

Network Working Group
Request for Comments: 1888
Category: Experimental

J. Bound
Digital Equipment Corporation
B. Carpenter
CERN
D. Harrington
Digital Equipment Corporation
J. Houldsworth
ICL Network Systems
A. Lloyd
Datacraft Technologies
August 1996

OSI NSAPs and IPv6

Status of this Memo

This memo defines an Experimental Protocol for the Internet community. This memo does not specify an Internet standard of any kind. Discussion and suggestions for improvement are requested. Distribution of this memo is unlimited.

Abstract

This document recommends that network implementors who have planned or deployed an OSI NSAP addressing plan, and who wish to deploy or transition to IPv6, should redesign a native IPv6 addressing plan to meet their needs. However, it also defines a set of mechanisms for the support of OSI NSAP addressing in an IPv6 network. These mechanisms are the ones that MUST be used if such support is required. This document also defines a mapping of IPv6 addresses within the OSI address format, should this be required.

Table of Contents

1. General recommendation on NSAP addressing plans.....	2
2. Summary of defined mechanisms.....	4
3. Restricted NSAPA in a 16-byte IPv6 address for ICD and DCC...4	
3.1 Routing restricted NSAPAs.....	5
4. Truncated NSAPA used as an IPv6 address.....	6
4.1 Routing truncated NSAPAs.....	8
5. Carriage of full NSAPAs in IPv6 destination option.....	9
6. IPv6 addresses inside an NSAPA.....	10
7. Security Considerations.....	11
Acknowledgements.....	11
References.....	12
Annex A: Summary of NSAP Allocations.....	13
Annex B: Additional Rationale.....	14
Authors' Addresses.....	16

1. General recommendation on NSAP addressing plans

This recommendation is addressed to network implementors who have already planned or deployed an OSI NSAP addressing plan for the usage of OSI CLNP [IS8473] according to the OSI network layer addressing plan [IS8348] using ES-IS and IS-IS routing [IS9542, IS10589]. It recommends how they should adapt their addressing plan for use with IPv6 [RFC1883].

The majority of known CLNP addressing plans use either the Digital Country Code (DCC) or the International Code Designator (ICD) formats defined in [IS8348]. A particular example of this is the US Government OSI Profile Version 2 (GOSIP) addressing plan [RFC1629]. The basic NSAP addressing scheme and current implementations are summarised in Annex A.

[IS8348] specifies a maximum NSAPA (NSAP address) size of 20 bytes and some network implementors have designed address allocation schemes which make use of this 20 byte address space.

Other NSAP addressing plans have been specified by the ITU-T for public data services, such as X.25 and ISDN, and these can also have addresses up to 20 bytes in length.

The general recommendation is that implementors SHOULD design native IPv6 addressing plans according to [RFC1884], but doing so as a natural re-mapping of their CLNP addressing plans. While it is impossible to give a general recipe for this, CLNP addresses in DCC or ICD format can normally be split into two parts: the high order part relating to the network service provider and the low order part relating to the user network topology and host computers.

For example, in some applications of US GOSIP the high order part is the AFI, ICD, DFI, AA and RD fields, together occupying 9 bytes. The low order part is the Area and ID fields, together occupying 8 bytes. (The selector byte and the two reserved bytes are not part of the addressing plan.) Thus, in such a case, the high-order part could be replaced by the provider part of an IPv6 provider-based addressing plan. An 8-byte prefix is recommended for this case and [RFC1884] MUST be followed in planning such a replacement. The low order part would then be mapped directly in the low-order half of the IPv6 address space, and user site address plans are unchanged. A 6-byte ID field, exactly as used in US GOSIP and other CLNP addressing plans, will be acceptable as the token for IPv6 autoconfiguration [RFC1971].

Analogous rules would be applied for other CLNP addressing plans similar to US GOSIP, which is used only as a well known example.

Three warnings must be carefully considered in every case:

1. The ES-IS/IS-IS model employs a routing hierarchy down to the Area level, but not all end systems in an Area need to be in the same physical subnet (on the same "wire" or "link"). IS routers on different links within a given Area exchange information about the end systems they can each reach directly. In contrast, the IPv6 routing model extends down to the subnet level and all hosts in the same subnet are assumed to be on the same link. In mapping a CLNP addressing plan into IPv6 format, without changing the physical topology, it may be necessary to add an extra level of hierarchy to cope with this mismatch. In other words, the Area number cannot blindly be mapped as a subnet number, unless the physical network topology corresponds to this mapping.

