

## Assignment of System Identifiers for TUBA/CLNP Hosts

### Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard. Distribution of this memo is unlimited.

### Abstract

This document describes conventions whereby the system identifier portion of an RFC 1237 style NSAP address may be guaranteed uniqueness within a routing domain for the purpose of autoconfiguration in TUBA/CLNP internets. The mechanism is extensible and can provide a basis for assigning system identifiers in a globally unique fashion.

### Introduction

This memo specifies methods for assigning a 6 octet system identifier portion of the OSI NSAP address formats described in "Guidelines for OSI NSAP Allocation in the Internet" [1], in a fashion that ensures that the ID is unique within a routing domain. It also recommends methods for assigning system identifiers having lengths other than 6 octets. The 6 octet system identifiers recommended in this RFC are assigned from 2 globally administered spaces (IEEE 802 or "Ethernet", and IP numbers, administered by the Internet Assigned Numbers Authority, IANA).

At this time, the primary purpose for assuring uniqueness of system identifiers is to aid in autoconfiguration of NSAP addresses in TUBA/CLNP internets [2]. The guidelines in this paper also establish an initial framework within which globally unique system identifiers, also called endpoint identifiers, may be assigned.

### Acknowledgments

Many thanks to Radia Perlman, Allison Mankin, and Ross Callon of for their insights and assistance. Thanks also to the Ethernet connector to my MAC, which conveniently and quite inobtrusively fell out, enabling me to get an entire day's worth of work done without email interruptions.

## 1. Background

The general format of OSI network service access point (NSAP) addresses is illustrated in Figure 1.

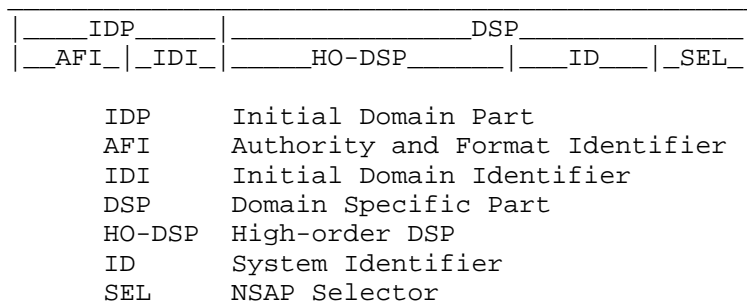


Figure 1: OSI NSAP Address Structure.

The recommended encoding and allocation of NSAP addresses in the Internet is specified in RFC 1237. RFC 1237 makes the following statements regarding the system identifier (ID) field of the NSAPA:

1. the ID field may be from one to eight octets in length
2. the ID must have a single known length in any particular routing domain
3. the ID field must be unique within an area for ESs and level 1 ISS, and unique within the routing domain for level 2 ISS.
4. the ID field is assumed to be flat

RFC 1237 further indicates that, within a routing domain that conforms to the OSI intradomain routing protocol [3] the lower-order octets of the NSAP should be structured as the ID and SEL fields shown in Figure 1 to take full advantage of intradomain IS-IS routing. (End systems with addresses which do not conform may require additional manual configuration and be subject to inferior routing performance.)

Both GOSIP Version 2 (under DFI-80h, see Figure 2a) and ANSI DCC NSAP addressing (Figure 2b) define a common DSP structure in which the system identifier is assumed to be a fixed length of 6 octets.

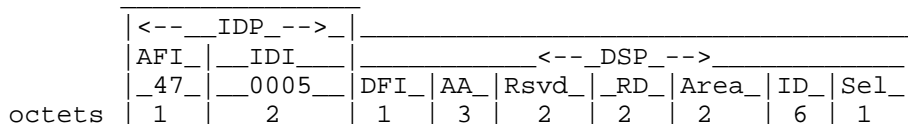
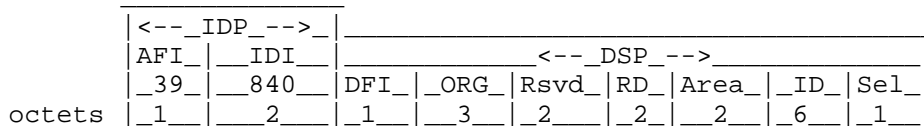


Figure 2 (a): GOSIP Version 2 NSAP structure.



- IDP Initial Domain Part
- AFI Authority and Format Identifier
- IDI Initial Domain Identifier
- DSP Domain Specific Part
- DFI DSP Format Identifier
- ORG Organization Name (numeric form)
- Rsvd Reserved
- RD Routing Domain Identifier
- Area Area Identifier
- ID System Identifier
- SEL NSAP Selector

Figure 2(b): ANSI NSAP address format for DCC=840

## 2. Autoconfiguration

There are provisions in OSI for the autoconfiguration of area addresses. OSI end systems may learn their area addresses automatically by observing area address identified in the IS-Hello packets transmitted by routers using the ISO 9542 End System to Intermediate System Routing Protocol, and may then insert their own system identifier. (In particular, RFC 1237 explains that when the ID portion of the address is assigned using IEEE style 48-bit identifiers, an end system can reconfigure its entire NSAP address automatically without the need for manual intervention.) End systems that have not been pre-configured with an NSAPA may also request an address from an intermediate system their area using [5].

