Workgroup: Network Working Group
Internet-Draft:
draft-templin-intarea-parcels2-00
Published: 15 February 2024
Intended Status: Standards Track
Expires: 18 August 2024
Authors: F. L. Templin, Ed.
Boeing Research & Technology

IPv4 Parcels and Advanced Jumbos (AJs)

Abstract

IPv6 Parcels and Advanced Jumbos (AJs) present new data packaging constructs and a new link model for Internetworking. As is often the case, technologies developed in the IPv6 space can also be applied in IPv4 and vice-versa. This document presents the adaptations necessary to support Parcels and AJs in IPv4.

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 18 August 2024.

Table of Contents

1. Introduction

IPv6 Parcels and Advanced Jumbos (AJs) [I-D.templin-6man-parcels] present new data packaging constructs and a new link model for Internetworking. As is often the case, technologies developed in the IPv6 space [RFC8200] can also be applied in IPv4 [RFC0791] and vice-versa. This document presents the differences that need to be addressed to adapt IPv6 Parcels and AJs to IPv4.

All aspects of IPv6 Parcels and AJs, including the use of IP extension headers and control messaging, apply also to IPv4. Only differences in the IP header format and some control option encapsulations need to be accounted for as discussed below. This document therefore specifies IPv4 parcels and AJs.

2. Requirements

IPv4 parcels and AJs observe all requirements established for IPv6 [I-D.templin-6man-parcels] including the use of IPv6 Hop-by-Hop Options headers. This means that nodes that recognize IPv4 parcels/AJs MUST recognize and correctly process IP protocol 0 (Hop-by-Hop) extension headers the same as for IPv6 when they occur in an extension header chain following the IPv4 header but before the upper layer payload.

When an IPv4 router or destination end system processes a parcel/AJ probe for which the IPv4 Protocol field encodes an unrecognized value (such as 0 for Hop-by-Hop Options), it drops the probe and returns an ICMPv4 "Destination Unreachable - Protocol Unreachable" message [RFC0792]. The source should regard any such messages as an advisory indication that OMNI protocol UDP encapsulation may be necessary in future probes.

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. IPv4 Parcels and Advanced Jumbos (AJs)

All aspects of [I-D.templin-6man-parcels] are imported as normative specifications for IPv4 parcels and AJs, with the exception of the following differences:

3.1. IPv4 Total Length

The IPv6 header includes a "Payload Length" field defined as the: "Length of the IPv6 payload, i.e., the rest of the packet following this IPv6 header, in octets". The IPv4 header instead includes a "Total Length" field defined as: "the length of the datagram, measured in octets, including internet header and data".

IPv4 parcels/AJs always set the Total Length to the minimum of 576 octets and the actual length of the packet. This length provides a truncation limit for routers that do not recognize the parcel/AJ format but instead attempt to forward as an ordinary IPv4 packet.

3.2. IPv4 Time To Live (TTL)

The IPv4 "Time To Live (TTL)" and IPv6 "Hop Limit" values are treated in exactly the same way in both protocol versions. In particular, the source sets the TTL/Hop Limit to an initial value and each router in the path to the destination decrements the TTL/Hop Limit by 1 when it forwards a parcel/AJ/probe. (Note that this represents a parcel/AJ-specific requirement for IPv4 routers, since [RFC1812] permits routers to decrement TTL by values other than 1.)

3.3. IPv4 Parcel/Jumbo Payload Length

The same as for IPv6, the Parcel Payload Length field in the Parcel Payload Option and the Jumbo Payload Length field in the Jumbo Payload Option of IPv4 parcels/AJs encode the length of the IPv6 extension headers plus the length of the {TCP,UDP} header plus the

combined length of all concatenated segments with their per-segment headers/trailers.

Therefore, the length of the IPv4 header itself is not included in the Parcel/Jumbo Payload Length field the same as for IPv6. The IPv4 header length for IPv4 parcels and AJs is instead available in both the IPv4 Total Length and Internet Header Length (IHL) fields.

3.4. IPv4-Compatible IPv6 Addresses

Whenever an IPv4 address needs to be coded in an IPv6 address field, the address is coded as an IPv4-compatible IPv6 address as specified in [RFC4291].

3.5. IPv4 Parcel Packetization/Restoration

When a node performs packetization on a {TCP,UDP}/IPv4 parcel, it inserts a Parcel Parameters {TCP,UDP} option the same as for IPv6 [I-D.templin-6man-parcels].