2. It is highly desirable that subnet addresses can be aggregated for wide area routing purposes, to minimise the size of routing tables. Thus network implementors should ensure that the address prefix used for all their subnets is the same, regardless of whether a particular subnet is using a pure IPv6 addressing scheme or one derived from a CLNP scheme as above.

3. Some hosts have more than one physical network interface. In the ES-IS model, an end system may have more than one NSAP address, each of which identifies the host as a whole. Such an end system with more than one physical interface may be referenced by any one of the NSAPs, and reached via any one of the physical connections. In the IPv6 model, a host may have multiple IPv6 addresses per interface, but each of its physical interfaces must have its own unique addresses. This restriction must be applied when mapping an NSAP addressing plan into an IPv6 addressing plan for such hosts.

This document does not address the issues associated with migrating the routing protocols used with CLNP (ES-IS or IS-IS) and transition of their network infrastructure.

2. Summary of defined mechanisms

This document defines four distinct mechanisms. All of 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.

1. Restricted NSAPA mapping into 16-byte IPv6 address
2. Truncated NSAPA for routing, full NSAPA in IPv6 option
3. Normal IPv6 address, full NSAPA in IPv6 option
4. IPv6 address carried as OSI address

To clarify the relationship between the first three mechanisms, note that:

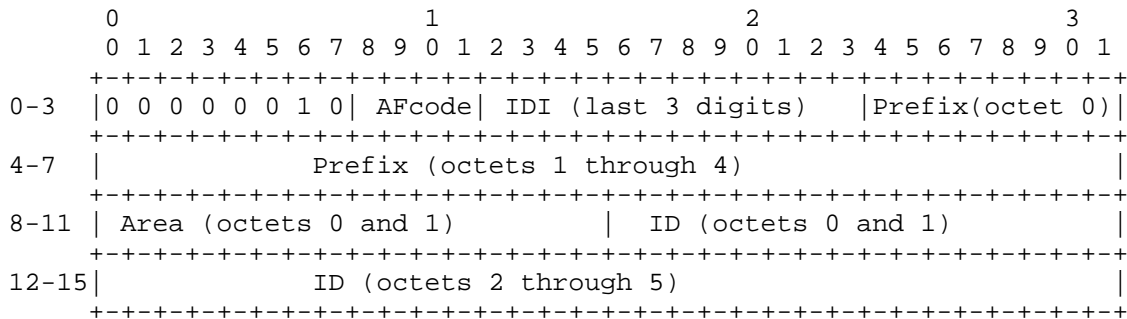
If the first byte of an IPv6 address is hexadecimal 0x02 (binary 00000010), then the remaining 15 bytes SHALL contain a restricted NSAPA mapped as in Chapter 3 below. The term "restricted" is used to indicate that this format is currently restricted to a subset of the ICD and DCC formats.

If the first byte of an IPv6 address is hexadecimal 0x03 (binary 00000011), then the remaining 15 bytes SHALL contain a truncated NSAPA as described in Chapter 4 below. EITHER a destination option containing the complete NSAPA of any format, as described in Chapter 5 below, OR an encapsulated CLNP packet, SHALL be present.

With any other format of IPv6 address, a destination option containing a complete NSAPA, as defined in Chapter 5 below, MAY be present.

3. Restricted NSAPA in a 16-byte IPv6 address for ICD and DCC

Some organizations may decide for various reasons not to follow the above general recommendation to redesign their addressing plan. They may wish to use their existing OSI NSAP addressing plan unchanged for IPv6. It should be noted that such a decision has serious implications for routing, since it means that routing between such organizations and the rest of the Internet is unlikely to be optimised. An organization using both native IPv6 addresses and NSAP addresses for IPv6 would be likely to have inefficient internal routing. Nevertheless, to cover this eventuality, the present document defines a way to map a subset of the NSAP address space into the IPv6 address space. The mapping is algorithmic and reversible within this subset of the NSAP address space.



