

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: 4 October 2026

M. Blanchet
Viagenie
2 April 2026

Deployment and Use of the Domain Name System(DNS) in Deep Space
draft-many-tiptop-dns-02

Abstract

Deep space communications involve long delays (e.g., Earth to Mars has one-way delays 4-24 minutes) and intermittent communications, mainly because of orbital dynamics. This document lists operational methods to enable local DNS name resolving on celestial body networks such that there are no real-time query and response flow to the authoritative name servers on (Earth) 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 4 October 2026.

Copyright Notice

Copyright (c) 2026 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
1.1. Requirements Language	3
2. Use Case and Requirements	3
3. Possible Approaches	4
3.1. Pre-walk of all needed names	4
3.2. Pre-fetch of all zones in the needed name hierarchy	5
3.3. Special zone	6
3.4. Local zones	6
4. Zone Transfer Considerations	7
5. DNSSEC Considerations	7
6. Cruising Spacecraft Considerations	7
7. Network Operations Considerations	8
8. IANA Considerations	8
9. Security Considerations	8
10. References	8
10.1. Normative References	8
10.2. Informative References	8
Appendix A. Alternative Name Resolution Approaches	9
Acknowledgements	10
Author's Address	10

1. Introduction

Deep space communications involve long delays (e.g., Earth to Mars has one-way delays 4-24 minutes) and intermittent communications, mainly because of orbital dynamics. To illustrate a simple example of intermittence, communications may be between a ground station on Earth to an asset on a planet through the use of an orbiter around that planet. That orbiter may have an overpass of 20 minutes every 2 hours. Therefore, continuous connectivity between the two endpoints cannot be assumed. However, as discussed in [I-D.many-tiptop-ip-architecture], orbiters may temporarily store IP packets until the next communication window with the next hop appears. From the standpoint of the transport or upper application protocols, the one-way delay and the RTT will be stable while direct connectivity happens, and then jump to a large value while the orbiter is not reachable anymore, until it becomes reachable again. In some cases, where the asset is on a shadow side of a planet or Moon and while there is no full constellation of orbiters providing connectivity all around, the asset and Earth ground may never or rarely have a uninterrupted end2end path. [I-D.ietf-tiptop-usecase] provides more details about key characteristics of deep space communications and the use case of networking.

[I-D.many-tiptop-ip-architecture] 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.

While this document uses deep space as the base use case, it may apply to other "mostly" isolated networks, such as IoT 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.

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. Use Case and Requirements

Given all the investment in technology, protocol, services, infrastructure, scrutiny and governance in place, the ability to use the current Internet DNS Root with DNSSEC[RFC9364] even in deep space has a lot of merit compared to starting another root with another trust chain. In order to achieve that goal, which is to keep the same DNS root and the current DNSSEC trust chain in deep space, domain data and all related keys necessary for local validation should also be stored locally.

The requirements and characteristics for this document use case are:

- * inability to do real-time DNS queries to the Internet DNS infrastructure from space
- * domains used in deep space are under the same unique DNS root[RFC2826], deployed under agreed community requirements[RFC7720]
- * ability to sometimes reach the celestial body network name servers from Earth Internet to update their data cache
- * multiple network and DNS operators may exist on the celestial body network, each managing their own namespace and/or caches

- * management of DNS servers and zones are done from Earth, since no human is expected to do server administration on celestial body network
- * Celestial body network autonomous DNS infrastructure to enable local resolution and DNSSEC validation

3. 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 celestial body 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.

"Target network" used in this document means the celestial body network.

It should be noted that, with proper planning and careful consideration, these approaches are not exclusive and may be deployed over time.

3.1. Pre-walk of all needed names

If one assumes that all names that will be used on the target 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 celestial body network. For example, an IP address record such as an A or AAAA record should resolve to an IP address relevant and reachable from the target network.

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

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 celestial body network by the hosts or applications is not in the uploaded file served by the local name servers, then the request will timeout.

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. Typical DNS zone transfer mechanism may not be suitable because TCP transport is likely not possible in deep space. In this case, other file transfer mechanisms may 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 network.

3.2. 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 target network and if the operator of the DNS infrastructure on the target 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 target network.

The authoritative name servers should serve the root zone and all zones 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 target 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 record from the zones before uploading them to the target network DNS infrastructure, but then if some names are missing in this removal, the same issues from the previous approach appear. Moreover, DNSSEC validation will fail if records are removed from a signed zone.

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 target network should be put in place, at the appropriate frequency. Typical DNS zone tranfer mechanism may not be suitable because TCP transport is likely not possible in deep space. In this case, other file transfer mechanisms may 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.

3.3. 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 target 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.

3.4. Local zones

Local zones can be deployed on the target networks. However, they may not be secured by DNSSEC. Moreover, the danger of name collision over time with the Earth Internet DNS can create a big operational issue later.

4. Zone Transfer Considerations

If DNS zone transfer is possible over the link between the Internet and the target 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.

Current zone transfer mechanisms, such as AXFR, IXFR, rsync or ftp, all use TCP as transport. However, TCP is not suitable for deep space. Therefore, it is recommended to use AXFR/IXFR over profiled[I-D.many-tiptop-quic-profile] QUIC[RFC9250].

