Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: 21 September 2025

S. Johnson, Ed.
Spacely Packets, LLC

20 March 2025

An Interplanetary DNS Model
draft-johnson-dtn-interplanetary-dns-04

Abstract

As human activity continues to spread beyond the Earth, so must the communications systems which support that activity continue to expand in scope. One such case, Internet naming, presents particular challenges when considering a multi-planet civilization. Proper operation of terrestrial DNS services and clients require constant, reasonably low latency connectivity to operate properly. These conditions are distinctly not present when considering interplanetary links; high latency and frequent disruption due to space weather events and general link availability prevail. To overcome these challenges in space networking, a delay and distruption tolerant (DTN) suite of protocols has been developed based on a store and forward mechanism, Bundle Protocol(BP). This DTN network, which optimizes to ensure bundle delivery, does not allow for end to end encapsulation of IP packets beyond terrestrial boundaries, as the latency still creates issues completing network transactions, even in the unlikely event of continuous link availability.

Table of Contents

1. Introduction

This writing will present extenstions to the DNS architecture providing a solution for isolated Internet segments, such as one would find on the Moon or Mars, having best effort interoperability with the terrestrial Internet via means of application layer gateway request parsing and transiting the DTN network to create a new request on the remote network. This is preferable to a DTN only network, as there does not exist a robust, standardized application stack based on Bundle Protocol. Logic dictates that the most likely Solar System Internet architecture deploys IP on terrestrial bodies (and in some cases, in their orbits, i.e. commercial satellite based ISPs) while deep space links generally deploy BP, connecting these isolated "islands" of IP.

2. 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.

3. IP/BP Gateway Architecture

Practically, an IP network on Mars or the Moon will behave similarly to an IP network on Earth. As such, local traffic remains local, and requires no special considerations. It is presumed that non-terrestrial IP networks will be numbered with IPv6, allowing universal (as opposed to global) uniqueness. It is further presumed that specified blocks of IPv6 addresses will be allocated for this purpose, enabling standard routing practices to determine paths to the BP network edge, if desired, and standard firewalling practices to control access. For example, a BGP speaker can announce a route to the Mars block of addresses. Said router can also run (or forward to a host running) the IP/BP gateway daemons, terminating the local IP traffic and generating BP traffic, as described below.

It is recognized that operational security policy and measures may prevent, on Earth, BGP announcement of Lunar, Martian, or other planetary IPv6 network addresses. The use of DNS naming of off-world assets in the manner described below, with those records resolving to terrestrial gateway nodes, allows the desired functionality without the potential security implications of exposing off-world network addresses or, by extension, direct IP connectivity. Thus, routes to these network addresses can be administratively prohibited terrestrially, while having no impact on the interoperability described herein.

To facilitate interoperability between isolated segments of the Internet connected via a BP network, one must consider that no actual IP packets need transverse the BP network. Instead, a https request destined for a remote IP network, for example, is completed locally by the BP/IP gateway where an application layer protocol appropriate response providing a URL and/or status message for the eventually returned resource is provided to the user/requesting application, if necessary. The request MAY be logged, MAY be assigned a unique transaction number (carried forward in a BP extension header or payload struct element), and becomes the payload of a bundle destined for the remote BP/IP gateway. Upon arrival, a new IP request, in the now local network, is generated from the data relayed in the bundle. Once the requested resource is returned, it becomes the payload of one or more bundles sent back to the requesting network, still

uniquely numbered in an extension header. This data is then published to the URL provided to the user. In this way, some (mostly non-interactive) services can be enabled in an interplanetary fashion. Some interactive processes, like ssh, may be viable in this fashion if the user is willing to tolerate the latency involved; once a ssh session is established on the remote body by the remote gateway, user i/o transit of the BP network becomes relatively simple in practice.

