

Network Working Group
Request for Comments: 1629
Obsoletes: 1237
Category: Standards Track

R. Colella
NIST
R. Callon
Wellfleet
E. Gardner
Mitre
Y. Rekhter
T.J. Watson Research Center, IBM Corp.
May 1994

Guidelines for OSI NSAP Allocation in the Internet

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Abstract

CLNP is currently being deployed in the Internet. This is useful to support OSI and DECnet(tm) traffic. In addition, CLNP has been proposed as a possible IPng candidate, to provide a long-term solution to IP address exhaustion. Required as part of the CLNP infrastructure are guidelines for network service access point (NSAP) address assignment. This paper provides guidelines for allocating NSAP addresses in the Internet.

The guidelines provided in this paper have been the basis for initial deployment of CLNP in the Internet, and have proven very valuable both as an aid to scaling of CLNP routing, and for address administration.

Table of Contents

Section 1. Introduction	4
Section 2. Scope	5
Section 3. Background	7
Section 3.1 OSI Routing Standards	7
Section 3.2 Overview of IS-IS (ISO/IEC 10589)	8
Section 3.3 Overview of IDRP (ISO/IEC 10747)	12
Section 3.3.1 Scaling Mechanisms in IDRP	14
Section 3.4 Requirements of IS-IS and IDRP on NSAPs ...	15
Section 4. NSAPs and Routing	16
Section 4.1 Routing Data Abstraction	16
Section 4.2 NSAP Administration and Efficiency	19
Section 5. NSAP Administration and Routing in the In- ternet	21
Section 5.1 Administration at the Area	23
Section 5.2 Administration at the Subscriber Routing Domain	24
Section 5.3 Administration at the Provider Routing Domain	24
Section 5.3.1 Direct Service Providers	25
Section 5.3.2 Indirect Providers	26
Section 5.4 Multi-homed Routing Domains	26
Section 5.5 Private Links	31
Section 5.6 Zero-Homed Routing Domains	33
Section 5.7 Address Transition Issues	33
Section 6. Recommendations	36
Section 6.1 Recommendations Specific to U.S. Parts of the Internet	37
Section 6.2 Recommendations Specific to European Parts of the Internet	39
Section 6.2.1 General NSAP Structure	40
Section 6.2.2 Structure of the Country Domain Part	40
Section 6.2.3 Structure of the Country Domain Specific Part	41
Section 6.3 Recommendations Specific to Other Parts of the Internet	41
Section 6.4 Recommendations for Multi-Homed Routing Domains	41
Section 6.5 Recommendations for RDI and RDCI assign- ment	42
Section 7. Security Considerations	42
Section 8. Authors' Addresses	43
Section 9. Acknowledgments	43
Section 10. References	44
Section A. Administration of NSAPs	46
Section A.1 GOSIP Version 2 NSAPs	47
Section A.1.1 Application for Administrative Authority	

Identifiers	48
Section A.1.2 Guidelines for NSAP Assignment	50
Section A.2 Data Country Code NSAPs	50
Section A.2.1 Application for Numeric Organization Name	51
Section A.3 Summary of Administrative Requirements ..	52

1. Introduction

The Internet is moving towards a multi-protocol environment that includes CLNP. To support CLNP in the Internet, an OSI lower layers infrastructure is required. This infrastructure comprises the connectionless network protocol (CLNP) [9] and supporting routing protocols. Also required as part of this infrastructure are guidelines for network service access point (NSAP) address assignment. This paper provides guidelines for allocating NSAP addresses in the Internet (the terms NSAP and NSAP address are used interchangeably throughout this paper in referring to NSAP addresses).

The guidelines presented in this document are quite similar to the guidelines that are proposed in the Internet for IP address allocation with CIDR (RFC 1519 [19]). The major difference between the two is the size of the addresses (4 octets for CIDR vs 20 octets for CLNP). The larger NSAP addresses allows considerably greater flexibility and scalability.

The remainder of this paper is organized into five major sections and an appendix. Section 2 defines the boundaries of the problem addressed in this paper and Section 3 provides background information on OSI routing and the implications for NSAP addresses.

