Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: 2 September 2025

M. Blanchet
Viagenie

1 March 2025

Domain Name System in Mostly Isolated Networks
draft-many-tiptop-dns-isolated-networks-00

Abstract

This document lists operational methods to enable local DNS name resolving on an isolated network, where that network have intermittent reachability to Internet and/or have very long delays, such as deep space networks, disabling the real-time query and response flow to the authoritative name servers on Internet.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

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."

This Internet-Draft will expire on 2 September 2025.

Copyright Notice

Table of Contents

1. Introduction

Deep space communications involve long delays (e.g. Earth to Mars is 4-20 minutes) and intermittent communications, because of orbital dynamics. [I-D.many-deepspace-ip-assessment] discusses the use of the whole IP stack in this context. Domain name requests and response over long delays generate timeouts and when there is no reachability to the DNS server, requests will not be answered. Therefore, on celestial bodies IP networks, a local DNS infrastructure with all the needed names and values stored locally is needed. Moreover, to keep the same DNS root and the current DNSSEC trust chain, all keys necessary for validation should also be stored locally. This document describes the different ways to accomplish this.

While this document uses deep space as the base use case, it applies to other "mostly" isolated networks. Mostly isolated means that most of the time the network is isolated, but there are times where it is not isolated and then may receive zone transfers or other means to populate or update its name caches. In case of deep space, the delays for those transfers is significant and the transport mechanisms are more limited, as discussed in [I-D.many-deepspace-ip-assessment].

The requirements and characteristics for this document use case are:

* domains under the unique DNS root[RFC2826]

* ability to sometimes reach the mostly isolated name servers to update their data cache

* most of the time, inability to do live DNS queries to the Internet DNS infrastructure

* multiple network and DNS operators may exist on the isolated network, each managing their own namespace

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Possible Approaches

This section presents various approaches that should meet the requirements set in the previous section. These approaches use the [RFC8806] approach for root zones, but augment it for the whole needed name hierarchy.

All approaches share similar naming infrastructure on the target isolated network:

* One or more authoritative name server.

* One or more resolvers using the above authoritative servers (or additional servers) in their hints/cache file.

* hosts using the above resolvers.

* DNSSEC verifying resolvers have the root trust anchor[trust-anchor] in their configuration.

2.1. Pre-walk of all needed names

If one assumes that all names that will be used on the isolated network are known in advance, then queries walking the tree from the root down to the final name of all needed names can be done on Internet and the responses saved in a file, together with the appropriate DNSSEC records (TBD: should we list those: aka RRSIG, DS, DNSKEY). The records should contain values that are relevant to the isolated network. For example, an IP address record such as an A or AAAA record should resolve to an IP address relevant and reachable on the target isolated network.

The resulting file containing all the records is uploaded to the authoritative name servers on the isolated network.

This method somewhat mimics the hosts.txt file used before the DNS was created.

The authoritative name servers should serve the root zone and all required domain tree records underneath as found above.

If a name used on the isolated network by the hosts or applications is not in the uploaded file served by the local name servers, then the request will leak and will timeout since the request will not reach the Internet DNS infrastructure.

If all needed DNSSEC material is not fully uploaded, then DNSSEC validation will fail.

A method for syncing and updating all the updated records to the isolated network should be put in place, at the appropriate frequency. It could be done using zone tranfer mechanism if TCP/IP reachability is possible but other file transfer mechanisms may also be used.

Some DNS records have values containing other names, such as the SRV and CNAME records. The referenced names should also be "walked".

This setup somewhat assumes that there is a single operator for the DNS authoritative infrastructure on the target isolated network.

2.2. Pre-fetch of all zones in the needed name hierarchy

If one assumes that the name hierarchy is known for all needed names used on the isolated network and if the operator of the DNS infrastructure on the isolated network has access to all the zones of the hierarchy, then these zones are saved. They may need to be modified so that the NS glue records point to the appropriate local authoritative name servers. These zones are then uploaded to the authoritative name servers on the isolated network.

The authoritative name servers should serve the root zone and all zones as discussed above.

