

# Draft-many-tiptop-ip-architecture Update

Marc Blanchet, Viagénie  
Wesley Eddy, Aalyria  
Tony Li, Hewlett-Packard Enterprise

IETF 125 Shenzhen, TIPTOP working group

March 17th 2026

# Status

- WG call for draft-many-tiptop-ip-architecture-02 adoption by tiptop wg, co-chairs on December 14th 2025
- Outcome sent on 2026-01-12 by co-chairs: "We received responses from 22 participants. A clear majority expressed support for adoption. Several participants did not support adoption at this stage and raised points requiring further work, and a few provided comments without stating an explicit position. It is important to restate the purpose of a Working Group adoption call. Support for adoption reflects Working Group interest in the topic and in progressing the work within the WG. It does not imply agreement with the current text, solution approaches, or conclusions, nor does it represent approval of a final document."
  - The summary included a table of comments
  - Request from co-chairs to update the document
- New version (-03) of the document posted on March 2nd
  - reply to the co-chairs table of comments
- This presentation is the same content as the reply
  - 10 people supported the adoption with no additional request/comment. Not shown in this presentation.
- Caveat: while we tried to resolve the comments, the outcome is not a document ready for publication, but (hopefully) ready for adoption. We still have work to do on it.

# Comments

- From: Carles Gomez Montenegro
- Support for adoption: Yes
- Co-chairs: "Considered the work aligned with the charter; suggested an editorial change to rename Section 6 and structure it by application protocol."
- Done:
  - Created an Application section, where HTTP is one of the section
  - Added generic text at the beginning of the Application section.

5.4. Other Transports . . . . .	12	6.4. Other Transports . . . . .	16
6. HTTP . . . . .	12	7. Applications . . . . .	16
6.1. CoAP . . . . .	13	7.1. HTTP . . . . .	16
		7.1.1. CoAP . . . . .	18
7. Network services . . . . .	13	8. Network services . . . . .	18

space.

## 7. Applications

While many Web applications already use REST, operating them in deep space requires a concrete deployment and design model that differs from terrestrial assumptions. A functional deep-space Web system would consist of RESTful HTTP (or CoAP) services with explicitly configured timeouts, caching behavior, and asynchronous request handling. Application endpoints are designed to accept long request-response delays, return idempotent resources, and avoid chatty multi-RTT interactions. Content is aggressively pre-cached and replicated at planetary or orbital gateways so that most REST interactions are satisfied locally, while end-to-end exchanges tolerate hours or days of delay. In this model, a "Web app" is not an interactive browser session, but a set of asynchronous REST resources whose state can be fetched, updated, or synchronized opportunistically as connectivity becomes available. This section describes the specific protocol configurations and application patterns needed to adapt existing REST-based systems into such a deep-space-capable Web architecture.

### 7.1. HTTP

# Comments (cont.)

- From: Meiling Chen
- Support for adoption: Yes
- Co-Chairs: "Requested a clearer end-to-end architecture view; clarification on protocols; expanded security discussion; store-and-forward mechanisms without protocol modification."
- Done: added a whole new section (now section 2) on architecture overview, defined node types, figure, new rewritten security section, enhanced introduction

## 2. Architecture Overview

This section provides a short overview of this architecture for IP in deep space.

One of the key reasons to use IP is to maintain compatibility and potential smooth interoperability with other IP networks that exist terrestrially and within space vehicle onboard networks. While some nodes will require specific deep space protocol support, many nodes within the network can have traditional unmodified IP stacks that remain configured as if for typical Internet uses.

Within an end-to-end deep space application data flow, several different types of nodes can be involved:

\* Unmodified / Internet nodes – These consists of existing "untuned" IP stacks (not necessarily modified in any way for deep space use versus typical Internet use). These can exist within planetary surface networks and spacecraft onboard networks, for example.

For instance, these may be the majority of nodes within terrestrial ground systems, mission operations centers, ground stations, and terrestrial backbone networks. Given use of 3GPP architecture in some agency plans, these nodes may be present within onboard or surface core networks.

\* TIPTOP end-hosts – These nodes include specially tuned transport/application protocol parameters in order to function in deep space scenarios. Within the network, these specially tuned nodes need only exist at application endpoints (or proxies).

