Internet Draft
draft-ietf-dnsop-dontpublish-unreachable-01.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 has two themes. Firstly, it reinforces the prohibition in [RFC 1918] about the appearance of private IP addresses in publicly visible DNS records. Secondly, the document discusses the problems that can be caused by 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 a blanket prohibition in the second case is difficult because inaccessibility may arise from many causes, some possibly legitimate. Instead, the document points out some of the problems that can arise, and suggests that other means of achieving the desired effects should be used wherever possible.

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 such private address ranges are designated in [RFC 1918], and others may be also assigned by IANA.

[Note: For example, there's 169.254/16, which is mentioned in draft-ietf-zeroconf-ipv4-linklocal-04.txt, but since that's still a draft, I can't cite it.]

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 never accessible from the public Internet have nevertheless been inserted into resource records in the public DNS. This document seeks to prohibit such behaviour in the case of truly private addresses, and to discourage it in the case of public, but unreachable, addresses.

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.

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. Consider this set of entries:

```
@       IN      MX  10  smtp
smtp    IN      A       10.1.2.3
smtp    IN      A       1.2.3.4
```

Zones set up in this way have been seen, and some administrators apparently believe this is useful, because it allows mail on their local network to be delivered straight to the internal server (the one with address 10.1.2.3). However, it all breaks down when a host on a foreign network that is also using the address 10.1.2.3 attempts to send mail to the domain.

In section 5 of [RFC 1918] there is a prohibition of the appearance of private addresses in publicly visible DNS records. It says:

> If an enterprise uses the private address space, or a mix of private and public address spaces, then DNS clients outside of the enterprise should not see addresses in the private address space used by the enterprise, since these addresses would be ambiguous.

The wording "should not" is not a very strong prohibition, considering the interworking problems that ignoring it can cause. Therefore, this document makes a stronger statement:

Public DNS zones MUST NOT contain [RFC 1918] addresses, or any other addresses designated by IANA as private, in any resource records.

3. Public network addresses that are inacessible

The situation with public network addresses is more complicated because the Internet cannot in general be cleanly divided into

"public" and "private" parts in this case. Examples of situations where the division is fuzzy are:

(1) 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.

(2) A host with a public address may make certain services available only to specific client hosts, for example, those in partner enterprises.

(3) A host might respond to incoming packets only if the client host is using IPsec.

When a host is providing any service at all over the public Internet, a publicly visible address record is of course required to give access to the host.

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.

If these services are always inaccessible over the public Internet, it is bad practice to include the MX or SRV records in public DNS zones, for the following reason:

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

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

Public DNS zones SHOULD NOT contain MX or SRV records that point to hosts for which the relevant services are never accessible over the public Internet. In other words, if there is no host that is able to make use of the service using the public Internet, the service SHOULD NOT be publicly advertised.

4. 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.

4.1 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.

4.2 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.

DNS zones MAY continue to make use of this technique.

4.3 Other uses of loopback networks

Apart from the exceptions mentioned in 4.1 and 4.2 above, the loopback addresses MUST NOT appear in address records in the public DNS.

4.4 References to loopback addresses

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

5. Alternative techniques

5.1 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, as in the example shown in section 2 above. 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.

5.2 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. 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, SHOULD NOT be done. Only the publicly accessible host should 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 need to use a DNS lookup.

5.3 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.

6. 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.

7. IANA Considerations

No IANA actions are required by this document.

8. Acknowledgements

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

Draft 01 has benefitted from comments made by Daniel Senie, John Schnizlein, Robert Elz, Bert Hubert, and Stuart Cheshire.

9. Author's Address

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

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

10. 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.

11. Changes made during development of this document

This section is provided for the convenients of those tracking the document. It will be removed from the final draft.

11.1 Changes made to the -00 version

. While leaving the MUSTs in for truly private addresses, I've tried to be more "educational" about the case of public addresses that are inaccessible, and backed down to SHOULD in those cases.

. I've pointed out the lack of a clear-cut public/private boundary, and tried to make the case for not advertising unavailable services without being so probititive in the wording. This includes using "never accessible" instead of "not accessible".

. Changed "hostmaster" to "zone" in a couple of cases.

. Included an example of bad MX practice with an [RFC 1918] address.

. Noted that [RFC 1918] is not the only list of private addresses.

. General tidying of the wording and rearrangement of the material.

. The Post Office changed our postcode!

Full Copyright Statement

Acknowledgement

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