Section 4 addresses the specific relationship between NSAP addresses and routing, especially with regard to hierarchical routing and data abstraction. This is followed in Section 5 with an application of these concepts to the Internet environment. Section 6 provides recommended guidelines for NSAP address allocation in the Internet. This includes recommendations for the U.S. and European parts of the Internet, as well as more general recommendations for any part of the Internet.

The Appendix contains a compendium of useful information concerning NSAP structure and allocation authorities. The GOSIP Version 2 NSAP structure is discussed in detail and the structure for U.S.-based DCC (Data Country Code) NSAPs is described. Contact information for the registration authorities for GOSIP and DCC-based NSAPs in the U.S., the General Services Administration (GSA) and the American National Standards Institute (ANSI), respectively, is provided.

This document obsoletes RFC 1237. The changes from RFC 1237 are minor, and primarily editorial in nature. The descriptions of OSI routing standards contained in Section 3 have been updated to reflect the current status of the relevant standards, and a description of the OSI Interdomain Routing Protocol (IDRP) has been added. Recommendations specific to the European part of the Internet have

been added in Section 6, along with recommendations for Routing Domain Identifiers and Routing Domain Confederation Identifiers needed for operation of IDRP.

2. Scope

Control over the collection of hosts and the transmission and switching facilities that compose the networking resources of the global Internet is not homogeneous, but is distributed among multiple administrative authorities. For the purposes of this paper, the term network service provider (or just provider) is defined to be an organization that is in the business of providing datagram switching services to customers. Organizations that are **only** customers (i.e., that do not provide datagram services to other organizations) are called network service subscribers (or simply subscribers).

In the current Internet, subscribers (e.g., campus and corporate site networks) attach to providers (e.g., regionals, commercial providers, and government backbones) in only one or a small number of carefully controlled access points. For discussion of OSI NSAP allocation in this paper, providers are treated as composing a mesh having no fixed hierarchy. Addressing solutions which require substantial changes or constraints on the current topology are not considered in this paper.

There are two aspects of interest when discussing OSI NSAP allocation within the Internet. The first is the set of administrative requirements for obtaining and allocating NSAP addresses; the second is the technical aspect of such assignments, having largely to do with routing, both within a routing domain (intra-domain routing) and between routing domains (inter-domain routing). This paper focuses on the technical issues.

The technical issues in NSAP allocation are mainly related to routing. This paper assumes that CLNP will be widely deployed in the Internet, and that the routing of CLNP traffic will normally be based on the OSI end-system to intermediate system routing protocol (ES-IS) [10], intra-domain IS-IS protocol [14], and inter-domain routing protocol (IDRP) [16]. It is expected that in the future the OSI routing architecture will be enhanced to include support for multicast, resource reservation, and other advanced services. The requirements for addressing for these future services is outside of the scope of this document.

The guidelines provided in this paper have been the basis for initial deployment of CLNP in the Internet, and have proven very valuable both as an aid to scaling of CLNP routing, and to address administration.

The guidelines in this paper are oriented primarily toward the large-scale division of NSAP address allocation in the Internet. Topics covered include:

- * Arrangement of parts of the NSAP for efficient operation of the IS-IS routing protocol;
- * Benefits of some topological information in NSAPs to reduce routing protocol overhead, and specifically the overhead on inter-domain routing (IDRP);
- * The anticipated need for additional levels of hierarchy in Internet addressing to support network growth and use of the Routing Domain Confederation mechanism of IDRP to provide support for additional levels of hierarchy;
- * The recommended mapping between Internet topological entities (i.e., service providers and service subscribers) and OSI addressing and routing components, such as areas, domains and confederations;
- * The recommended division of NSAP address assignment authority among service providers and service subscribers;
- * Background information on administrative procedures for registration of administrative authorities immediately below the national level (GOSIP administrative authorities and ANSI organization identifiers); and,
- * Choice of the high-order portion of the NSAP in subscriber routing domains that are connected to more than one service provider.

