

Embedding Globally-Routable Internet Addresses Considered Harmful

Status of This Memo

This document specifies an Internet Best Current Practices for the Internet Community, and requests discussion and suggestions for improvements. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2005).

Abstract

This document discourages the practice of embedding references to unique, globally-routable IP addresses in Internet hosts, describes some of the resulting problems, and considers selected alternatives. This document is intended to clarify best current practices in this regard.

Table of Contents

1. Introduction	2
2. Problems	2
3. Recommendations	4
3.1. Disable Unused Features	4
3.2. Provide User Interface for IP Features	4
3.3. Use Domain Names as Service Identifiers	4
3.4. Use Special-Purpose, Reserved IP Addresses When Available ..	5
3.5. Discover and Utilize Local Services	6
3.6. Avoid Mentioning the IP Addresses of Services	6
4. Security Considerations	6
5. Conclusion	7
6. Acknowledgements	7
7. References	7
Appendix A. Background	9

1. Introduction

Some vendors of consumer electronics and network gear have unfortunately chosen to embed, or "hard-code", globally-routable Internet Protocol addresses within their products' firmware. These embedded IP addresses are typically individual server IP addresses or IP subnet prefixes. Thus, they are sometimes used as service identifiers, to which unsolicited requests are directed, or as subnet identifiers, specifying sets of Internet addresses that the given product somehow treats specially.

One recent example was the embedding of the globally-routable IP address of a Network Time Protocol server in the firmware of hundreds of thousands of Internet hosts that are now in operation worldwide. The hosts are primarily, but are not necessarily, limited to low-cost routers and middleboxes for personal or residential use. In another case, IP address prefixes that had once been reserved by the Internet Assigned Numbers Authority (IANA) were embedded in a router product so that it can automatically discard packets that appear to have invalid source IP addresses.

Such "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 IP host configuration by pre-loading hosts with IP addresses. Products that rely on such embedded IP addresses initially may appear to be convenient to the product's designer and to 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; the assignment and use of IP addresses as identifiers is temporary and therefore should not be used in fixed configurations.

2. Problems

The embedding of IP addresses in products has caused an increasing number of Internet hosts 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 access a service using its IP address directly. Instead, they typically rely on a level of indirection provided by the Domain Name System, RFC 2219 [6]. When appropriately utilized, the DNS permits the service operator to reconfigure the resources for maintenance and to perform load balancing, without the participation of the users and without a requirement for configuration changes in the client hosts. For instance, one common load-balancing technique employs multiple DNS records with the same name; the set of answers that is returned is rotated in a round-robin fashion in successive queries. 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 mobility of those IP address blocks and increasing the cost of operation. Unsolicited traffic may continue to be delivered to the embedded address 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 [7] 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 embedded 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 cannot be relied upon to fix problems, if and when they arise. Therefore, a significant responsibility lies with the manufacturer or vendor of an Internet host to avoid embedding IP addresses in ways that 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 the single global Internet that they happen to observe today. A myriad of private or future internetworks in which these products will be used may not allow those hosts to establish communications with arbitrary hosts on the global Internet. Since the product failure modes resulting from an unknown future internetwork environment cannot be fully explored, one should avoid assumptions regarding the longevity of our current Internet.

The following recommendations are presented as best practice today.

3.1. Disable Unused Features

Vendors should, by default, disable unnecessary features in their products. This is especially true of features that generate unsolicited Internet 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 known public Simple Network Time Protocol (SNTP RFC 2030 [8]) servers in products. However, only a small fraction of users will benefit from these products having some notion of the current date and time.

3.2. Provide User Interface for IP Features

Vendors should provide an operator interface for every feature that generates unsolicited Internet traffic. A prime example is this: the Domain Name System resolver should have an interface enabling the operator to either explicitly set the choice of servers or enable a standard automated configuration protocol such as DHCP, defined by RFC 2132 [9]. These features should originally be disabled within the operator interface, and subsequently enabling these features should alert the operator that the feature exists. This will make it more likely that the product's owner or operator can participate in problem determination and mitigation when problems arise.

RFC 2606 [2] defines the IANA-reserved "example.com", "example.net", and "example.org" domains for use in example configurations and documentation. These are candidate examples to be used in user interface documentation.

3.3. Use Domain Names as Service Identifiers

Internet hosts should use the Domain Name System to determine the IP addresses associated with the Internet services they require.

When using domain names as service identifiers in the configurations of deployed Internet hosts, designers and vendors are encouraged to introduce service names. These names should be within a domain that they either control or are permitted to utilize by an agreement with its operator (such as for public services provided by the Internet community). This is commonly done by introducing a service-specific prefix to the domain name.