The AFcode nibble is overloaded, and encoded as follows

- 0000-1001 (0-9 decimal) Implied AFI value is 47 (ICD)
AFcode is first BCD digit of the ICD
IDI is last three BCD digits of the ICD
- 1010 (hex. A) Implied AFI value is 39 (DCC)
IDI is the three BCD digits of the DCC
- 1011-1111 (hex. B-F) Reserved, not to be used.

The NSEL octet is not included. It is of no use for TCP and UDP traffic. In any case where it is needed, the mechanism described in the next chapter should be used.

The longest CLNP routing prefixes known to be in active use today are 5 octets (subdivided into AA and RD fields in US GOSIP version 2). Thus the semantics of existing 20-octet NSAPAs can be fully mapped. DECnet/OSI (Registered Trade Mark) address semantics are also fully mapped.

It is expected that hosts using restricted NSAPAs could be configured using IPv6 auto-configuration [RFC1971], and that they could use normal IPv6 neighbour discovery mechanisms [RFC1970].

Restricted NSAPAs, assuming that they can be fully routed using IPv6 routing protocols, may be used in IPv6 routing headers.

3.1 Routing restricted NSAPAs

As mentioned in Chapter 1, there is a mismatch between the OSI or GOSIP routing model and the IPv6 routing model. Restricted NSAPAs can be routed hierarchically down to the Area level but must be flat-routed within an Area. Normal IPv6 addresses can be routed

hierarchically down to physical subnet (link) level and only have to be flat-routed on the physical subnet.

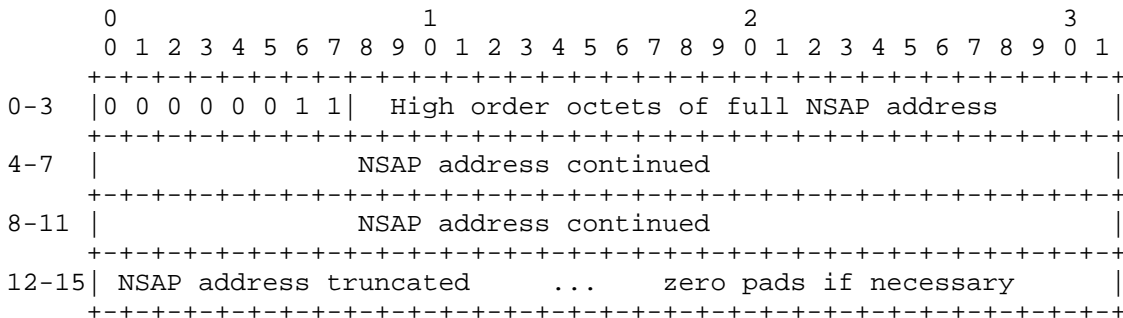
Thus, packets whose destination address is a restricted NSAPA can be routed using any normal IPv6 routing protocol only as far as the Area. If the Area contains more than one physical subnet reached by more than one router, no IPv6 routing protocol can route the packet to the correct final router. There is no solution to this problem within the existing IPv6 mechanisms. Presumably a flooding algorithm, or a suitably adapted implementation of ES-IS, could solve this problem.

In the absence of such a routing protocol, either the Area number must be hierarchically structured to correspond to physical subnets, or each Area must be limited to one physical subnet.

It is necessary in an IPv6 network that routes may be aggregated to minimise the size of routing tables. If a subscriber is using both normal IPv6 addresses [RFC1884] and restricted NSAPAs, these two types of address will certainly not aggregate with each other, since they differ from the second most significant bit onwards. This means that there may be a significant operational penalty for using both types of address with currently known routing technology.

4. Truncated NSAPA used as an IPv6 address

An NSAP address contains routing information (e.g. Routing Domain and area/subnet identifiers) in the form of the Area Address (as defined in [IS10589]). The format and length of this routing information are typically compatible with a 16 byte IPv6 address, and may be represented as such using the following format:



If appropriate, when used as a destination IPv6 address, the truncated NSAPA may be interpreted as an IPv6 anycast address. An anycast address may be used to identify either an IPv6 node, or potentially even an OSI End System or Intermediate System. For

example, it might be configured to identify the endpoints of a CLNP tunnel, or it might identify a particular OSI capable system in a particular subnet.

