

IP Parcels Implementation Report

IETF115 Intarea Working Group Session (November 9, 2022)

Fred L. Templin (The Boeing Company)

fred.l.templin@boeing.com

Implementation Elements

- Parcel framing in ION-DTN LTP convergence layer (ion-open-source-4.1.0)
- Parcel extensions for (UDP_SEGMENT; UDP_GRO) socket API (linux-5.10.67)
- Linux kernel UDP/IPv4 parcel support, with “IPv4 Jumbo Payload” option
- IPv6 fragmentation in ip6_tunnel.c → RFC2473 surrogate OMNI interface
- SO_NO_CHECK for LTP/kernel parcel segment checksum selection (Rx/Tx)
- New skb_checksum_parcel() for kernel parcel segment checksums

Modules Updated

Linux kernel:

- ./net/core/dev.c
- ./net/core/skbuff.c
- ./net/core/sock.c
- ./net/ipv4/ip_input.c
- ./net/ipv4/ip_options.c
- ./net/ipv4/ip_output.c
- ./net/ipv4/udp.c
- ./net/ipv6/ip6_tunnel.c
- ./net/ipv6/udp.c
- ./include/linux/udp.h
- ./include/linux/skbuff.h
- ./include/net/udp.h
- ./include/net/ipv6.h
- ./include/net/inet_sock.h
- ./include/net/ip.h
- ./include/uapi/linux/ip.h

ION-DTN:

- ./ltp/udp/libudplsa.c
- ./ltp/udp/udplsa.h
- ./ltp/udp/udplsi.c
- ./ltp/udp/udplso.c
- ./ltp/library/ltpP.h

OMNI Interface

- RFC2473 IPv6 encapsulation/fragmentation (Adaptation Layer)
- Configured over 10Gbps P2P Ethernet (1500 MTU)

IPv4 in IPv6
Encapsulation

Large data received for small
of packets (58KB/pkt average)

Large Maximum Transmission Unit
(MTU = ((64KB-1) - 128))

```
fltemplin@barrier:~/DTN-OMNI/configs/25-ipn-ltp$ ifconfig omni0
omni0: flags=209<UP,POINTOPOINT,RUNNING,NOARP> mtu 65407
inet 192.168.0.2 netmask 255.255.255.0 destination 192.168.0.2
inet6 fe80::28e9:6eff:fe27:73c3 prefixlen 64 scopeid 0x20<link>
unspec FE-80-00-00-00-00-00-00-00-00-00-00-00-00 txqueuelen 1000 (UNSPEC)
RX packets 14562 bytes 853236422 (853.2 MB)
RX errors 0 dropped 0 overruns 0 frame 0
TX packets 83 bytes 21335 (21.3 KB)
TX errors 6 dropped 6 overruns 0 carrier 0 collisions 0
```

UDP/IPv4 Parcel

- Produced by LTP/UDP source; captured at destination OMNI interface
- Transparent IPv6 fragmentation over underlying 10Gbps Ethernet
- 64052-octet UDP/IPv4 parcel with 16 (3999-octet) segments:

IP version 4 with 28 octet header IPv4 Length = first segment length only IPv4 Jumbo Payload option:
Type = 0x0b (RFC1063); Len = 0x06; NSegs = 0x10;
Jumbo Payload Len = 0x00fa34; padding = 0x0000

```
|15:16:29.027313 IP 192.168.0.1.45080 > harrier.25113: truncated-udplength 0
UDP checksum (header only)  0x0000: 4700 0f9f acbc 0000 4011 2603 c0a8 0001 G.....@.&.....
                                0x0010: c0a8 0002 0b06 1000 fa34 0000 b018 6219 .....4....b.
                                0x0020: 0000 c33f 1957 cead b88e a26f 8c50 c1e5 ...?.W.....o.P..
UDP Length = 0              0x0030: a250 833b 6426 c510 a5fb 86e6 6751 c83b .P.;d&.....gQ.;
                                0x0040: a926 81cd 0014 0200 0100 9f17 0681 1011 .&.....
                                0x0050: 1901 1401 1401 0000 82d7 8ce7 4001 822c .....@...
                                0x0060: 0005 1009 6970 6e00 3230 2e30 0014 0101 ....ipn.20.0....
                                0x0070: 0001 09bd 8440 7465 7374 2e2e 2e00 0000 .....@test.....
```