It is noted that there are other aspects of NSAP allocation, both technical and administrative, that are not covered in this paper. Topics not covered or mentioned only superficially include:

- * Identification of specific administrative domains in the Internet;
- * Policy or mechanisms for making registered information known to third parties (such as the entity to which a specific NSAP or a portion of the NSAP address space has been allocated);

- * How a routing domain (especially a site) should organize its internal topology of areas or allocate portions of its NSAP address space; the relationship between topology and addresses is discussed, but the method of deciding on a particular topology or internal addressing plan is not; and,
- * Procedures for assigning the System Identifier (ID) portion of the NSAP. A method for assignment of System IDs is presented in [18].

3. Background

Some background information is provided in this section that is helpful in understanding the issues involved in NSAP allocation. A brief discussion of OSI routing is provided, followed by a review of the intra-domain and inter-domain protocols in sufficient detail to understand the issues involved in NSAP allocation. Finally, the specific constraints that the routing protocols place on NSAPs are listed.

3.1. OSI Routing Standards

OSI partitions the routing problem into three parts:

- * routing exchanges between hosts (a.k.a., end systems or ESs) and routers (a.k.a., intermediate systems or ISs) (ES-IS);
- * routing exchanges between routers in the same routing domain (intra-domain IS-IS); and,
- * routing among routing domains (inter-domain IS-IS).

ES-IS (international standard ISO 9542) advanced to international standard (IS) status within ISO in 1987. Intra-domain IS-IS advanced to IS status within ISO in 1992. Inter-Domain Routing Protocol (IDRP) advanced to IS status within ISO in October 1993. CLNP, ES-IS, and IS-IS are all widely available in vendor products, and have been deployed in the Internet for several years. IDRP is currently being implemented in vendor products.

This paper examines the technical implications of NSAP assignment under the assumption that ES-IS, intra-domain IS-IS, and IDRP routing are deployed to support CLNP.

3.2. Overview of ISIS (ISO/IEC 10589)

The IS-IS intra-domain routing protocol, ISO/IEC 10589, provides routing for OSI environments. In particular, IS-IS is designed to work in conjunction with CLNP, ES-IS, and IDRP. This section briefly describes the manner in which IS-IS operates.

In IS-IS, the internetwork is partitioned into routing domains. A routing domain is a collection of ESs and ISs that operate common routing protocols and are under the control of a single administration (throughout this paper, "domain" and "routing domain" are used interchangeably). Typically, a routing domain may consist of a corporate network, a university campus network, a regional network, a backbone, or a similar contiguous network under control of a single administrative organization. The boundaries of routing domains are defined by network management by setting some links to be exterior, or inter-domain, links. If a link is marked as exterior, no intra-domain IS-IS routing messages are sent on that link.

IS-IS routing makes use of two-level hierarchical routing. A routing domain is subdivided into areas (also known as level 1 subdomains). Level 1 routers know the topology in their area, including all routers and hosts. However, level 1 routers do not know the identity of routers or destinations outside of their area. Level 1 routers forward all traffic for destinations outside of their area to a level 2 router within their area.

Similarly, level 2 routers know the level 2 topology and know which addresses are reachable via each level 2 router. The set of all level 2 routers in a routing domain are known as the level 2 subdomain, which can be thought of as a backbone for interconnecting the areas. Level 2 routers do not need to know the topology within any level 1 area, except to the extent that a level 2 router may also be a level 1 router within a single area. Only level 2 routers can exchange data packets or routing information directly with routers located outside of their routing domain.

NSAP addresses provide a flexible, variable length addressing format, which allows for multi-level hierarchical address assignment. These addresses provide the flexibility needed to solve two critical problems simultaneously: (i) How to administer a worldwide address space; and (ii) How to assign addresses in a manner which makes routing scale well in a worldwide Internet.

As illustrated in Figure 1, ISO addresses are subdivided into the Initial Domain Part (IDP) and the Domain Specific Part (DSP). The IDP is the part which is standardized by ISO, and specifies the format and authority responsible for assigning the rest of the

address. The DSP is assigned by whatever addressing authority is specified by the IDP (see Appendix A for more discussion on the top level NSAP addressing authorities). It is expected that the authority specified by the IDP may further sub-divide the DSP, and may assign sub-authorities responsible for parts of the DSP.