It is potentially useful that multicast address space be reserved for numbering of said gateways; ff0b:: being one possibility. This architecture does not preclude the growth of BP based networks on other worlds, in favor of IP only networks. Indeed, early efforts will be largely BP based, as the needs of robotic missions can largely be met by the per-mission customized applications created for that purpose. Human exploration and colonization at scale, however, will introduce new network requirements that are best served by the existing, richly robust operability provided by IP based applications, given that high latency and disruptive conditions do not generally exist on a terrestrial network. It is likely that networks on other worlds will consist of a combination of BP only, dual stacked IP/BP, and IP only nodes. Each of these would deploy the appropriate network connections based on the utility required and access control status of indivudial devices or networks.

4. Additions to DNS

The latencies and distances associated with interplanetary networking are fundamentally incompatible with the query/response nature of DNS as presently implemented. As well, the ephmeral nature of DNS records presents the possibility of queried records becoming stale in transit over great distances. It is reasonable to assume that astronauts and eventual colonists on other planetary bodies will want to enjoy local DNS service. A solution to this employs discrete root server networks on such planetary bodies, allowing local traffic to stay local, while enabling such fundamental local tools as http, smtp, ssh, ntp, etc. The question becomes: "How to we delineate between these disparate networks, such that resources located in any of them are accessible to all of them?" In order to achieve naming interoperability between multiple isolated DNS root systems, a set of new TLDs will be required. Specifically, one or more TLD for each remote world’s network SHOULDl be added to each network’s root zone. These new TLDs need not be populated with complete, real data, as all IP related entries therein will resolve to IP/BP gateways eliminating the need for long distance synchronization of DNS data. Root server networks publishing TLDs MUST publish valid A or AAAA records only on the world where they represent local assets. Sub-domains in stub domains MAY populate records to include IPN records containing BP

Node Numbers and MX records pointing to terrestrial Interplanetary Mail gateways. Similarly, CNAME records can exist in stubs, directing requests to appropriate application layer gateway hosts The existing terrestrial TLDs will maintain their scope for Earth based communications. DNSSEC related records, such as RRSIG or NSEC, can protect stub domain data representing off-world resources.

Consider the example of Earth, Moon, and Mars as populated worlds. Each has a robust local Internet, is numbered in a universally unique manner, and has it’s own DNS root server network. The contents of these three naming databases will vary greatly, and that is no cause for concern, as each is limited in scope to it’s local planetary body. TLD pointers to the local IP/BP gateway, or "InterPlanetary Exchange Point", also exist to allow traffic to named destinations in other planet’s networks. On Earth, at least one per planetary body networked additional (.luna and .mars or similar) TLD would exist for this purpose; on Mars, stubs for Earth based TLDs and .luna; on the Moon, stubs for Earth based TLDs and .mars. In this example, a user on Mars may request the following resource on their device:

https://foo.org/researchpaper.pdf

The IP/BP gateway collects the request and creates a bundle containing request data, which is sent to Earth. A command to fetch

https://foo.org/researchpaper.pdf

is generated in the IP/BP gateway on Earth from said data in said bundle. researchpaper.pdf is then placed into one or more bundles, and returned to the Mars IP/BP Gateway. At no time has either DNS or HTTPS traffic passed through the bundle network. All DNS and HTTPS requests remained local to their own networks. Only the metadata required to specify and identify a request in the remote network, and the result of that request, if applicable, have passed through the BP network. As a matter of operational concern, application services will need to be handled on a service by service basis. The above example examines a simple HTTPS document request, handled in a particular way, algorithmically. Similar considerations will need to be me made for other non-interactive or batchable protocols, and some interactive protocols which can be made to function in a delay/ disruption-tolerant fashion. Either a daemon similar to an inetd superserver, or other modular listeners, will be required to parse and bundleize requests to applicable services integrated with a bundle protocol application capable of making generic IP client requests, storing, and forwarding the results. The details of these are beyond the scope of this document, and are a subject of active and potentially expanded future research and development. It is understood that there are political and economic ramifications

associated with the deployment of TLDs, or other specified domains, in the existing Earth-based DNS network. It is also recognized that any domain can also serve this function, provided that it has in its hierarchy foo.mars., such as foo.mars.sol.int, for example, when considering an Earth to Mars connection similar to the above.