Integrity block w/16 2-octet checksums (one per segment)

Parcel (continued)

0x0000:	4700	0f9f	acbc	0000	4011	2603	c0a8	0001	G.....@.&.....
0x0010:	c0a8	0002	0b06	1000	fa34	0000	b018	62194....b.
0x0020:	0000	c33f	1957	cead	b88e	a26f	8c50	c1e5	...?.W.....o.P..
0x0030:	a250	833b	6426	c510	a5fb	86e6	6751	c83b	.P.;d&.....gO.;
0x0040:	a926	81cd	0014	0200	0100	9f17	0681	1011	.&.....
0x0050:	1901	1401	1401	0000	82d7	8ce7	4001	822c@..,
0x0060:	0005	1009	6970	6e00	3230	2e30	0014	0101ipn.20.0....
0x0070:	0001	09bd	8440	7465	7374	2e2e	2e00	0000@test.....
0x0080:	0000	0000	0000	0000	0000	0000	0000	0000
0x0090:	0000	0000	0000	0000	0000	0000	0000	0000
0x00a0:	0000	0000	0000	0000	0000	0000	0000	0000
0x00b0:	0000	0000	0000	0000	0000	0000	0000	0000
0x00c0:	0000	0000	0000	0000	0000	0000	0000	0000
0x00d0:	0000	0000	0000	0000	0000	0000	0000	0000
0x00e0:	0000	0000	0000	0000	0000	0000	0000	0000
0x00f0:	0000	0000	0000	0000	0000	0000	0000	0000
0x0100:	0000	0000	0000	0000	0000	0000	0000	0000
0x0110:	0000	0000	0000	0000	0000	0000	0000	0000
0x0120:	0000	0000	0000	0000	0000	0000	0000	0000
0x0130:	0000	0000	0000	0000	0000	0000	0000	0000
0x0140:	0000	0000	0000	0000	0000	0000	0000	0000
0x0150:	0000	0000	0000	0000	0000	0000	0000	0000
0x0160:	0000	0000	0000	0000	0000	0000	0000	0000
0x0170:	0000	0000	0000	0000	0000	0000	0000	0000
0x0180:	0000	0000	0000	0000	0000	0000	0000	0000
0x0190:	0000	0000	0000	0000	0000	0000	0000	0000

← Parcel Headers

Start first segment

Parcel (continued)

0x5d00:	0000	0000	0000	0000	0000	0000	0000	0000
0x5d10:	0000	0000	0000	0000	0000	0000	0000	0000
0x5d20:	0000	0000	0000	0000	0000	0000	0000	0000
0x5d30:	0000	0000	0000	0000	0000	0000	0000	0000
0x5d40:	0000	0000	0000	0000	0000	0000	0000	0000
0x5d50:	0000	0000	0000	0000	0000	0000	0000	0000
0x5d60:	0000	0000	0000	0000	0000	0000	0000	0000
0x5d70:	0000	0000	0000	0000	0000	0000	0000	0000
0x5d80:	0000	0000	0000	0000	0000	0000	0000	0000
0x5d90:	0000	0000	0000	0000	0000	0000	0000	0000
0x5da0:	0000	0000	0000	0000	0000	0000	0000	0000
0x5db0:	0000	0000	0000	0000	0000	0000	0000	0000
0x5dc0:	0000	0000	0000	0000	0000	0000	0000	0000
0x5dd0:	0000	0000	0000	0000	0000	0000	0000	0000
0x5de0:	0000	0000	0000	0000	0000	0000	0000	0000
0x5df0:	0000	0000	0000	0000	0000	0000	0000	0014
0x5e00:	0200	0181	bb04	9f15	0000	0000	0000	0000
0x5e10:	0000	0000	0000	0000	0000	0000	0000	0000
0x5e20:	0000	0000	0000	0000	0000	0000	0000	0000
0x5e30:	0000	0000	0000	0000	0000	0000	0000	0000
0x5e40:	0000	0000	0000	0000	0000	0000	0000	0000
0x5e50:	0000	0000	0000	0000	0000	0000	0000	0000
0x5e60:	0000	0000	0000	0000	0000	0000	0000	0000
0x5e70:	0000	0000	0000	0000	0000	0000	0000	0000
0x5e80:	0000	0000	0000	0000	0000	0000	0000	0000
0x5e90:	0000	0000	0000	0000	0000	0000	0000	0000

