Network Working Group B. Carpenter
INTERNET-DRAFT CERN
Expires: December 22nd, 1995

D. Katz
cisco Systems, Inc.

S. Thomas
AT&T Tridom

K. Sklower
University of California,
Berkeley

Mechanisms for OSI CLNP and TP over IPv6

draft-carpenter-ipv6-osi-01.txt

Status of This Memo

Abstract

This document defines a set of mechanisms for the support of OSI CLNP, and Transport Protocols over an IPv6 network. These mechanisms are the ones that MUST be used if such support is required.

Acknowledgements

All direct contributors of text are listed below as authors. The writers are also pleased to acknowledge the suggestions and comments of Richard Collella, Dirk Fieldhouse, Denise Heagerty, Cyndi Jung, Yakov Rekhter, and many other members of the former TUBA and new IPv6 working groups of the IETF. The support of Scott Bradner and Allison Mankin of the IESG was essential.

Conventions

The following language conventions are used in the items of specification in this document:

- o MUST, SHALL or MANDATORY -- the item is an absolute requirement of the specification.
- o SHOULD or RECOMMENDED -- the item should generally be followed for all but exceptional circumstances.
- o MAY or OPTIONAL -- the item is truly optional and may be followed or ignored according to the needs of the implementor.

Table of Contents

1. Summary of defined mechanisms

This document defines two mechanisms for carrying OSI traffic over an IPv6 network:

1. CLNP encapsulated in IPv6
2. Transport Protocol carried over IPv6

These are ELECTIVE mechanisms, i.e. they are not mandatory parts of an IPv6 implementation, but if such mechanisms are needed they MUST

be implemented as defined in this document.

Note in addition that an Internet Standard STD-35 "ISO Transport Service on top of the TCP" exists already [RFC1006]. There is a also a Proposed Standard for "OSI Connectionless Transport Service over UDP" [RFC1240]. Both of these documents may need revision for IPv6. All of these mechanism may co-exist.

2. CLNP encapsulated in IPv6

If it is required to tunnel CLNP [IS8473] through an IPv6 network, then the upper layer header SHALL be a CLNP PDU, and the final IPv6 Next Header field SHALL have the value 80 decimal (as defined for ISO-IP in [assigned]).

Mechanisms for the creation of CLNP tunnels and their management are outside the scope of this document.

Note that the tunnelling of CLNP over the Internet is discussed in detail in [RFC1070], but that document has no standards status and makes different assumptions about address mapping. In contrast to [RFC1070], CLNP tunnels through an IPv6 network are simply a virtual point-to-point encapsulation technology, using statically configured tunnel endpoints. There is no support for simulating a multipoint subnetwork, nor for dynamic mapping between NSAP addresses and IP addresses. Instead, IP addresses are simply viewed as Subnetwork Point of Attachment (SNPA) addresses that must be statically configured to create the tunnel.

Once a tunnel is established, data is transmitted using CLNP [IS8473]. The ES-IS [IS9542], IS-IS [IS10589], and IDRP [IS10747] protocols may be used to dynamically establish neighbor adjacencies and routing. Any NSAP addresses may be assigned to the systems at either end of the tunnel. There is no need to constrain the NSAP address format as documented in [RFC1070], since there is no need to perform dynamic address mapping. The "EON" header of [RFC1070] is not present.

No attempt is required to implement feedback of error indications from ICMP in the IP subnetwork into CLNP error PDUs. The tunnel is ignorant of problems in the IP subnetwork, and depends upon mechanisms in the OSI routing protocols to detect connectivity failures.

If a CLNP tunnel has an anycast destination, i.e. the packets are decapsulated by any one of a set of decapsulators, and if an IPv6 packet needs to be fragmented to get through the tunnel, the fragments may not be sent via same path. If this happens the original CLNP packet can never be decapsulated, since its fragments have arrived at different decapsulators. To avoid this problem, CLNP PDUs must be segmented as defined in [IS8473] if their size would create IPv6 packets exceeding the IPv6 path MTU. Reassembly will take place at the final destination according to [IS8473].

3. ISO Transport Protocols over IPv6

If it is required to carry ISO Transport Protocols [ISO8072, ISO8073] over an IPv6 network, then the IPv6 transport header SHALL be a TP PDU, and the final IPv6 Next Header field SHALL have the value 29 decimal (as defined for ISO-TP in [assigned]).

```
+---------------+-----------------------
|  IPv6 header  | TP PDU
|               |
| Next Header = |
|       ISO-TP  |
+---------------+-----------------------


+---------------+----------------+-----------------------
|  IPv6 header  | Routing header | TP PDU
|               |                |
| Next Header = | Next Header =  |
|       Routing |       ISO-TP   |
+---------------+----------------+-----------------------
```

3.1. Protocol Classes

The ISO connection-oriented transport protocol [ISO8073] supports five different classes of service. Only one such class, class 4 (TP4), is suitable for use on a connectionless network service such as provided by IPv6. Transport classes 0 through 3 should not carried over an IPv6 network in this manner.

Note that the connectionless transport protocol [ISO8072] has no such restriction. Its PDUs should be carried exactly as described above. There is no conflict inherent in using the same IPv6 Next Header value for both connection-oriented and connectionless protocols. ISO transport implementations can distinguish the two protocols by their different PDU types.

