Global Routing Operations
Internet-Draft
Expires: December 7, 2004

D. Plonka
University of Wisconsin

June 8, 2004

Embedding Globally Routable Internet Addresses Considered Harmful
draft-ietf-grow-embed-addr-02

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 December 7, 2004.

Copyright Notice

Abstract

This document means to clarify best current practices in the Internet community. Internet hosts should not contain globally routable Internet Protocol addresses embedded within firmware or elsewhere as part of their default configuration such that it influences run-time

behavior.

Revision History

RFC-EDITOR: PLEASE REMOVE REVISION HISTORY BEFORE PUBLICATION. The following is the revision history of this document

$Log: draft-ietf-grow-embed-addr.xml,v $
Revision 1.15 2004/06/08 14:16:45 plonka
revised conclusion based on input from Geoff Huston

added netgear-sntp technical report to list of informative references

Revision 1.14 2004/06/07 18:16:27 plonka
split references into normative and informative sections

Revision 1.13 2004/06/07 16:32:10 plonka
Set category to BCP.

Rewrote/resized abstract and introduction as suggested by Pekka Savola.

Improved section about using DNS names, re; hard-coding caveats, as suggested by Pekka Savola.

Encouraged use of IPv4 documentation/example prefix 192.0.2.0/24 rather than private addresses, as noted by Pekka Savola.

Mentioned IPv6 2001:DB8::/32 documentation prefix, as noted by Tom Petch.

Added note for RFC-editor requesting that revision history be removed.

Reworded various portions.

Renamed from "-00" to "-01" and updated date.

Revision 1.12 2003/12/05 15:51:23 plonka
typo fixes and updates from Michael Patton

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

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.

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.

This document 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 also reminds 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.

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 typically rely on a level of indirection provided by the Domain Name System, RFC 2219 [6]. DNS permits the service operator to reconfigure the resources for maintenance and to load-balance 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 many DNS server implementations. Upon receiving such a response to a query, resolvers typically will try the answers in order, until one succeeds, 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 well after the IP address or block has been reassigned and no longer hosts the service for which that traffic was intended. Circa 1997, the authors of RFC 2101 [5] 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.

When IP addresses are used as service identifiers in the configuration of many Internet hosts, the IP address blocks 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. Likewise, to facilitate IP address reuse, RFC 2050 [1], encourages Internet Service Providers (ISPs) to treat address assignments as "loans".

Because consumers are not necessarily experienced in the operation 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 in ways which cause the aforementioned problems.

3. Recommendations

Internet host and router designers, including network product manufacturers, should not assume that their products will be deployed and used in only a single global Internet, that they happen to observe today. A myriad of private or future internets in which these products will be used may not allow those hosts to establish end-to-end communications with arbitrary hosts on the global Internet. Since the product failure modes resulting from unknown future states cannot be fully explored, one should avoid assumptions regarding the longevity of our current 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 [7]) 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 is that the Domain Name System resolver should have an interface enabling the operator to either explicitly set the servers of his choosing or to enable the use of a standard automated configuration protocol such as DHCP, defined by RFC 2132 [8]. Within the operator interface, these features should originally be disabled so that one consequence of subsequently 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 IP addresses associated with the Internet services they require. However, 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 acquisitions and litigation. A given vendor ought not assume that anyone will retain control of a given zone indefinitely. RFC 2606 [2] defines the IANA-reserved "example.com", "example.net", and "example.org" domains for use in example configurations and documentation.

Default configurations, documentation, and example configurations for Internet hosts should use Internet addresses that reside with special blocks that have been reserved for these purposes, rather than unique, globally routable IP addresses. For IPv4, RFC 3330 [3]

states that the 192.0.2.0/24 block has been assigned for use in documentation and example code. The IPv6 global unicast address prefix 2001:DB8::/32 has been similarly reserved for documentation purposes. Private Internet Addresses, as defined by RFC 1918 [4], should not be used for such purposes.

Service providers and enterprise network operators should advertise the identities of suitable local services. For instance, the DHCP protocol, as defined by RFC 2132 [8], enables one to configure a server to answer queries for service identitifiers to clients that ask for them. When local services are available but not pervasively advertised using such common protocols, designers are more likely deploy ad hoc initialization mechanisms that unnecessarily rely on central services.

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

4. Security Considerations

Embedding or "hard-coding" IP addresses within a host's configuration often means that a 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 in that unknown parties may inadvertently be trusted.

Internet host designers 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 decisions can be made, perhaps 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. IANA Considerations

This document creates no new requirements on IANA namespaces.

6. Conclusion

When large numbers of homogenous Internet hosts are deployed, it is particularly important that both their designers and other members of the Internet community diligently assess host implementation quality and reconfigurability.

Implementors of host services should avoid any kind of use of unique globally routable IP addresses within a fixed configuration part of the service implementation. If there is a requirement for pre-configured state then care should be taken to use an appropriate service identifier and use standard resolution mechanisms to dynamically resolve the identifier into an IP address. Also, any such identifiers should be alterable in the field through a conventional command and control interface for the service.

7. Acknowledgements

The author thanks the following reviewers for their contributions to this document: Paul Barford, Geoff Huston, David Meyer, Mike O'Connor, Michael Patton, Tom Petch, and Pekka Savola.

8. References

8.1 Normative References

[1] Hubbard, K., "INTERNET REGISTRY IP ALLOCATION GUIDELINES", RFC 2050, BCP 12, November 1996.

[2] Eastlake, D., "Reserved Top Level DNS Names", RFC 2606, BCP 32, June 1999.

[3] Internet Assigned Numbers Authority, "Special-Use IPv4 Addresses", RFC 3330, September 2002.

[4] Rekhter, Y., "Address Allocation for Private Internets", RFC 1918, BCP 5, February 1996.

8.2 Informative References

[5] Carpenter, B., "IPv4 Address Behaviour Today", RFC 2101, February 1997.

[6] Hamilton, M., "Use of DNS Aliases for Network Services", RFC 2219, BCP 17, October 1997.

[7] Mills, D., "Simple Network Time Protocol (SNTP) Version 4 for IPv4, IPv6 and OSI", RFC 2030, October 1996.

[8] Alexander, S., "DHCP Options and BOOTP Vendor Extensions", RFC 2132, March 1997.

[9] Plonka, D., "Flawed Routers Flood University of Wisconsin Internet Time Server", August 2003, <http://www.cs.wisc.edu/~plonka/netgear-sntp/>.

Author's Address

David Plonka
University of Wisconsin - Madison

EMail: plonka AT doit DOT wisc DOT edu
URI: http://net.doit.wisc.edu/~plonka/

Appendix A. Background

In June 2003, the University of Wisconsin discovered that a 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 an unrelated bug in the SNMP client, the affected products occasionally exhibit 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 source addresses, 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 [9].

Intellectual Property Statement

Full Copyright Statement

Acknowledgment

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