If a truncated NSAPA is used as a source address, it must be interpreted as a unicast address and must therefore be uniquely assigned within the IPv6 address space.

If a truncated NSAPA is used as either the source or destination IPv6 address (or both), EITHER an NSAPA destination option OR an encapsulated CLNP packet MUST be present. It is the responsibility of the destination system to take the appropriate action for each IPv6 packet received (e.g. forward, decapsulate, discard) and, if necessary, return to the originating host an appropriate ICMP error message.

If the truncated NSAPA is used to identify a router, and an NSAPA destination option is present, then it is the responsibility of that router to forward the complete IPv6 packet to the appropriate host based upon the Destination NSAP field in the NSAPA option. This forwarding process may be based upon static routing information (i.e. a manual mapping of NSAPs to IPv6 unicast addresses), or it may be gathered in an automated fashion analogous to the ES-IS mechanism, perhaps using extensions to the Neighbor Discovery protocol [RFC1970]. The details of such a mechanism are beyond the scope of this document.

This document does not restrict the formats of NSAP address that may be used in truncated NSAPAs, but it is apparent that binary ICD or DCC formats will be much easier to accommodate in an IPv6 routing infrastructure than the other formats defined in [IS8348].

It is not expected that IPv6 autoconfiguration [RFC1971] and discovery [RFC1970] will work unchanged for truncated NSAPAs.

Truncated NSAPAs are not meaningful within IPv6 routing headers, and there is no way to include full NSAPAs in routing headers.

If a packet whose source address is a truncated NSAPA causes an ICMP message to be returned for whatever reason, this ICMP message may be discarded rather than being returned to the true source of the packet.

4.1 Routing truncated NSAPAs

This is a grey area. If the truncated NSAPA retains a hierarchical structure, it can be routed like a restricted NSAPA, subject to the same problem concerning the mismatch between Areas and subnets. If possible, in the case of a GOSIP-like NSAPA, it should be truncated immediately after the Area number. In this case the routing considerations will be similar to those for restricted NSAPAs, except that final delivery of the packet will depend on the last IPv6 router being able to interpret the NSAPA destination option (or an encapsulated CLNP packet).

In the general case, nothing can be said since the NSAPA could have almost any format and might have very little hierarchical content after truncation. There may be many cases in which truncated NSAPAs cannot be routed across large regions of the IPv6 network.

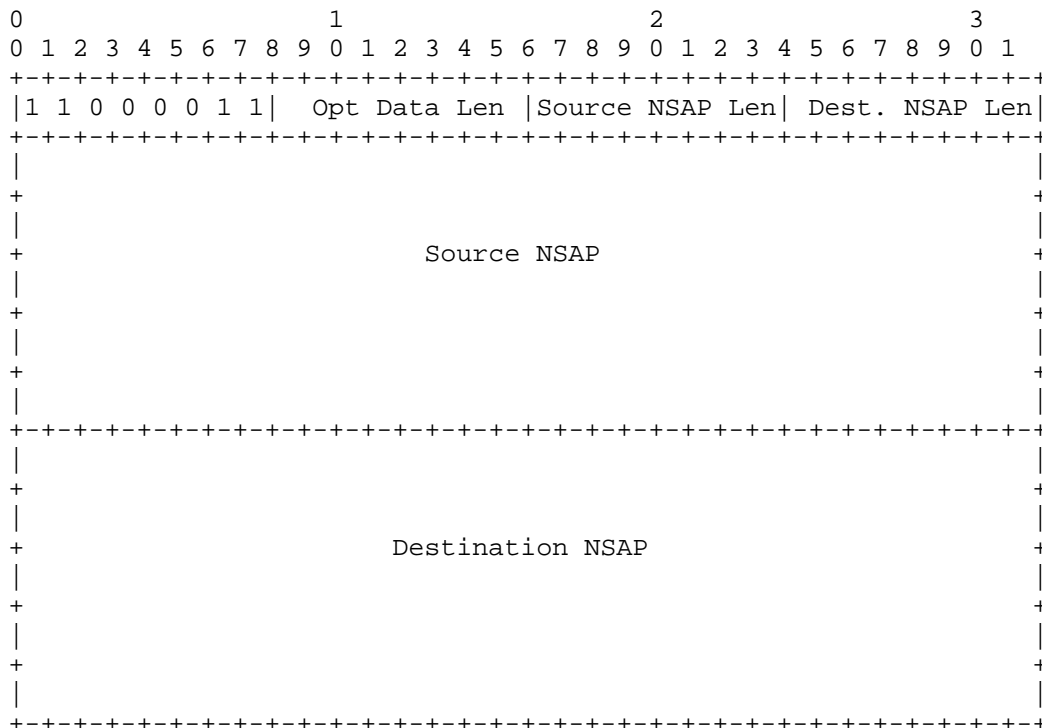
The situation for route aggregation is similar to that described in Section 3.1 as long as the truncated NSAPAs have ICD or DCC format. However, if arbitrary NSAPAs are used nothing can be predicted about route aggregation and we must assume that it will be poor.