For routing purposes, ISO addresses are subdivided by IS-IS into the area address, the system identifier (ID), and the NSAP selector (SEL). The area address identifies both the routing domain and the area within the routing domain. Generally, the area address corresponds to the IDP plus a high-order part of the DSP (HO-DSP).

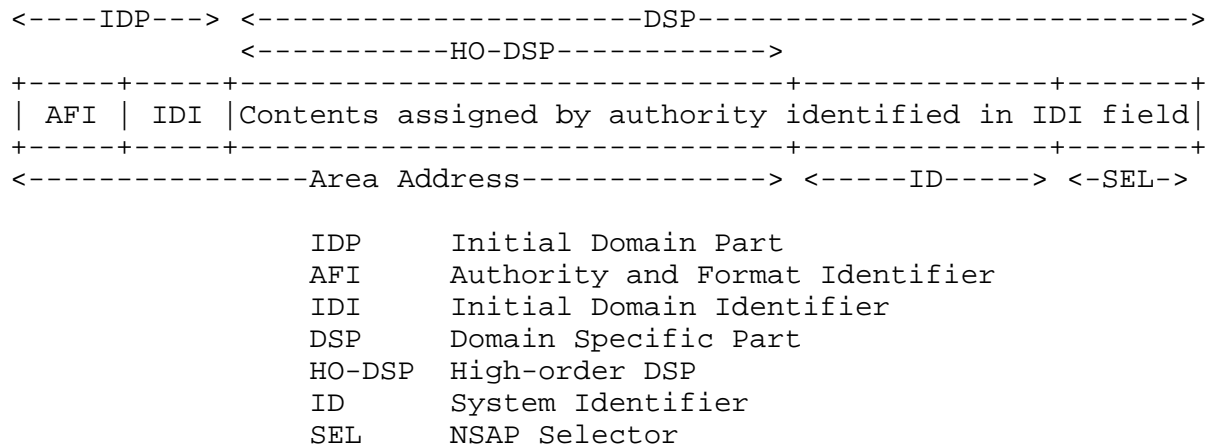


Figure 1: OSI Hierarchical Address Structure.

The ID field may be from one to eight octets in length, but must have a single known length in any particular routing domain. Each router is configured to know what length is used in its domain. The SEL field is always one octet in length. Each router is therefore able to identify the ID and SEL fields as a known number of trailing octets of the NSAP address. The area address can be identified as the remainder of the address (after truncation of the ID and SEL fields). It is therefore not necessary for the area address to have any particular length -- the length of the area address could vary between different area addresses in a given routing domain.

Usually, all nodes in an area have the same area address. However, sometimes an area might have multiple addresses. Motivations for allowing this are several:

- * It might be desirable to change the address of an area. The most graceful way of changing an area address from A to B is to first allow it to have both addresses A and B, and then after all nodes in the area have been modified to recognize both addresses, one by one the nodes can be modified to forget address A.
- * It might be desirable to merge areas A and B into one area. The method for accomplishing this is to, one by one, add knowledge of address B into the A partition, and similarly add knowledge of address A into the B partition.
- * It might be desirable to partition an area C into two areas, A and B (where A might equal C, in which case this example becomes one of removing a portion of an area). This would be accomplished by first introducing knowledge of address A into the appropriate nodes (those destined to become area A), and knowledge of address B into the appropriate nodes, and then one by one removing knowledge of address C.

Since the addressing explicitly identifies the area, it is very easy for level 1 routers to identify packets going to destinations outside of their area, which need to be forwarded to level 2 routers. Thus, in IS-IS routers perform as follows:

- * Level 1 intermediate systems route within an area based on the ID portion of the ISO address. Level 1 routers recognize, based on the destination address in a packet, whether the destination is within the area. If so, they route towards the destination. If not, they route to the nearest level 2 router.
- * Level 2 intermediate systems route based on address prefixes, preferring the longest matching prefix, and preferring internal routes over external routes. They route towards areas, without regard to the internal structure of an area; or towards level 2 routers on the routing domain boundary that have advertised external address prefixes into the level 2 subdomain. A level 2 router may also be operating as a level 1 router in one area.