### 2.1 Autoconfiguration and 48-bit addresses

There is a general misassumption that the 6-octet system identifier must be a 48-bit IEEE assigned (ethernet) address. Generally speaking, autoconfiguration does not rely on the use of a 48-bit ethernet style address; any system identifier that is globally

administered and is unique will do. The use of 48-bit/6 octet system identifiers is "convenient...because it is the same length as an 802 address", but more importantly, choice of a single, uniform ID length allows for "efficient packet forwarding", since routers won't have to make on the fly decisions about ID length (see [6], pages 156-157). Still, it is not a requirement that system identifiers be 6 octets to operate the the IS-IS protocol, and IS-IS allows for the use of IDs with lengths from 1 to 8 octets.

3. System Identifiers for TUBA/CLNP

Autoconfiguration is a desirable feature for TUBA/CLNP, and is viewed by some as "essential if a network is to scale to a truly large size" [6].

For this purpose, and to accommodate communities who do not wish to use ethernet style addresses, a generalized format that satisfies the following criteria is desired:

- o the format is compatible with installed end systems complying to RFC 1237
- o the format accommodates 6 octet, globally unique system identifiers that do not come from the ethernet address space
- o the format accommodates globally unique system identifiers having lengths other than 6 octets

The format and encoding of a globally unique system identifier that meets these requirements is illustrated in Figure 3:

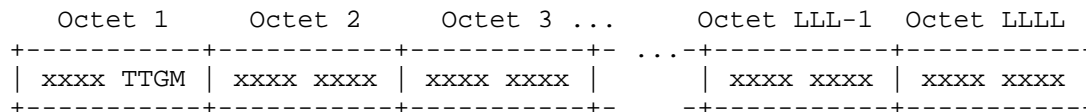


Figure 3. General format of the system identifier

3.1 IEEE 802 Form of System Identifier

The format is compatible with globally assigned IEEE 802 addresses, since it carefully preserves the semantics of the global/local and group/individual bits. Octet 1 identifies 2 qualifier bits, G and M, and a subtype (TT) field whose semantics are associated with the qualifier bits. When a globally assigned IEEE 802 address is used as a system identifier, the qualifier bit M, representing the multicast/unicast bit, must always be set to zero to denote a unicast address. The qualifier bit G may be either 0 or 1, depending on

whether the individual address is globally or locally assigned. In these circumstances, the subtype bits are "don't care", and the system identifier shall be interpreted as a 48-bit, globally unique identifier assigned from the IEEE 802 committee (an ethernet address). The remaining bits in octet 1, together with octets 2 and 3 are the vendor code or OUI (organizationally unique identifier), as illustrated in Figure 4. The ID is encoded in IEEE 802 canonical form (low order bit of low order hex digit of leftmost octet is the first bit transmitted).

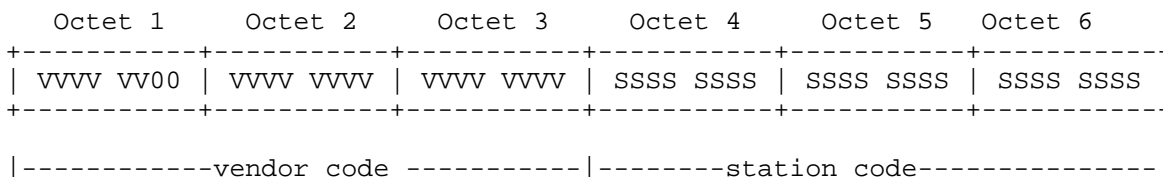


Figure 4. IEEE 802 form of system identifier

#### 4. Embedded IP Address as System Identifier

To distinguish 48-bit IEEE 802 addresses used as system identifiers from other forms of globally administered system identifiers, the qualifier bit M shall be set to 1. The correct interpretation of the M bit set to 1 should be, "this can't be an IEEE 802 multicast address, since use of multicast addresses is by convention illegal, so it must be some other form of system identifier". The subtype (TT) bits illustrated in Figure 3 thus become relevant.

When the subtype bits (TT) are set to a value of 0, the system identifier contains an embedded IP address. The remainder of the 48-bit system identifier is encoded as follows. The remaining nibble in octet 1 shall be set to zero. Octet 2 is reserved and shall be set to a pre-assigned value (see Figure 5). Octets 3 through 6 shall contain a valid IP address, assigned by IANA. Each octet of the IP address is encoded in binary, in internet canonical form, i.e., the leftmost bit of the network number first.