For instance, a vendor named "Example, Inc." with the domain "example.com" might configure its product to find its SNTP server by the name "sntp-server.config.example.com" or even by a name that is specific to the product and version, such as "sntp-server.v1.widget.config.example.com". Here the "config.example.com" namespace is dedicated to that vendor's product configuration, with subdomains introduced as deemed necessary. Such special-purpose domain names enable ongoing maintenance and reconfiguration of the services for their client hosts and can aid in the ongoing measurement of service usage throughout the product's lifetime.

An alternative to inventing vendor-specific domain naming conventions for a product's service identifiers is to utilize SRV resource records (RRs), defined by RFC 2782 [10]. SRV records are a generic type of RR that uses a service-specific prefix in combination with a base domain name. For example, an SRV-cognizant SNTP client might discover Example, Inc.'s suggested NTP server by performing an SRV-type query to lookup for "_ntp._udp.example.com".

However, note that simply hard-coding DNS name service identifiers 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. As such, developers and vendors should explore a product's potential failure modes resulting from the loss of administrative control of a given domain for whatever reason.

3.4. Use Special-Purpose, Reserved IP Addresses When Available

Default configurations, documentation, and example configurations for Internet hosts should use Internet addresses that reside within 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 RFC 3849 [4]. Private Internet Addresses, as defined by RFC 1918 [5], should not be used for such purposes.

3.5. Discover and Utilize Local Services

Service providers and enterprise network operators should advertise the identities of suitable local services, such as NTP. Very often these services exist, but the advertisement and automated configuration of their use is missing. For instance, the DHCP protocol, as defined by RFC 2132 [9], enables one to configure a server to answer client queries for service identifiers. When local services, including NTP, are available but not pervasively advertised using such common protocols, designers are more likely to deploy ad hoc initialization mechanisms that unnecessarily rely on central services.

3.6. Avoid Mentioning the IP Addresses of Services

Operators who provide public services on the global Internet, such as those in 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 when necessary to address unexpected, increased traffic and the aforementioned problems.

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 globally-routable IP addresses, the practice of embedding them within products' firmware or default configurations presents a security risk in which unknown parties may be trusted inadvertently.

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, its presence should be fully disclosed to the customer, operator, and user, so that an informed decision 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. Therefore, remote control mechanisms should be disabled by default, to be subsequently enabled and disabled by the user.

5. Conclusion

When large numbers of homogeneous 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 to use standard mechanisms for dynamically resolving 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.

6. 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. Harald Alvestrand, Spencer Dawkins, Ted Hardie, David Kessens, and Thomas Narten provided valuable feedback during AD and IESG review.

7. References

7.1. Normative References

- [1] Hubbard, K., Kusters, M., Conrad, D., Karrenberg, D., and J. Postel, "Internet Registry IP Allocation Guidelines", BCP 12, RFC 2050, November 1996.
- [2] Eastlake 3rd, D. and A. Panitz, "Reserved Top Level DNS Names", BCP 32, RFC 2606, June 1999.
- [3] IANA, "Special-Use IPv4 Addresses", RFC 3330, September 2002.
- [4] Huston, G., Lord, A., and P. Smith, "IPv6 Address Prefix Reserved for Documentation", RFC 3849, July 2004.
- [5] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.

7.2. Informative References

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

- [7] Carpenter, B., Crowcroft, J., and Y. Rekhter, "IPv4 Address Behaviour Today", RFC 2101, February 1997.
- [8] Mills, D., "Simple Network Time Protocol (SNTP) Version 4 for IPv4, IPv6 and OSI", RFC 2030, October 1996.
- [9] Alexander, S. and R. Droms, "DHCP Options and BOOTP Vendor Extensions", RFC 2132, March 1997.
- [10] Gulbrandsen, A., Vixie, P., and L. Esibov, "A DNS RR for specifying the location of services (DNS SRV)", RFC 2782, February 2000.
- [11] Plonka, D., "Flawed Routers Flood University of Wisconsin Internet Time Server", August 2003.
<http://www.cs.wisc.edu/~plonka/netgear-sntp/>

Appendix A. Background

In May 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 SNTP 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 that 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 [11].

Author's Address

David Plonka
University of Wisconsin - Madison

Email: plonka@doit.wisc.edu
URI: <http://net.doit.wisc.edu/~plonka/>

Full Copyright Statement

Copyright (C) The Internet Society (2005).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM 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.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights 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; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat 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 on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

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