This approach have less risk of missing a name since all names under the hierarchy are uploaded. However, if the zones are too big compared to the transfer capacity to the isolated network, then this solution is not appropriate. Moreover, it may be possible that most of the names in the uploaded zones will not be used, therefore it is a possible waste of resources (bandwidth, memory/cpu on server, ...).

Therefore, careful consideration on the chosen hierarchy, specially the top-level domain, is relevant. Given deep space relative limited use networks, it would make sense to dedicate some top-level domain or subdomain for its needs. However, it is possible to remove all the non-needed recoard from the zones before uploading them to the isolated network DNS infrastructure, but then if some names are missing in this removal, the same issues from the previous approach appear.

If all needed DNSSEC material is not fully uploaded, then DNSSEC validation will fail.

A method for syncing and updating all the updated records to the isolated network should be put in place, at the appropriate frequency. It could be done using zone tranfer mechanism if TCP/IP reachability is possible but other file transfer mechanisms may also be used.

In the context of multiple operators on the target network, each one may do this process independently for its own zones, without having to rely on another party.

2.3. Special zone

Instead of fetching a whole zone containing a lot of non useful records, the manager of that zone creates a special version of the zone containing only the useful records and sign it. It is then sent to the isolated network DNS infrastructure. This approach is a combination of the previous approaches, but require careful management of the two versions of the zone. In terms of deployment and operations, it has the same properties as the zone pre-fetch approach.

3. Zone Transfer Conideration

If DNS zone transfer is possible over the link between the Internet and the isolated network, then incremental zone transfer (aka IXFR) might be advised to minimize the use of the bandwidth and also minimize the data merge on the target DNS server.

If DNS zone transfer is not possible or not optimal, than various file transfer mechanisms such as FTP, ssh, git, rsync may be used.

4. DNSSEC Considerations

Zones are signed at various frequencies based on the operator policies. If a signature on a record has expired, then DNSSEC validation will fail. Therefore, the frequency of uploading updated records should be higher than the frequency of the signing of the uploaded zones.

Similarly, the key lifetimes, including the root zone anchor, should be monitored to make sure that new keys are uploaded before the old ones expire.

Finally, the DNSSEC RR TTL values need to be longer than the update times.

5. Network Operations Considerations

Even with careful management, there is some probability that some applications or host on the isolated network will query names that were uploaded to the local DNS infrastructure, but refer to services or IP addresses that are not reachable from the isolated network. If the isolated network do have intermittent IP connectivity to Internet but the link is not appropriate for live queries, such as long delays in deep space, costly bandwidth or very small time window of reachability, then the network may try to route the packets to the Internet. Therefore, a default route pointing to null or other mechanisms to signal unreachability may be appropriate to be setup at the edge of the isolated network.

6. IANA Considerations

This memo includes no request to IANA.

7. Security Considerations

By expanding the use of the same Internet DNS root to space, the space IP network naming infrastructure is then secured at the same level as on Internet.

8. References

8.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

8.2. Informative References

[RFC2826] IAB, "IAB Technical Comment on the Unique DNS Root", RFC 2826, DOI 10.17487/RFC2826, May 2000, <https://www.rfc-editor.org/info/rfc2826>.

[RFC8806] Kumari, W. and P. Hoffman, "Running a Root Server Local to a Resolver", RFC 8806, DOI 10.17487/RFC8806, June 2020, <https://www.rfc-editor.org/info/rfc8806>.

[trust-anchor] IANA, "Trust Anchors and Keys", <https://www.iana.org/dnssec/files>.

[I-D.many-deepspace-ip-assessment] Blanchet, M., Huitema, C., and D. Bogdanovi, "Revisiting the Use of the IP Protocol Stack in Deep Space: Assessment and Possible Solutions", Work in Progress, Internet-Draft, draft-many-deepspace-ip-assessment-02, 10 September 2024, <https://datatracker.ietf.org/doc/html/draft-many-deepspace-ip-assessment-02>.

Acknowledgements

The idea of the pre-walk was suggested by Warren Kumari. The idea of a special zone was suggested by Mark Andrews. All errors are authors'.

Author's Address

Marc Blanchet
Viagenie
Canada
Email: marc.blanchet@viagenie.ca