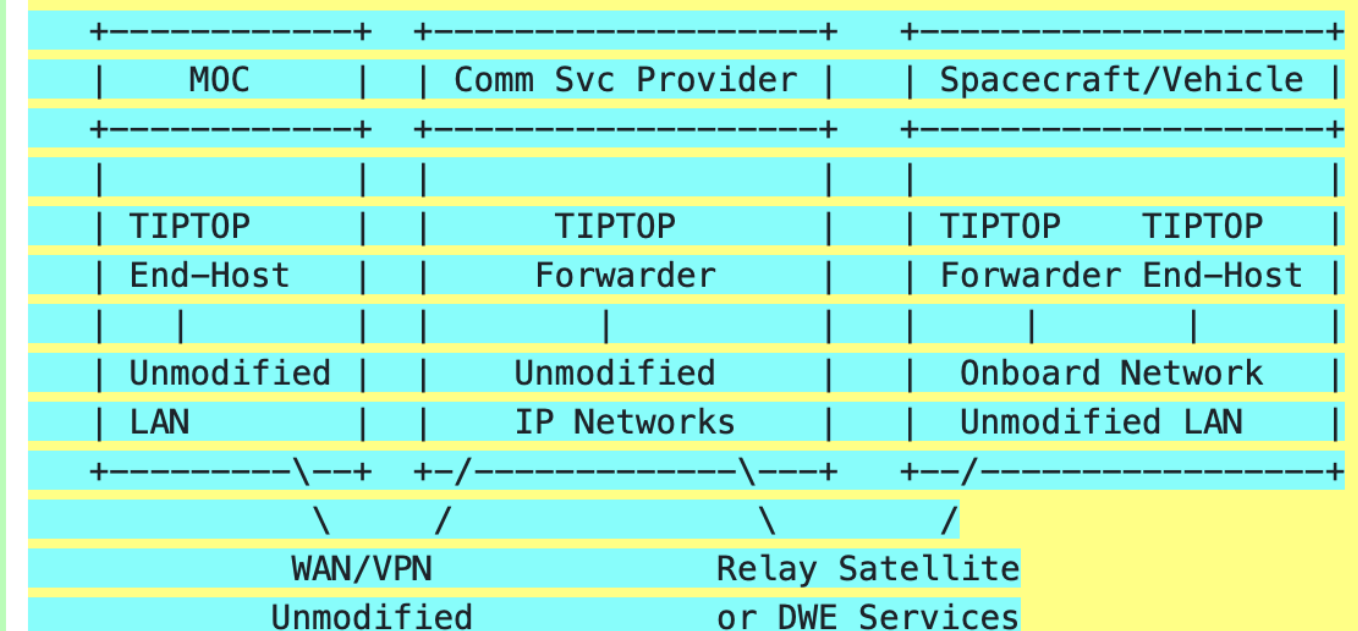
\* TIPTOP forwarding nodes – These are like typical IP forwarding nodes, except include queuing capabilities for IP packets that go beyond normal Internet use, as described more in Section 4.

Within the total deep space network, these need only exist at transition points between well-connected regimes and the hops through deep space communication system links. For instance, these may be at the edge of onboard networks or ground station networks.

\* TIPTOP-full nodes – These nodes combine aspects of TIPTOP end hosts and TIPTOP forwarding nodes. This combination may be useful in cases such as small spacecraft or other vehicles that consist of only a single node that must roam autonomously rather than functioning within a well-connected onboard or surface network regime.

\* Scheduling/orchestration systems – These systems may typically operate out-of-band from application data flows, but are responsible for creating and distributing coordinated network plans based on projected application needs (e.g. to support planned mission operations schedules). These may only exist within terrestrial or other planetary core networks, but are essential, especially in early growth of deep space networks, in order to provide knowledge of expected connectivity schedules and associated stack parameters.

The figure below illustrates a simplified example set of end-to-end of mission network elements based on this taxonomy, with the data flow described following.



An example end-to-end command application flow might pass through many nodes, starting within a mission operations center (MOC) at (1) a TIPTOP end-host running command workstation software, and then traversing (2) an unmodified LAN within the MOC, which routes the command packet via (3) an unmodified WAN or VPN connection to a communication service provider (CSP) network. Within the CSP network, there may be (4) unmodified IP forwarding in other WAN, VPN, and LAN nodes, and (5) a TIPTOP forwarder that queues packets and routes according to schedule via (6) unmodified relay or direct communications services towards the destination spacecraft node. At the destination spacecraft/vehicle, there may be (7) a TIPTOP forwarder, (8) an unmodified onboard network, and (9) a TIPTOP end-host.

