Internet Draft
draft-ietf-dnsop-dontpublish-unreachable-00.txt
Valid for six months
Category: Best Current Practice

Philip Hazel
University of Cambridge

September 2001

IP Addresses that should never appear in the public DNS

Status of this Memo

Abstract

This document specifies an Internet Best Current Practice for the Internet Community. It prohibits the appearance of private IP addresses in publicly visible DNS records. It also prohibits the appearance of public addresses, or indirect references to them, when the service implied by the address or reference is inaccessible from the public Internet. Specifying the second prohibition is more difficult because inaccessibility may arise from many causes, some possibly legitimate.

1. Introduction

The increasing use of firewalls, NAT boxes, and similar technology has resulted in the fragmentation of the Internet into regions whose boundaries do not allow general connectivity. There are two primary reasons for this:

(1) The perceived shortage of IPv4 addresses has caused increasing use of private IP network addresses such as 10.0.0.0/8 on LANs. A number of private address ranges are designated in [RFC 1918]. Hosts using private addresses that wish to communicate with the public Internet must do so via an address translation mechanism such as a NAT box. This allows a host with a private address to send packets to public Internet hosts, and to receive replies. However, unsolicited incoming packets cannot reach these hosts from outside their own private network.

(2) Increasing security concerns have caused many sites to install firewalls or to implement restrictions in their boundary routers in order to lock out certain kinds of connection to their hosts, even when the hosts are using public Internet addresses, though in many cases firewalls also provide NAT functionality.

Thus, there are two classes of host which some or all types of unexpected incoming packet from the public Internet cannot reach.

A number of instances have been observed where IP addresses that are not accessible from the public Internet have nevertheless been inserted into resource records in the public DNS. This document seeks to prohibit such behaviour.

The keywords "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119].

The phrase "address record" means an A record or an AAAA record, or any other kind of name-to-address record that may come into use.

2. Private network addresses

Examples of [RFC 1918] private host addresses are 10.0.0.1 and 172.16.42.53. Packets cannot be routed to such addresses from the public Internet. [RFC 1918] explains this in section 3, from where this paragraph is taken:

> Because private addresses have no global meaning, routing information about private networks shall not be propagated on inter-enterprise links, and packets with private source or destination addresses should not be forwarded across such links. Routers in networks not using private address space, especially those of Internet service providers, are expected to be configured to reject (filter out) routing information about private networks.

In section 5 of [RFC 1918] there is already a prohibition of the appearance of private addresses in publicly visible DNS records. However, the wording is merely "should not". This document makes a stronger statement:

Public DNS zones MUST NOT contain [RFC 1918] private addresses in any resource records.

Because the same private addresses are in use in many different organizations, they are ambiguous. The appearance of private addresses in the DNS could therefore lead to unpredictable and unwanted behaviour.

3. Public network addresses that are blocked

The situation with public network addresses is more complicated. For example, a host with a public address that is behind a firewall may be accessible for SSH sessions, but not for SMTP sessions. That is, the blocking may apply only to certain ports. A publicly visible address record is therefore required to give access to those ports that are accessible, and there can be no blanket prohibition.

However, for some protocols and services, additional DNS records are defined that reference hosts' address records. These are the MX record for SMTP, and the SRV record for other services. The existence of such indirect records advertises the availability of the relevant service.

Public DNS zones MUST NOT contain MX or SRV records that point to hosts for which the relevant services are not accessible from the public Internet. In other words, if a DNS resource record that yields an IP address is visible to some part of the Internet, the IP address yielded must be reachable by the protocol(s) implied by the resource record type from the parts of the Internet where the record is visible.

4. Why publishing public but unreachable addresses is bad

A host that tries to connect to an unreachable address (or port) may not receive an immediate rejection; in many cases the connection will only fail after a timeout expires. The wasted effort ties up resources on the calling host and the network, possibly for some considerable time (SMTP timeouts are of the order of minutes). It also causes a gratuitous slowing down of the application.