A level 1 router will have the area portion of its address manually configured. It will refuse to become a neighbor with a router whose area addresses do not overlap its own area addresses. However, if a level 1 router has area addresses A, B, and C, and a neighbor has area addresses B and D, then the level 1 IS will accept the other IS as a level 1 neighbor.

A level 2 router will accept another level 2 router as a neighbor, regardless of area address. However, if the area addresses do not overlap, the link would be considered by both routers to be level 2

only, and only level 2 routing packets would flow on the link. External links (i.e., to other routing domains) must be between level 2 routers in different routing domains.

IS-IS provides an optional partition repair function. If a level 1 area becomes partitioned, this function, if implemented, allows the partition to be repaired via use of level 2 routes.

IS-IS requires that the set of level 2 routers be connected. Should the level 2 backbone become partitioned, there is no provision for use of level 1 links to repair a level 2 partition.

Occasionally a single level 2 router may lose connectivity to the level 2 backbone. In this case the level 2 router will indicate in its level 1 routing packets that it is not "attached", thereby allowing level 1 routers in the area to route traffic for outside of the area to a different level 2 router. Level 1 routers therefore route traffic to destinations outside of their area only to level 2 routers which indicate in their level 1 routing packets that they are "attached".

A host may autoconfigure the area portion of its address by extracting the area portion of a neighboring router's address. If this is the case, then a host will always accept a router as a neighbor. Since the standard does not specify that the host **must** autoconfigure its area address, a host may be pre-configured with an area address.

Special treatment is necessary for broadcast subnetworks, such as LANs. This solves two sets of issues: (i) In the absence of special treatment, each router on the subnetwork would announce a link to every other router on the subnetwork, resulting in $O(n^2)$ links reported; (ii) Again, in the absence of special treatment, each router on the LAN would report the same identical list of end systems on the LAN, resulting in substantial duplication.

These problems are avoided by use of a "pseudonode", which represents the LAN. Each router on the LAN reports that it has a link to the pseudonode (rather than reporting a link to every other router on the LAN). One of the routers on the LAN is elected "designated router". The designated router then sends out a Link State Packet (LSP) on behalf of the pseudonode, reporting links to all of the routers on the LAN. This reduces the potential n^2 links to n links. In addition, only the pseudonode LSP includes the list of end systems on the LAN, thereby eliminating the potential duplication.

The IS-IS provides for optional Quality of Service (QOS) routing, based on throughput (the default metric), delay, expense, or residual error probability.

IS-IS has a provision for authentication information to be carried in all IS-IS PDUs. Currently the only form of authentication which is defined is a simple password. A password may be associated with each link, each area, and with the level 2 subdomain. A router not in possession of the appropriate password(s) is prohibited from participating in the corresponding function (i.e., may not initialize a link, be a member of the area, or a member of the level 2 subdomain, respectively).

Procedures are provided to allow graceful migration of passwords without disrupting operation of the routing protocol. The authentication functions are extensible so that a stronger, cryptographically-based security scheme may be added in an upwardly compatible fashion at a future date.

3.3. Overview of IDRP (ISO/IEC 10747)

The Inter-Domain Routing Protocol (IDRP, ISO/IEC 10747), developed in ISO, provides routing for OSI environments. In particular, IDRP is designed to work in conjunction with CLNP, ES-IS, and IS-IS. This section briefly describes the manner in which IDRP operates.

Consistent with the OSI Routing Framework [13], in IDRP the internetwork is partitioned into routing domains. IDRP places no restrictions on the inter-domain topology. A router that participates in IDRP is called a Boundary Intermediate System (BIS). Routing domains that participate in IDRP are not allowed to overlap - a BIS may belong to only one domain.