The IPv4 destination then performs restoration by gathering up IPv4 packets that arrive with the same upper layer 5-tuple and with Parcel Parameters information including the same Identification. The Parcel Parameters Index then determines the ordinal position of each packet segment to be concatenated into the restoration buffer, i.e., the same as for IPv6. (Note: if the IPv4 destination does not recognize the {TCP,UDP} Parcel Parameters option, it simply processes the packet as a singleton IPv4 packet. This would result in correct behavior, but with Generic Receive Offload (GRO) disabled.)

3.6. Parcel Probing

When an IPv4 router or destination receives an intact parcel probe, it processes the probe the same as specified for IPv6 [I-D.templin-6man-parcels].

When an IPv4 router forwards the parcel probe, it MUST decrement the TTL by exactly 1 the same as specified for the IPv6 Hop Limit.

When an IPv4 destination receives a parcel probe, it should return a Parcel Parameters option in any {TCP,UDP}/IPv4 packet to be returned to the source the same as for IPv6.

When the IPv4 source receives a {TCP,UDP} packet that includes a Parcel Parameters option it matches the Identification value with its recently-transmitted probes. If there is a match, the source accepts the MTU and Parcel Limit values found in the Parcel Parameters option.

The same as for IPv6, if the source or destination is located outside of a controlled environment / limited domain [RFC8799] the source should send probes including the IPv4 header followed by an OMNI UDP encapsulation header followed by the Hop-by-Hop Options header and finally followed by the {TCP,UDP} header plus protocol data.

3.7. Parcel/Jumbo Replys

When an IPv4 router returns a Parcel/Jumbo Reply, it prepares an ICMPv4 message according to [RFC0792] with Type set to TBD1 and with Code set to TBD2 for a Parcel Reply or TBD3 for a Jumbo Reply (see: IANA Considerations).

The router populates the ICMPv4 message header fields, then copies up to 576 octets from the parcel/AJ into the ICMPv4 message body. The router then calculates and sets the ICMPv4 Checksum and returns the message to the parcel/AJ source.

3.8. Advanced Jumbos (AJs)

All aspects of IPv4 Advanced Jumbos (AJ) are processed the same as for IPv6 AJs.

3.9. Jumbo-in-Jumbo Encapsulation

Original IPv4 parcels/AJs can follow the "e(X)treme" forwarding paths across successive OMNI links in the path using "jumbo-in-jumbo" encapsulation the same as for IPv6. The OMNI link ingress encapsulates each IPv4 parcel/AJ in an OMNI IPv6 header plus any outer L2 encapsulations which may include an IPv4 header with an Advanced Jumbo option Hop-By-Hop extension header. All aspects of this "jumbo-in-jumbo" encapsulation are the same as for IPv6.

3.10. Integrity

To support the IPv4 parcel/AJ header checksum calculation, the network layer uses modified versions of the {TCP,UDP}/IPv4 pseudo-header found in [RFC9293]. Note that while the contents of the two IP protocol version-specific pseudo-headers beyond the address fields are the same, the order in which the contents are arranged differs and must be honored according to the specific IP protocol version.

The IPv6 pseudo-header is found in [I-D.templin-6man-parcels], while the IPv4 pseudo-header is shown in Figure 1. The similarities between the two pseudo-headers allows for maximal reuse of widely deployed code while ensuring interoperability.

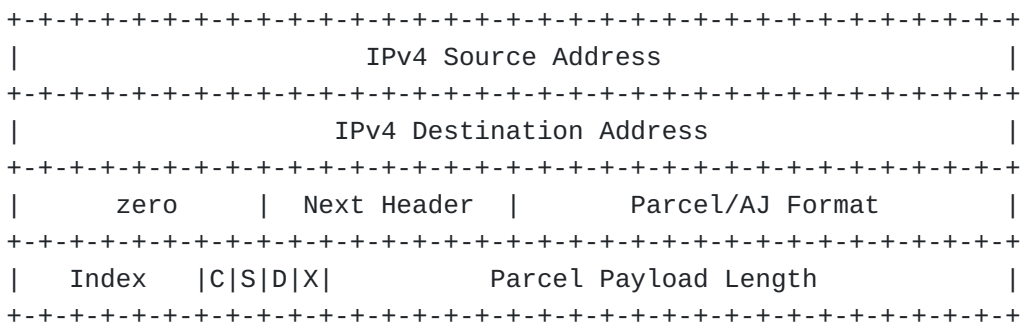

Figure 1: {TCP,UDP}/IPv4 Parcel/AJ Pseudo-Header Format

Note: The same as for IPv6, the "Index/C/S/D/X" and Parcel Payload Length fields in the IPv4 parcel pseudo-header are replaced by the single 4-octet Jumbo Payload Length field in the IPv4 AJ pseudo-header.