Within this example command path, there are 2 TIPTOP end hosts, 2 TIPTOP forwarders, and 5+ unmodified typical Internet stack nodes (since multiple hops can be involved in the several LAN and WAN traversals). This deep space IP architecture is incrementally deployable.

Making the end-to-end example command data flow possible, an orchestration system on Earth can use knowledge of the spacecraft mission and its commanding needs in order to ensure that at proper times there are scheduled communication system resources (antennas, modems, etc.) that are physically pointed, tracking, and otherwise configured to close a forward link to the spacecraft at the required data rate. The orchestration system can also set scheduled routing table entries for the relevant nodes within the path.

# Comments (cont.)

- From: Britta Hale
- Support for adoption: No
- Co-Chairs: "Raised concerns about architectural scope, solution/protocol specificity, and treatment of security considerations."
- Done: removed QUIC as implicit default, and write it as just an example, added other transports, rewrote 0-RTT text being subject to replay attacks and should be disabled depending on the security policy, rewrote text on key establishment at launch, new rewritten security section.

established connections long-term. However, this will not be possible in all cases for all applications, if flows can't be planned pre-launch. Due to the need to be robust to stale packets, errors, and denial of service attacks, historically, Internet transports have included handshaking state machines, such as 3-way and 4-way handshakes for connection establishment (the case can be even worse for some transport protocol and TLS combinations), although QUIC can establish secured connections with only 1 round-trip time. Since the interplanetary round trip times may be larger than the duration of contact periods, these handshaking mechanisms are very inefficient. Though they are tolerable in cases such as between Earth and lunar networks, they are stifling for Earth-Mars and other interplanetary network paths. Transports, such as QUIC, that have the ability to resume based on shared state from prior application connections or to rapidly start transferring data ("0-RTT") can be suitably efficient, once initial state is obtained; however, they still require a complete initial handshake with the full set of round trip times imposed. For interplanetary use, it may be beneficial to find ways to securely pre-set information to allow this more efficient startup, without requiring the full initial handshake

planned pre-launch. Due to the need to be robust to stale packets, errors, and denial of service attacks, historically, Internet transports and security protocols have included handshaking state machines, such as 3-way and 4-way handshakes for connection establishment (needing 1.5 or 2 RTTs, if client authentication is needed and pre-shared keys are not available). Since the interplanetary round trip times may be larger than the duration of contact periods, these handshaking mechanisms are very inefficient. Though they are tolerable in cases such as between Earth and lunar networks, they are stifling for Earth-Mars and other interplanetary network paths. Transports, such as QUIC, that have the ability to resume based on shared state from prior application connections or to rapidly start transferring data ("0-RTT") can be suitably efficient, once initial state is obtained; however, they still require a complete initial handshake with the full set of round trip times imposed. The use of 0-RTT is subject to replay attacks[RFC9001] and therefore should be considered to be disabled depending on the security policy of the mission. For interplanetary use, it may be beneficial to find ways to securely pre-set information to allow this more efficient startup, without requiring the full initial handshake even once.

established connections long-term. However, this will not be possible in all cases for all applications, if flows can't be planned pre-launch. Due to the need to be robust to stale packets, errors, and denial of service attacks, historically, Internet transports and security protocols have included handshaking state machines, such as 3-way and 4-way handshakes for connection establishment (needing 1.5 or 2 RTTs, if client authentication is needed and pre-shared keys are not available). Since the interplanetary round trip times may be larger than the duration of contact periods, these handshaking mechanisms are very inefficient. Though they are tolerable in cases such as between Earth and lunar networks, they are stifling for Earth-Mars and other interplanetary network paths. Transports, such as QUIC, that have the ability to resume based on shared state from prior application connections or to rapidly start transferring data ("0-RTT") can be suitably efficient, once initial state is obtained; however, they still require a complete initial handshake with the full set of round trip times imposed. The use of 0-RTT is subject to replay attacks[RFC9001] and therefore should be considered to be disabled depending on the security policy of the mission. For interplanetary use, it may be beneficial to find ways to securely pre-set information to allow this more efficient startup, without requiring the full initial handshake even once.

# Comments (cont.)

- From: Britta Hale (cont.)
- Done: new rewritten security section

## 10. Security Considerations

Using the current IP protocol stack in deep space inherits all the work on privacy, cryptography, key management, firewalls, and scrutiny of protocols that are deployed on the Internet. Since no

change is made in the protocols, this architecture does not bring new security issues on the protocols themselves.