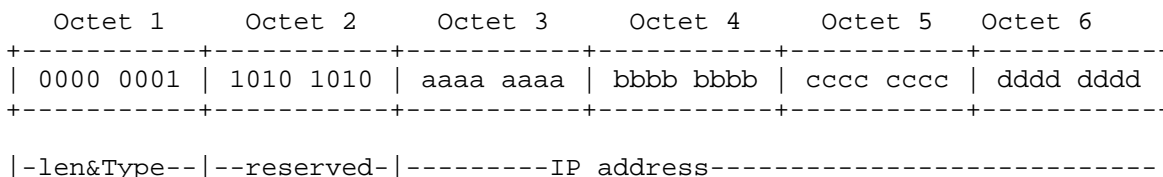


Figure 5. Embedded IP address as system identifier

As an example, the host "eve.bellcore.com = 128.96.90.55" could retain its IP address as a system identifier in a TUBA/CLNP network. The encoded ID is illustrated in Figure 6.

Octet 1	Octet 2	Octet 3	Octet 4	Octet 5	Octet 6
0000 0001	1010 1010	1000 0000	0110 0000	0101 1010	0011 0111
-len&Type--		--reserved-		-----IP address-----	

Figure 6. Example of IP address encoded as ID

## H 2 "Other forms of System Identifiers"

To allow for the future definition of additional 6-octet system identifiers, the remaining subtype values are reserved.

It is also possible to identify system identifiers with lengths other than 6 octets. Communities who wish to use 8 octet identifiers (for example, embedded E.164 international numbers for the ISDN ERA) must use a GOSIP/ANSI DSP format that allows for the specification of 2 additional octets in the ID field, perhaps at the expense of the "Rsvd" fields; this document recommends that a separate Domain Format Indicator value be assigned for such purposes; i.e., a DFI value that is interpreted as saying, among other things, "the system identifier encoded in this DSP is 64-bits/8 octets. The resulting ANSI/GOSIP DSP formats under such circumstances are illustrated in Figure 7:

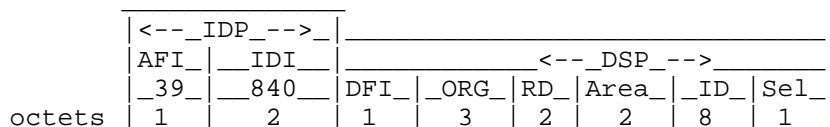
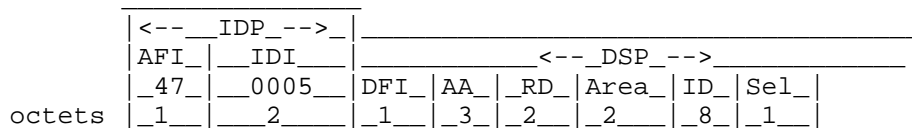


Figure 7a: ANSI NSAP address format for DCC=840, DFI=foo



IDP Initial Domain Part  
 AFI Authority and Format Identifier  
 IDI Initial Domain Identifier  
 DSP Domain Specific Part  
 DFI DSP Format Identifier  
 AA Administrative Authority  
 RD Routing Domain Identifier  
 Area Area Identifier  
 ID System Identifier  
 SEL NSAP Selector

Figure 7b: GOSIP Version 2 NSAP structure, DFI=bar

Similar address engineering can be applied for those communities who wish to have shorter system identifiers; have another DFI assigned, and expand the reserved field.

## 5. Conclusions

This proposal should debunk the "if it's 48-bits, it's gotta be an ethernet address" myth. It demonstrates how IP addresses may be encoded within the 48-bit system identifier field in a compatible fashion with IEEE 802 addresses, and offers guidelines for those who wish to use system identifiers other than those enumerated here.

## 6. References

- [1] Callon, R., Gardner, E., and R. Colella, "Guidelines for OSI NSAP Allocation in the Internet", RFC 1237, NIST, Mitre, DEC, June 1991.
- [2] Callon, R., "TCP and UDP with Bigger Addresses (TUBA), A Simple Proposal for Internet Addressing and Routing", RFC 1347, DEC, June 1992.
- [3] ISO, "Intradomain routing protocol for use in conjunction with ISO 8473, Protocol for providing the OSI connectionless network service", ISO 10589.
- [4] ISO, "End-system and intermediate-system routing protocol for use in conjunction with ISO 8473, Protocol for providing the OSI connectionless network service", ISO 9542.
- [5] ISO, "End-system and intermediate-system routing protocol for use in conjunction with ISO 8473, Protocol for providing the OSI connectionless network service. Amendment 1: Dynamic Discovery of OSI NSAP Addresses End Systems", ISO 9542/DAM1.
- [6] Perlman, R., "Interconnections: Bridges and Routers", Addison-Wesley Publishers, Reading, MA. 1992.

## 7. Security Considerations

Security issues are not discussed in this memo.

## 8. Author's Address

David M. Piscitello  
Bell Communications Research  
NVC 1C322  
331 Newman Springs Road  
Red Bank, NJ 07701

EMail: dave@mail.bellcore.com