4. Implementation Status

An early prototype of UDP/IPv4 parcels (draft version -15) has been implemented relative to the linux-5.10.67 kernel and ION-DTN ion-open-source-4.1.0 source distributions. Patch distribution found at: "https://github.com/fltemplin/ip-parcels.git".

5. IANA Considerations

The IANA is instructed to assign a new Type number TBD1 in the 'icmp-parameters' registry "ICMP Type Numbers" table (registration procedures IESG Approval or Standards Action). The entry should set "Type" to TBD1, "Name" to "Packet Too Big (PTB)" and "Reference" to [RFCXXXX] (i.e., this document).

The IANA is further instructed to create a new table titled: "Type TBD1 - Packet Too Big (PTB)" in the 'icmp-parameters' Code tables, with registration procedures IESG Approval or Standards Action. The table should have the following initial format:

Code	Name	Reference
TBD2 (3 suggested)	Parcel Report	[RFCXXXX]
TBD3 (4 suggested)	Jumbo Report	[RFCXXXX]

Figure 2: Type TBD1 - Packet Too Big (PTB)

6. Security Considerations

Security Considerations are the same as for IPv6 as found in [I-D.templin-6man-parcels].

7. Acknowledgements

This work was inspired by ongoing AERO/OMNI/DTN investigations. The concepts were further motivated through discussions with colleagues.

Honoring life, liberty and the pursuit of happiness.

8. References

8.1. Normative References

[I-D.templin-6man-parcels]
Templin, F., "IPv6 Parcels and Advanced Jumbos (AJs)", Work in Progress, Internet-Draft, draft-templin-6man-parcels-19, 12 February 2024, <https://datatracker.ietf.org/doc/html/draft-templin-6man-parcels-19>.

[I-D.templin-intarea-omni]
Templin, F., "Transmission of IP Packets over Overlay Multilink Network (OMNI) Interfaces", Work in Progress, Internet-Draft, draft-templin-intarea-omni-66, 12 February 2024, <https://datatracker.ietf.org/doc/html/draft-templin-intarea-omni-66>.

[RFC0768] Postel, J., "User Datagram Protocol", STD 6, RFC 768, DOI 10.17487/RFC0768, August 1980, <https://www.rfc-editor.org/info/rfc768>.

[RFC0791] Postel, J., "Internet Protocol", STD 5, RFC 791, DOI 10.17487/RFC0791, September 1981, <https://www.rfc-editor.org/info/rfc791>.

[RFC0792] Postel, J., "Internet Control Message Protocol", STD 5, RFC 792, DOI 10.17487/RFC0792, September 1981, <https://www.rfc-editor.org/info/rfc792>.

[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>.

[RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <https://www.rfc-editor.org/info/rfc4291>.

[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>.

[RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <https://www.rfc-editor.org/info/rfc8200>.

[RFC9293] Eddy, W., Ed., "Transmission Control Protocol (TCP)", STD 7, RFC 9293, DOI 10.17487/RFC9293, August 2022, <https://www.rfc-editor.org/info/rfc9293>.

8.2. Informative References

[RFC1812] Baker, F., Ed., "Requirements for IP Version 4 Routers", RFC 1812, DOI 10.17487/RFC1812, June 1995, <https://www.rfc-editor.org/info/rfc1812>.

[RFC8799] Carpenter, B. and B. Liu, "Limited Domains and Internet Protocols", RFC 8799, DOI 10.17487/RFC8799, July 2020, <https://www.rfc-editor.org/info/rfc8799>.

Appendix A. Change Log

<< RFC Editor - remove prior to publication >>

Changes from earlier versions:

*Submit for review.

Author's Address

Fred L. Templin (editor)
Boeing Research & Technology
P.O. Box 3707
Seattle, WA 98124
United States of America

Email: fltemplin@acm.org