However, with longer delays and intermittence, deep space networking brings additional considerations. Security designs must tolerate long validation delays and potentially operate in a disconnected trust model.

Certificates and keys need to be renewed before their expiration, taking into account the delay to send, receive and confirm. Protocols such as OCSP[RFC6960] providing on-line real-time validation and revocation check will likely not work given the too long delays, therefore certificates need to be validated using local trust anchors.

The use of long term keys, such as ones set prior to launch, may create exposure, therefore keys should be renewed at appropriate frequency.

Given possible lower frequency of time synchronization, clock drifts may affect expiration and validation.

Intermediate forwarding nodes may buffer packets for a significant time. While it is presumed that most IP packet payloads will be encrypted by IP or transport security, data-at-rest becomes a possible exposure. If the intermediate node is compromised, the data-at-rest becomes available for the bad actor. This is new and not expected on terrestrial Internet intermediate nodes. If bad traffic is injected sufficiently to fill out the intermediate nodes buffer, then this becomes a denial-of-service attack.

If some packets are buffered in one intermediate node, If multiple paths are possible to a destination, then bad packet injection may use an alternate faster path that would result in the destination receiving the injected bad packets before the proper ones, which may become denial-of-service attacks of various kinds (destination endpoints buffer full, connection resets, ...). While this scenario can happen on terrestrial Internet, the bad traffic level and the duration window may be orders of magnitude larger than on Internet, therefore having a potentially much larger impact.

Given possible short contact windows and relatively few alternate paths, an attacker may flood a link during the whole contact window, disabling any remediation during that contact window, which means the actual impact will remain longer than the actual attack duration.

Given relatively few alternate paths, malicious injection of routes may have a larger impact than on terrestrial Internet.

Given a sparse network of few nodes and more predictable traffic patterns than terrestrial Internet, traffic analysis may become more effective, even for encrypted traffic.

Security considerations of each transport protocol are discussed in their respective transport protocol profile document.

Given low bandwidth, low number of alternate paths, high costs of links and nodes high value, access control to the deep space network and related policies should be in place. For example, at the beginning of the deployment, the deep space network shall be isolated from the current Internet by an "air gap", to disable any direct communications from the Internet to deep space. Moreover, destination IP prefix filtering shall be used to restrict the traffic to only the relevant one for each link. Note that this shall also be implemented in the routing control plane, but additional security might be appropriate to further protect the deep space links.

Each tiptop forwarding node, such as celestial network edge device, shall have firewall rules to prevent inappropriate traffic from entering deep space links. If communications from Mars may only occur to Earth, but not to the Moon, then appropriate filtering based on destination IP prefixes shall be used.

# Comments (cont.)

- From: Marshall Eubanks
- Support for adoption: Yes
- Co-Chairs: "Suggested clarifying and retaining surface networking as part of an end-to-end architecture."
- Done: section 1 has been enhanced and also referred to use case doc., section 2 added

Currently, space agencies have published plans that include deploying IP networks on celestial bodies, such as the Moon [ioag] and Mars [ioag-mars], both on the surfaces and orbiting constellations, including link layer technologies such as Wi-Fi or 5G. On the surface, the plans involve dense networking around facilities and habitats. New mission concepts are also including clusters of multiple networked nodes co-located at Lagrange points.

Given the evolution of modern IP application protocol stacks and the new needs of deep space missions, this document describes an architecture for the use of IP in deep space, meeting needs described in [I-D.ietf-tiptop-usecase]. This includes:

\* Unmodified / Internet nodes – These consists of existing "untuned" IP stacks (not necessarily modified in any way for deep space use versus typical Internet use). These can exist within planetary surface networks and spacecraft onboard networks, for example. For instance, these may be the majority of nodes within terrestrial ground systems, mission operations centers, ground stations, and terrestrial backbone networks. Given use of 3GPP architecture in some agency plans, these nodes may be present within onboard or surface core networks.

\* TIPTOP end-hosts – These nodes include specially tuned transport/application protocol parameters in order to function in deep space scenarios. Within the network, these specially tuned nodes need only exist at application endpoints (or proxies).

# Comments (cont.)

- From: Dongjie (Jimmy) Jie
- Support for adoption: Yes
- Co-Chairs: "Suggested clarifying WG scope regarding surface networking and softening solution-specific guidance."
- Done: see last slide

# Comments (cont.)

- From: Sean Turner
- Support for adoption: No
- Co-Chairs: "Suggested the draft be more requirements- and constraints-focused and less solution-centric."
- Done: new section 2 with node types, requirements. Quic removed as implicit default.