3.2. Maximum TPDU Size

When negotiating a maximum TPDU size, TP4 implementations may consider the services available from the network layer. Unlike IPv4 or CLNP, IPv6 only permits fragmentation by the originating system. TP4 may use its knowledge of the capabilities of the local system to maximize the efficiency of data transfer.

3.2.1. Path MTU Discovery and Fragmentation

If the TP4 implementation can accept Path MTU Discovery [RFC1191] information, and if the TP4 implementation can efficiently invoke the IPv6 fragmentation function, then the TP4 may propose the largest TPDU size and/or preferred maximum TPDU size that the implementation can support.

If, during the life of the connection, IPv6 reports PMTU information to the TP4 implementation, TP4 should adjust its local TPDU size accordingly. Note that the original TPDU (the one which solicited the PMTU) cannot be repacketized; TP4 must instead rely on IPv6 fragmentation for that PDU's retransmission.

3.2.2. No Path MTU Discovery or Fragmentation

If the TP4 implementation cannot accept Path MTU Discovery information from IPv6, or if it cannot efficiently invoke the IPv6 fragmentation function, then TP4 may propose a TPDU size of 512 octets and a preferred maximum TPDU size of 512 octets. These sizes will ensure that TPDUs are no larger than the IPv6 minimum MTU of 576 bytes [IPv6].

3.3. PDU Lifetime

Unlike IPv4 and CLNP, IPv6 nodes are not required to enforce PDU lifetimes. Any transport protocol that relies on the network protocol to limit packet lifetime ought to be upgraded to provide its own mechanisms for detecting and discarding obsolete packets.

3.4. Related work

The carriage of OSI Connectionless Transport Services over UDP is described in [RFC1240], which is currently a Proposed Standard. The present proposal is independent of that one.

4. Security condiderations

Security issues are not specifically addressed in this document, but it is compatible with the IPv6 security mechanisms [security]. Note, however, that when CLNP is tunnelled through IPv6 the IPv6 security mechanisms can at best protect the tunnel, but not the end-to-end CLNP service.

5. References

[IS08072] International Organisation for Standardization, "Transport Service Definition", International Standard 8072, 1987.

[IS08073] International Organisation for Standardization, "Protocol for providing the connection-mode transport service", International Standard 8073 (2nd ed.), 1992.

[RFC1191] Mogul, J., and S. Deering, "Path MTU Discovery", DECWRL and Stanford University, November 1990.

[IS8473] International Organisation for Standardization, "Data communications protocol for providing the connectionless-mode network service", International Standard 8473, 1988.

[IS9542] International Organisation for Standardization, "ISO, "End system to Intermediate system routeing exchange protocol for use in conjunction with the Protocol for providing the connectionless-mode network service (ISO 8473)," International Standard 9542, 1988.

[IS10589] International Organisation for Standardization, "Intermediate system to Intermediate system routeing information exchange protocol for use in conjunction with the Protocol for providing the Connectionless-mode Network Service (ISO 8473)," International Standard 10589, 1992.

[IS10747] International Organisation for Standardization, "Intermediate system to Intermediate system interdomain routeing information exchange protocol for use in conjunction with the Protocol for providing the Connectionless-mode Network Service (ISO 8473)," International Standard 10747, 1993.

[IPv6] The IPv6 base documents, especially S. Deering, R. Hinden, Internet Protocol, Version 6 (IPv6) Specification, work in progress, draft-hinden-ipng-ipv6-spec-01.txt, March 1995.

[RFC1006] Rose, M., and D. Cass, "ISO Transport Service on top of the TCP", STD-35, RFC 1006, Northrop Research and Technology Center, May 1987.

[RFC1070] Hagens, R., Hall, N., and M. Rose, "Use of the Internet as a Subnetwork for Experimentation with the OSI Network Layer", RFC 1070, University of Wisconsin, February 1989.

[RFC1240] Shue, C., Haggerty, W., and K. Dobbins, "OSI Connectionless Transport Services on top of UDP", RFC 1240, Open

Software Foundation, June 1991

[assigned] Reynolds, J., and J. Postel, "Assigned Numbers", STD 2, RFC 1700, USC/Information Sciences Institute, October 1994.

[security] IPv6 security spec, especially, R. Atkinson, "Security Architecture for the Internet Protocol", work in progress, draft-ietf-ipsec-arch-02.txt, May 1995.

6. Authors' Addresses

Brian E. Carpenter
Group Leader, Communications Systems
Computing and Networks Division
European Laboratory for Particle Physics (CERN)
1211 Geneva 23, Switzerland

Phone: +41 22 767-4967
Fax: +41 22 767-7155
Email: brian@dxcoms.cern.ch

Dave Katz
cisco Systems, Inc.
1525 O'Brien Dr.
Menlo Park, CA 94025

Phone: (415) 688-8284
EMail: dkatz@cisco.com

Stephen Thomas
Associate Principal Engineer
AT&T Tridom
840 Franklin Court
Marietta, GA 30067 USA

Phone: (404) 514-3522
Fax: (404) 514-3491
Email: stephen.thomas@tridom.com

Keith Sklower
Computer Science Department
384 Soda Hall, Mail Stop 1776
University of California
Berkeley, CA 94720-1776

Phone: (510) 642-9587
EMail: sklower@CS.Berkeley.EDU

7. Expiration Date of this Draft

December 22nd, 1995