5. Carriage of full NSAPAs in IPv6 destination option

In the case of a truncated NSAPA used as an IPv6 address other than for a CLNP tunnel, the full NSAPA must be carried in a destination option. Any format defined in [IS8348] is allowed.

The NSAPA destination option is illustrated below. It has no alignment requirement.

The option type code is 11-0-00011 = 195 decimal.



The length fields are each one octet long and are expressed in octets. The destination node should check the consistency of the length fields (Option Data Length = Source NSAP Length + Dest. NSAP Length +2). In case of inconsistency the destination node shall discard the packet and send an ICMP Parameter Problem, Code 2, message to the packet's source address, pointing to the Option Data Length field.

The boundary between the source NSAP and the destination NSAP is simply aligned on an octet boundary. With standard 20 octet NSAPs the total option length is 44 bytes and the Option Data Length is 42.

The NSAP encodings follow [IS8348] exactly.

If this option is used, both end systems concerned SHOULD use NSAP addresses. In the exceptional case that only one of the end systems uses NSAP addresses, the NSAP Length field of the other SHALL be set to zero in the NSAP destination option.

This destination option is used in two cases. Firstly, an IPv6 source node using normal IPv6 addresses (unicast address or anycast address) MAY supply an NSAP destination option header for interpretation by the IPv6 destination node. Secondly, an IPv6 node MAY use a truncated NSAP address in place of a normal IPv6 address.

IPv6 nodes are not required to implement this option, except for nodes using truncated NSAPAs other than for CLNP tunnels.

6. IPv6 addresses inside an NSAPA

If it is required, for whatever reason, to embed an IPv6 address inside a 20-octet NSAP address, then the following format MUST be used.

A specific possible use of this embedding is to express an IP address within the ATM Forum address format. Another possible use would be to allow CLNP packets that encapsulate IPv6 packets to be routed in a CLNP network using the IPv6 address architecture. Several leading bytes of the IPv6 address could be used as a CLNP routing prefix.

The first three octets are an IDP in binary format, using the AFI code in the process of being allocated to the IANA. The AFI value provisionally allocated is 35, but this requires a formal modification to [IS8348]. The encoding format is as for AFI value 47 [IS8348]. The third octet of the IDP is known as the ICP (Internet Code Point) and its value must be zero. All other values are reserved for allocation by the IANA.

Thus an AFI value of 35 with an ICP value of zero means that "this NSAPA embeds a 16 byte IPv6 address".

The last octet is a selector. To maintain compatibility with both NSAP format and IPv6 addressing, this octet must be present, but it has no significance for IPv6. Its default value is zero.

	0									1									2									3								
	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1				
0-3	AFI = 35									ICP = 0000									IPv6 (byte 0)																	
4-7	IPv6 (bytes 1-4)																																			
8-11	IPv6 (bytes 5-8)																																			
12-15	IPv6 (bytes 9-12)																																			
16-19	IPv6 (bytes 13-15)																		0 0 0 0 0 0 0 0																	

Theoretically this format would allow recursive address embedding.

However, this is considered dangerous since it might lead to routing table anomalies or to loops (compare [RFC1326]). Thus embedded IPv6 address MUST NOT have the prefixes 0x02 or 0x03, and an NSAPA with the IANA AFI code MUST NOT be embedded in an IPv6 header.

An NSAPA with the IANA AFI code and ICP set to zero is converted to an IPv6 address by stripping off the first three and the twentieth octets. All other formats of NSAPA are handled according to the previous Chapters of this document.