# Comments (cont.)

- From: Eric Rescorla
- Support for adoption: No
- Co-Chairs: "Suggested clarifying architectural intent, normative expectations, and removing vague examples."
- Done: new section 2 with node types, requirements. Quic removed as implicit default, removed WASM, new rewritten security section

# Comments (cont.)

- From: Johnson Liu(刘珺)
- Support for adoption: Yes
- Co-Chairs: "Suggested strengthening architectural functions beyond conceptual description."
- Done: new section on Architectural overview (section 2) (shown before). additional text regarding PCE.

Existing routing protocols require proof of liveness between protocol partners, implemented through the periodic exchange of packets between partners. This is impractical on long-delay or intermittent links, so a Path Computation Engine (PCE) [RFC4655] based approach seems appropriate for those domains possibly supplemented by contact plan schedules[I-D.ietf-tvr-schedule-yang]. Interconnection between domains can still be done with BGP [RFC4271], but long-delay or intermittent links should be avoided. Domains straddling such links must provide proxy advertisements for prefixes reachable across such links.

Optimal routing for domains with intermittent links is out of scope for this document.

On the surface of celestial bodies and in proximal orbit, traditional protocols are applicable and should be used (e.g., [RFC9717]).

Because of bandwidth limitations, the PCE for a domain may require the ability to specify an explicit path, but using an explicit path across an intermittent link may be problematic. If a packet is bound to an explicit path and stored in an intermediate node, then the path may be invalid at the end of the storage period. Conveying a backup path within the packet itself would incur a large performance penalty and is not recommended. Allowing the intermediate node itself to compute the remainder of the path would seem to be the most robust solution. Conventional traffic engineering techniques for getting to storage nodes seem to be appropriate.

# Comments (cont.)

- From: Russ Housley
- Support for adoption: No
- Co-Chairs: "Requested focus on architecture-level requirements rather than solution selection."
- Done: replied on the mailing list. RIR person comment did not require any change. Addressing is enhanced and referred to the other document

## 5. IP Addressing and Routing

### 5.1. Addressing

The IP address space is a hierarchical namespace where ranges of addresses are encoded as "prefixes". Individual domains advertise prefixes to the broader Internet and assign these addresses internally. Prefixes may be aggregated into less-specific prefixes, which makes the routing subsystem more efficient by decreasing overhead.

Space networks provide a unique opportunity to provide extremely efficient routing by assigning a unique prefix or block of addresses per celestial body and its proximal orbits. Management of the IP address space is currently documented in [RFC7020], but this only covers continental regions and does not provide for addressing for space. The depletion of the IPv4 address space means that there is little that can be done if agencies decide to use IPv4. However, if agencies decide to utilize IPv6, it would be deeply beneficial if addressing was designed to allow for the maximum aggregation of routing prefixes. Aggregation of prefixes assigned to celestial bodies would minimize the overhead incurred by relaying spacecraft, minimizing expensive hardware requirements and would help minimize route flap. Thus, to help achieve maximum aggregation, the address space for outer space should be managed by a Regional Internet Registry (RIR) and blocks of address space should be allocated for each celestial body of interest. Space service providers should use prefixes assigned by this RIR. This is discussed in more detail in [I-D.li-tiptop-address-space].

# Comments (cont.)

- From: Benjamin Dowling
- Support for adoption: No
- Co-Chairs: "Requested clearer articulation of security goals for space networking."
- Done: rewrote security section (see before). Did not add new security requirements/arch as they should be first defined in the use case-requirements document.

# Comments (cont.)

- From: Jad El Cham
- Support for adoption: No position
- Co-Chairs: "Suggested engaging RIR community and ensuring realistic addressing proposals."
- Done: Addressing section enhanced. Reference to the other draft. RIR discussions have been started.

# Comments (cont.)

- From: Xisen Tian
- Support for adoption: No
- Co-Chairs: "Raised security-focused concerns regarding protocol profiling and threat assumptions."
- Done: removed QUIC as implicit default, and write it as just an example, added other transports, rewrote 0-RTT text being subject to replay attacks and should be disabled depending on the security policy, rewrote text on key establishment at launch, new rewritten security section.

# Summary

- Editors hopefully have addressed the wg adoption call comments.
- This does not say the document is ready for publication, but for adoption
- Requesting wg adoption of draft-many-tiptop-ip-architecture-03