If DNS zone transfer is not possible or not optimal, than another file transfer mechanism such as HTTP over QUIC should be used.

5. 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.

6. Cruising Spacecraft Considerations

This document assumes a minimal infrastructure on a celestial body surface and vicinity network. As discussed in [I-D.ietf-tiptop-usecase], there are other instances of IP networks in deep space. One of them is a cruising spacecraft which has non-continuous connectivity to Earth using a direct point-to-point link to Earth ground stations. A cruising spacecraft typically have a local IP network to connect the various payloads and computers together. Names may be used on the onboard network, either by a simple infrastructure as described in this document, or by other simpler means such as a static file distributed on all on-board computers. When the spacecraft lands or connects to an infrastructure network such as a celestial body surface network, some names may need to be updated to use the local DNS infrastructure. The transition is not discussed in this document.

7. Network Operations Considerations

Even with careful management, there is some probability that some applications or host on the target 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 target network. If the target 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 mechanisms to signal unreachability may be appropriate to be setup at the edge of the target network.

8. IANA Considerations

This memo includes no request to IANA.

9. 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.

10. References

10.1. 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>>.
- [RFC9364] Hoffman, P., "DNS Security Extensions (DNSSEC)", BCP 237, RFC 9364, DOI 10.17487/RFC9364, February 2023, <<https://www.rfc-editor.org/info/rfc9364>>.

10.2. 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>>.

- [RFC7720] Blanchet, M. and L. Liman, "DNS Root Name Service Protocol and Deployment Requirements", BCP 40, RFC 7720, DOI 10.17487/RFC7720, December 2015, <<https://www.rfc-editor.org/info/rfc7720>>.
- [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>>.
- [RFC9250] Huitema, C., Dickinson, S., and A. Mankin, "DNS over Dedicated QUIC Connections", RFC 9250, DOI 10.17487/RFC9250, May 2022, <<https://www.rfc-editor.org/info/rfc9250>>.
- [trust-anchor]
IANA, "Trust Anchors and Keys",
<<https://www.iana.org/dnssec/files>>.
- [I-D.ietf-tiptop-usecase]
Blanchet, M., Eddy, W., and M. Eubanks, "IP in Deep Space: Key Characteristics, Use Cases and Requirements", Work in Progress, Internet-Draft, draft-ietf-tiptop-usecase-01, 21 January 2026, <<https://datatracker.ietf.org/doc/html/draft-ietf-tiptop-usecase-01>>.
- [I-D.many-tiptop-quic-profile]
Blanchet, M. and W. Eddy, "QUIC Profile for Deep Space", Work in Progress, Internet-Draft, draft-many-tiptop-quic-profile-02, 1 March 2026, <<https://datatracker.ietf.org/doc/html/draft-many-tiptop-quic-profile-02>>.
- [I-D.many-tiptop-ip-architecture]
Blanchet, M., Eddy, W., and T. Li, "An Architecture for IP in Deep Space", Work in Progress, Internet-Draft, draft-many-tiptop-ip-architecture-03, 2 March 2026, <<https://datatracker.ietf.org/doc/html/draft-many-tiptop-ip-architecture-03>>.

Appendix A. Alternative Name Resolution Approaches

Alternatives of using DNS include static distribution of name-to-address mappings and the direct use of IP addresses without names.

TIPTOP nodes MAY be provisioned with static host name mappings (e.g., "hosts.txt"-style files), enabling local resolution without network queries. Updates are distributed periodically through TBD out-of-band mechanisms between celestial networks and Earth-based infrastructure.

In small or constrained celestial networks, applications MAY use IP addresses directly, avoiding any naming system. This removes resolution dependencies entirely but reduces flexibility, human usability, error prone by typing IPv6 addresses, and abstraction from addressing changes. It will also disable the use of a large variety of DNS record types.

These approaches are simple and robust to disconnection, but do not scale to large or dynamic TIPTOP deployments, introduce update latency (for static mappings), and lack DNS capabilities such as delegation, hierarchical naming, record types, caching control, and DNSSEC.

Operators of TIPTOP deployments SHOULD consider namespace size and dynamics, update distribution frequency between celestial networks and Earth-based infrastructure, administrative boundaries, and data integrity mechanisms. Static or IP-only approaches MAY be sufficient for small, stable celestial networks, but DNS-based solutions provide greater flexibility for larger or evolving TIPTOP deployments.

Hybrid models can combine static distribution of a small set of nodes in a static mapping file such as /etc/hosts with DNS-based mechanisms. For example, nodes may be provisioned with hosts files for critical infrastructure, while DNS data is distributed as discussed in this document, with a name resolution preference set properly.

Static distribution mechanisms provide simplicity and robustness under extreme delay and intermittence, but at the cost of scalability and functionality. DNS, when adapted for deep space conditions, offers a more general and extensible solution.

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'. The following people have provided valuable feedback and comments, in no specific order: Lloyd Wood, Warren Kumari.

Author's Address

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