7. Security Considerations

Security issues are not specifically addressed in this document, but it is compatible with the IPv6 security mechanisms [RFC1825].

Acknowledgements


The authors are pleased to acknowledge the suggestions and comments of Ross Callon, Richard Collella, Steve Deering, Dirk Fieldhouse, Joel Halpern, Denise Heagerty, Cyndi Jung, Yakov Rekhter, and members of the former TUBA and current IPNG working groups of the IETF. The support of Scott Bradner and Allison Mankin of the IESG was essential.

Herb Bertine, Alan Chambers, Dave Marlow, and Jack Wheeler were all active in arranging the AFI allocation by ISO/IEC JTC1/SC6.

References

- [IS8473] Data communications protocol for providing the connectionless-mode network service, ISO/IEC 8473, 1988.
- [IS8348] Annex A, Network Layer Addressing, and Annex B, Rationale for the material in Annex A, of ISO/IEC 8348, 1993 (identical to CCITT Recommendation X.213, 1992).
- [IS10589] Intermediate system to intermediate system intra-domain-routeing routine information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode Network Service (ISO 8473), ISO 10589, 1992.
- [IS9542] End system to Intermediate system routeing exchange protocol for use in conjunction with the Protocol for providing the connectionless-mode network service (ISO 8473), ISO 9542, 1988.
- [RFC1629] Colella, R., Callon, R., Gardner, E., and Y. Rekhter, "Guidelines for OSI NSAP Allocation in the Internet", RFC 1629, May 1994.
- [RFC1326] Tsuchiya, P., "Mutual Encapsulation Considered Dangerous", RFC 1326, May 1992.
- [RFC1883] Deering, S., and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 1883, December 1995.
- [RFC1884] Hinden, R., and S. Deering, "IP Version 6 Addressing Architecture", RFC 1884, December 1995.
- [RFC1971] Thompson, S., and T. Narten, "IPv6 Stateless Address Autoconfiguration", RFC1971, August 1996.
- [RFC1970] Narten, T., Nordmark, E., and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", RFC1970, August 1996.
- [RFC1825] Atkinson, R., "Security Architecture for the Internet Protocol", RFC 1825, August 1995.

Annex A: Summary of NSAP Allocations



The Initial Domain Part (IDP) is split into Authority and Format Identifier (AFI) followed by the Initial Domain Identifier (IDI). This combination is followed by the Domain Specific Part and allocation within that part is domain specific.

The following is a summary of current allocations:

ISO DCC Scheme

AFI = decimal 38 or binary 39 = ISO Data Country Code Scheme. IDI = 3 decimal or binary digits specifying the country. ISO allocate the country codes. The DSP is administered by the standards authority for each country. In the UK, the British Standards Institution have delegated administration to the Federation of Electronics Industries - FEI

The UK DSP is split into a single digit UK Format Indicator (UKFI) which indicates large, medium or small organisation rather like IP addressing and a UK Domain Identifier (UKDI). Using binary coded decimal examples only (there are binary equivalents):

UKFI = 0 is reserved UKFI = 1, UKDI = nnn, UK Domain Specific Part = 31 digits. UKFI = 2, UKDI = nnnnn, UKDSP = 29 digits max. UKFI = 3, UKDI = nnnnnnnn, UKDSP = 26 digits max.

UKFI = 4 to 9 reserved

The UK Government have been allocated a UKDI in the UKFI = 1 (large organisation) format and have specified the breakdown of the Government Domain Specific Part with sub domain addresses followed by a station ID (which could be a MAC address) and a selector (which could be a TSAP selection).

ITU-T X.121

AFI = decimal 36 or 52, binary 37 or 53 indicates that the IDI is a 14 digit max X.121 International Numbering Plan address (prefix, 3 digit Data Country Code, dial up data network number). The full X.121 address indicates who controls the formatting of the DSP.

ITU-T F.69

AFI = 40,54 or binary 41,55 indicates that the IDI is a telex number up to 8 digits long.

ITU-T E.163

AFI = 42,56 or binary 43,57 indicates that the IDI is a normal telephone number up to 12 digits long.

ITU-T E.164

AFI = 44,58 or binary 45,59 indicates that the IDI is an ISDN number up to 15 digits long.