5. DNSSEC and other Cryptographic Considerations

With the addition of new root server networks, as relates to DNSSEC operation, nothing changes; it simply needs to be instantiated with new, discrete local trust anchors in each new instance. Since there is no "cross pollination" of DNS data, there is no need for syncronization between discrete systems, as all signed records are limited to their respective networks.

In order to secure each phase of an Interplanetary Internet transaction, certain accomdations must be made in respective systems of the gateway node. While these considerations, in detail, are beyond the scope of this document, generalities can be explored. As off-world domains are intended to resolve to local gateways, those gateways can complete crytographically dependent requests, like https and smtps, with wildcard certificates applied to off-world TLDs, and optionally, dedicated certificates can be applied to more specific domain names. DMARC components can be handled in a similar fashion, ensuring that only network-local resource s are required to complete any given IP request.

Using the 3rd level domain mapping method described above, however, does present the challenge of lacking end-to-end cryptographic integrity and confidentiality assurance using a single cryptographic system. While cryptographic protection of this data can exist in a segmented fashion, (TLS-BPSEC-TLS), no solution to maintain end-to-end TLS integrity/confidentiality presently exists. In consideration of the design of any such end-to-end solution, per the https example above, it must be recognized that the cryptographic protection presumably ends inside any webserver which receives a request, so it can process the request and respond over the encrypted channel. The webserver is presumed trusted in this scenario. Simularly, in a segmented confidentiality and integrity approach, the primary output from the https listener is not directed back through the original encrypted channel initially, but instead, is transmitted using different protocols,over a different channel, encrypted with appropriate methods for the network to be traversed.

6. IANA Considerations

This memo includes no request to IANA.

7. Security Considerations

Significant security considerations must always be made when making network connections to assets in space or on other worlds. Chiefly among these are the desire to prevent direct IP connections to off-world assets. This prevents a variety of attacks. As has been noted concerning the ongoing widespread July 2024 outage, a host which is not accessible over the Internet is impervious to network based attacks. The store and forward nature of BP allows the opportunity for greater scrutiny of traffic entering a BP network. Fundamentally, the attack surface presented by the gateway/exchange, and those services made available thereby, are limited in scope as relates to the off-world network. DDoS attack reach would be limited to terrestrial gateway nodes, and could be mitigated by whitelisted access at the network layer.

8. Conclusion

It is the opinion of the author that the most effective architecture for expansion of networking to support Human exploration of the solar system, i.e. the creation of a "Solar System Internet" requires appropriate use of both the Bundle family of protocols and the IP family of protocols. The former, used for most space based and some planetary based assets, provides a delay/disruption tolerant interplanetary backbone and fine tailored control over space hardware systems, while the latter provides the robust range of functionality that underpins the utility of the modern Internet in terrestrial and some orbital systems. A BP application method for DNS interoperability between multi-terrestrial Internet based networks has been presented, but is intended by the author to be viewed as a guideline towards a fully integrated and cohesive Solar System Internet.

9. References

9.1. Normative References

[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>.

[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>.

Acknowledgements

Thanks are due to Vint Cerf and Scott Burleigh, both of whom made meaningful contributions to the initial and ongoing discussions which led to this document and the concepts herein.

Jim Schier and Leigh Torgerson provided useful feedback on a previous , privately published version of this document, which is greatly appreciated.

Thank you to Shota Suzuki Yoshiki Uchida, and Tony Chen from Keio University for their suggestion of discrete TLD’s "per world," or the scope of a given TLD being limited to a single world, in addition to successfully replicating and validating the results of lab experiments employing this architecture for Interplanetary DNS

Thank you to Nathan Luu from California State University, Los Angeles for suggesting treatment of IPN records in context of this writing.

Author’s Address

Scott M. Johnson (editor)
Spacely Packets, LLC
46 High Ridge Road
Daytona Beach, FL 32117
United States of America
Phone: 386-888-7311
Email: scott@spacelypackets.com