disclosed to the customer, operator, and user so that an informed decision can be made, in accordance with local security or privacy policy. Furthermore, the significant possibility of malicious parties exploiting such a remote control mechanism may completely negate any potential benefit of the remote control scheme.



## **5. Conclusion**

As larger numbers of homogenous hosts continue to be deployed, it is particularly important that both their designers and other members of the Internet community are diligent in assessing host implementation quality and reconfigurability. Unique, globally routable IP addresses should not be embedded within a host's fixed configuration because doing so excludes the ability to remotely influence hosts when the unsolicited IP traffic they generate causes problems for the for those operating the IP addresses to which the traffic is destined.





## **6. Acknowledgements**

Thanks go to the following folks for providing input during the preparation of this document: Paul Barford and Mike O'Connor.

## References

- [1] Hamilton, M., "Use of DNS Aliases for Network Services", [RFC 2219](#), [BCP 17](#), October 1997.
- [2] Hubbard, K., "INTERNET REGISTRY IP ALLOCATION GUIDELINES", [RFC 2050](#), [BCP 12](#), November 1996.
- [3] Carpenter, B., "IPv4 Address Behaviour Today", [RFC 2101](#), February 1997.
- [4] Mills, D., "Simple Network Time Protocol (SNTP) Version 4 for IPv4, IPv6 and OSI", [RFC 2030](#), October 1996.
- [5] Alexander, S., "DHCP Options and BOOTP Vendor Extensions", [RFC 2132](#), March 1997.
- [6] Rekhter, Y., "Address Allocation for Private Internets", [RFC 1918](#), [BCP 5](#), February 1996.

## Author's Address

David J. Plonka  
University of Wisconsin - Madison  
DoIT, room b116  
1210 W. Dayton Street  
Madison, WI 53705  
US

Phone: +1 608 265 5184  
EMail: [plonka@doit.wisc.edu](mailto:plonka@doit.wisc.edu)  
URI: <http://net.doit.wisc.edu/~plonka/>



## [Appendix A](#). Background

In June 2003, the University of Wisconsin discovered that the network product vendor named NetGear had manufactured and shipped over 700,000 routers with firmware containing a hard-coded reference to the IP address of one of the University's NTP servers: 128.105.39.11, which was also known as "ntp1.cs.wisc.edu", a public stratum-2 NTP server.

Due to that embedded fixed configuration and a bug in the implementation, the NetGear SNTP client has a failure mode in which each flawed router produces one query per second destined for the IP address 128.105.39.11, and hence produces a large-scale flood of Internet traffic from hundreds-of-thousands of legitimate source addresses and destined for the University's network resulting in significant operational problems.

These flawed routers are widely deployed throughout the global Internet and are likely to remain in use for years to come. As such, the University of Wisconsin with the cooperation of NetGear will build a new anycast time service which aims to mitigate the damage caused by the misbehavior of these flawed routers.

A technical report regarding the details of this situation is available on the world-wide-web: Flawed Routers Flood University of Wisconsin Internet Time Server [7]



## Intellectual Property Statement

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 of the 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 implementors 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.

## Full Copyright Statement

Copyright (C) The Internet Society (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.

#### Acknowledgment

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