ISO 6523-ICD

AFI = 46 or binary 47 indicates that the IDI is an International Code Designator allocated according to ISO 6523. You have to be a global organisation to get one of these. The Organisation to which the ISO 6523 designator is issued specifies the DSP allocation.

Annex B: Additional Rationale

This annex is intended to give additional rationale, motivation and justification for the support of NSAPAs in an IPv6 network.

There are several models for OSI-IPv6 convergence, of which address mapping is only one. The other models can be identified as

1. Dual stack coexistence, in which a CLNP network and an IPv6 network exist side by side indefinitely using multiprotocol routers.
2. CLNP tunnels over IPv6.
3. OSI transport over IPv6.
4. OSI transport over UDP.
5. OSI transport over TCP (compare RFC 1006)

The present model is more fundamental, as it attempts to unify and reconcile the OSI and IPv6 addressing and routing schemes, and replace CLNP by IPv6 at the network level. The rationale for this choice is to preserve investment in NSAPA allocation schemes, and to open the door for peer-to-peer routing models between IPv6 and bearer services (such as ATM) using NSAPA addressing. It should be noted

that such peer-to-peer models are contentious at the time of writing, but in any case a consistent address mapping is preferable to multiple mappings.

In addition to their use to retain an existing addressing plan, certain other uses of restricted NSAPAs could be envisaged. They could be used as an intermediate addressing plan for a network making a transition from CLNP to IPv6. They could be used in a header translation scheme for dynamic translation between IPv6 and CLNP. They could be used to allow CLNP and IPv6 traffic to share the same routing architecture within an organization ("Ships in the Day").

It should be noted that the use of full NSAPA addresses in end systems impacts many things. The most obvious are the API and DNS. If applications are to work normally, everything that has to be modified to cope with IPv6 addresses has to be further modified for full NSAPAs. The mechanisms defined in the present document are only a small part of the whole.

A destination option was chosen to carry full NSAPAs, in preference to a dedicated extension header. In the case of an extension header, all IPv6 nodes would have needed to understand its syntax merely in order to ignore it. In contrast, intermediate nodes can ignore the destination option without any knowledge of its syntax. Thus only nodes interested in NSAPAs need to know anything about them.

Thus we end up with two classes of IPv6 nodes:

1. Nodes knowing only about 16 byte addresses (including restricted NSAPAs, which behave largely like any other IPv6 addresses).
2. Nodes also knowing about 20 byte NSAPAs, either as an extension of the IPv6 address space or as the CLNP address space. In either case, regions of the network containing such nodes are connected to each other by unicast or anycast tunnels through the 16 byte address space. Routing, system configuration, and neighbour discovery in the NSAPA regions are outside the scope of the normal IPv6 mechanisms.

Authors' Addresses

Jim Bound
Member Technical Staff
Network Operating Systems
Digital Equipment Corporation
110 Spitbrook Road, ZK03-3/U14
Nashua, NH 03062
Phone: (603) 881-0400
Fax: (603) 881-0120
Email: bound@zk3.dec.com

Brian E. Carpenter
Group Leader, Communications Systems
Computing and Networks Division
CERN
European Laboratory for Particle Physics
1211 Geneva 23, Switzerland
Phone: +41 22 767-4967
Fax: +41 22 767-7155
Telex: 419000 cer ch
Email: brian@dxcoms.cern.ch

Dan Harrington
Digital Equipment Corp.
550 King Street (LKG2-2/Q9)
Littleton, MA 01460
Phone: (508) 486-7643
Email: dan@netrix.lkg.dec.com

Jack Houldsworth
ICL Network Systems
Cavendish Road
Stevenage
Herts
UK SG1 4BQ
Phone- ICL: +44 438 786112
Home: +44 438 352997
Fax: +44 438 786150
Email: j.houldsworth@ste0906.wins.icl.co.uk

Alan Lloyd
Datacraft Technologies
252 Maroondah Highway
Mooroolbark 3138
Victoria Australia
Phone: +61 3 727 9222
Fax: +61 3 727 1557
Email: alan.lloyd@datacraft.com.au

X.400- G=alan;S=lloyd;O=dcthq;P=datacraft;A=telememo;C=au