← Later segment break

Parcel (continued)

```
0xf900: 0000 0000 0000 0000 0000 0000 0000 0000 0000 .....
0xf910: 0000 0000 0000 0000 0000 0000 0000 0000 0000 .....
0xf920: 0000 0000 0000 0000 0000 0000 0000 0000 0000 .....
0xf930: 0000 0000 0000 0000 0000 0000 0000 0000 0000 .....
0xf940: 0000 0000 0000 0000 0000 0000 0000 0000 0000 .....
0xf950: 0000 0000 0000 0000 0000 0000 0000 0000 0000 .....
0xf960: 0000 0000 0000 0000 0000 0000 0000 0000 0000 .....
0xf970: 0000 0000 0000 0000 0000 0000 0000 0000 0000 .....
0xf980: 0000 0000 0000 0000 0000 0000 0000 0000 0000 .....
0xf990: 0000 0000 0000 0000 0000 0000 0000 0000 0000 .....
0xf9a0: 0000 0000 0000 0000 0000 0000 0000 0000 0000 .....
0xf9b0: 0000 0000 0000 0000 0000 0000 0000 0000 0000 .....
0xf9c0: 0000 0000 0000 0000 0000 0000 0000 0000 0000 .....
0xf9d0: 0000 0000 0000 0000 0000 0000 0000 0000 0000 .....
0xf9e0: 0000 0000 0000 0000 0000 0000 0000 0000 0000 .....
0xf9f0: 0000 0000 0000 0000 0000 0000 0000 0000 0000 .....
0xfa00: 0000 0000 0000 0000 0000 0000 0000 0000 0000 .....
0xfa10: 0000 0000 0000 0000 0000 0000 0000 0000 0000 .....
0xfa20: 0000 0000 0000 0000 0000 0000 0000 0000 0000 .....
0xfa30: 0000 0000 .....

```

end of parcel (~4000 lines of text)

Single-segment Performance

- 1280: 609Mbps
 - 1460: 684Mbps
- } Simple UDP/IPv4 Packets Over Native Ethernet Interface (1500 MTU)

- 2000: 775Mbps
 - 4000: 1450Mbps
 - 8000: 2335Mbps
 - 9000: 2537Mbps
 - 10000: 2664Mbps
 - 12000: 3042Mbps
 - 14000: 3288Mbps
- } UDP/IPv4 Parcels Over OMNI Interface ((64KB-1)-128 MTU)

- 18000: 3668Mbps
- 24000: 4130Mbps
- 28000: 4268Mbps
- 30000: 4359Mbps
- 32000: 4616Mbps

758% increase from 1280 octets to 32000 octets

596% increase from 2000 octets to 32000 octets

- Proof that increasing segment size multiplies
- performance even with IPv6 fragmentation

Multi-segment Performance

- 1 Segment: 2000: 775Mbps / 4000: 1450Mbps
- 4 Segment: 2000: 965Mbps / 4000: 1636Mbps
- 6 Segment: 2000: 1029Mbps / 4000: 1730Mbps
- 8 Segment: 2000: 1037Mbps / 4000: 1810Mbps
- 10 Segment: 2000: 1078Mbps / 4000: 1754Mbps
- 12 Segment: 2000: 1095Mbps / 4000: 1812Mbps
- 16 Segment: 2000: 1072Mbps / 4000: 1862Mbps
- 20 Segment: 2000: 1076Mbps
- 24 Segment: 2000: 1078Mbps
- 28 Segment: 2000: 1108Mbps
- 30 Segment: 2000: 1144Mbps
- 32 Segment: 2000: 1137Mbps

28% increase for 16 segments @ 4000 octets

38% increase for 16 segments @ 2000 octets

47% increase for 32 segments @ 2000 octets

Next Steps

- Improve ION-DTN parcel integration robustness
- Support full (64KB-1) OMNI interface MTU
- System call; kernel segmentation support for large parcels (64KB+)
- Kernel UDP/IPv6 parcel support with “IPv6 Jumbo Payload” option
- Path qualification for parcels/jumbos over 64KB+ MTU links
- TCP/IP parcels and performance analysis

Backups

Github

<https://github.com/fltemplin/ip-parcels>

Drafts

[IP Parcels](#)

[OMNI](#)

[AERO](#)

[ATN/BGP](#)