Furthermore, in the case of dial-up, ISDN, or other kinds of usage-based charged network connection, the wasted network resources may cost real money.

5. Loopback addresses

The loopback addresses (127.0.0.1 for IPv4 and ::1 for IPv6) are another form of private address. There has been a practice of including them in DNS zones for two entirely different reasons.

5.1 The name "localhost"

Some hostmasters include records of this type in their zones:

```
  localhost.some.domain.example.  A  127.0.0.1
```

The reason for doing this is so that other hosts in the domain that use the DNS for all their name resolution can make use of the unqualified name "localhost". This works because DNS resolvers normally add the local enclosing domain to unqualified names.

DNS zones MAY make use of this technique for the name "localhost" only, if it is required in their environment, but SHOULD avoid it if possible.

5.3 DNS "black lists"

There is an increasingly popular practice of creating "black lists" of misbehaving hosts (for example, open mail relays) in the DNS. The first of these was the "Realtime Blackhole List" (RBL). Such lists make use of addresses in the 127.0.0.0/8 network in DNS address records to give information about listed hosts (which are looked up via their inverted IP addresses).

Such records are in specific "black list" domains, and are well understood not to be invitations to attempt connections to the addresses they publish.

Hostmasters MAY continue to make use of this technique.

5.4 Other uses of loopback networks

Apart from the exceptions mentioned in 5.2 and 5.3, the loopback addresses MUST NOT appear in address records in the public DNS.

5.5 References to loopback addresses

When address records that contain loopback addresses do exist, hostmasters MUST NOT create indirect records (MX or SRV) that reference them.

6. Alternative techniques

6.1 Splitting DNS zones

A site that is using private addresses may well want to use DNS lookups for address resolution on its hosts. The lazy way approach is simply to put the data into the public DNS zone. Because this can cause problems for external hosts, this MUST NOT be done.

One approach that is commonly taken is to run a so-called "split DNS". Two different authoritative servers are created: one containing all the zone data is accessible only from within the private network. External DNS queries are directed to the second server, which contains a filtered version of the zone, without the private addresses.

6.2 SMTP servers behind firewalls

The complication of a split DNS is not normally needed if it is only SMTP traffic that is being blocked to a public address on a host behind a firewall. Public MX records must always point to publicly accessible hosts. Setting up MX records like this:

```
  plc.example.   MX   5   mail.plc.example.
                 MX  10   public.plc.example.
```

where both hosts have public IP addresses, but the first is blocked at the firewall, MUST NOT be done. Only the publicly accessible host must be used:

```
  plc.example.   MX  10   public.plc.example.
```

If a split DNS is in use, the host public.plc.example can use the internal version to route the mail onwards. However, most MTAs have configuration facilities to allow for explicit routing of mail, without the use of the DNS.

6.3 Specification of no SMTP service

MX records that point to host names whose address records specify the loopback address have been seen in the DNS. This seems to be a misguided attempt to specify "no SMTP service for this domain".

If such a facility is required, it should instead be done by arranging for the hosts in question to return

```
  554 No SMTP service here
```

to all SMTP connections.

7. Security Considerations

This document is not known to create new security issues in the DNS, mail agents, etc. In some sense, it may reduce security exposure by insisting that a site's inappropriate internal data not be exposed.

8. IANA Considerations

No IANA actions are required by this document.

9. Acknowledgements

Randy Bush read an early draft of this document and suggested several improvements.

10. Author's Address

Philip Hazel
University of Cambridge Computing Service
New Museums Site, Pembroke Street
Cambridge CB2 3QG, England

Phone: + 44 1223 334714
Email: ph10@cam.ac.uk

11. References

[RFC 1918] Rekhter, Y. et al "Address allocation for Private Internets", BCP 5, RFC 1918, February 1996.

[RFC 2119] Bradner, S."Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

Full Copyright Statement

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.