A pair of BISs are called external neighbors if these BISs belong to different domains but share a common subnetwork (i.e., a BIS can reach its external neighbor in a single network layer hop). Two domains are said to be adjacent if they have BISs that are external neighbors of each other. A pair of BISs are called internal neighbors if these BISs belong to the same domain. In contrast with external neighbors, internal neighbors don't have to share a common subnetwork -- IDRP assumes that a BIS should be able to exchange Network Protocol Data Units (NPDUs) with any of its internal neighbors by relying solely on intra-domain routing procedures.

IDRP governs the exchange of routing information between a pair of neighbors, either external or internal. IDRP is self-contained with respect to the exchange of information between external neighbors. Exchange of information between internal neighbors relies on

additional support provided by intra-domain routing (unless internal neighbors share a common subnetwork).

To facilitate routing information aggregation/abstraction, IDRP allows grouping of a set of connected domains into a Routing Domain Confederation (RDC). A given domain may belong to more than one RDC. There are no restrictions on how many RDCs a given domain may simultaneously belong to, and no preconditions on how RDCs should be formed -- RDCs may be either nested, or disjoint, or may overlap. One RDC is nested within another RDC if all members (RDs) of the former are also members of the latter, but not vice versa. Two RDCs overlap if they have members in common and also each has members that are not in the other. Two RDCs are disjoint if they have no members in common.

Each domain participating in IDRP is assigned a unique Routing Domain Identifier (RDI). Syntactically an RDI is represented as an OSI network layer address. Each RDC is assigned a unique Routing Domain Confederation Identifier (RDCI). RDCIs are assigned out of the address space allocated for RDIs -- RDCIs and RDIs are syntactically indistinguishable. Procedures for assigning and managing RDIs and RDCIs are outside the scope of the protocol. However, since RDIs are syntactically nothing more than network layer addresses, and RDCIs are syntactically nothing more than RDIs, it is expected that RDI and RDCI assignment and management would be part of the network layer assignment and management procedures. Recommendations for RDI and RDCI assignment are provided in Section 6.5.

IDRP requires a BIS to be preconfigured with the RDI of the domain to which the BIS belongs. If a BIS belongs to a domain that is a member of one or more RDCs, then the BIS has to be preconfigured with RDCIs of all the RDCs the domain is in, and the information about relations between the RDCs - nested or overlapped.

IDRP doesn't assume or require any particular internal structure for the addresses. The protocol provides correct routing as long as the following guidelines are met:

- * End systems and intermediate systems may use any NSAP address or Network Entity Title (NET -- i.e., an NSAP address without the selector) that has been assigned under ISO 8348 [11] guidelines;
- * An NSAP prefix carried in the Network Layer Reachability Information (NLRI) field for a route originated by a BIS in a given routing domain should be associated with only that routing domain; that is, no system identified by the prefix should reside in a different routing domain; ambiguous routing may result if several routing domains originate routes whose

NLRI field contain identical NSAP address prefixes, since this would imply that the same system(s) is simultaneously located in several routing domains;

- * Several different NSAP prefixes may be associated with a single routing domain which contains a mix of systems which use NSAP addresses assigned by several different addressing authorities.

IDRP assumes that the above guidelines have been satisfied, but it contains no means to verify that this is so. Therefore, such verification is assumed to be the responsibility of the administrators of routing domains.

IDRP provides mandatory support for data integrity and optional support for data origin authentication for all of its messages. Each message carries a 16-octet digital signature that is computed by applying the MD-4 algorithm (RFC 1320) to the context of the message itself. This signature provides support for data integrity. To support data origin authentication a BIS, when computing a digital signature of a message, may prepend and append additional information to the message. This information is not passed as part of the message but is known to the receiver.

3.3.1. Scaling Mechanisms in IDRP

The ability to group domains in RDCs provides a simple, yet powerful mechanism for routing information aggregation and abstraction. It allows reduction of topological information by replacing a sequence of RDIs carried by the RD_PATH attribute with a single RDCI. It also allows reduction of the amount of information related to transit policies, since the policies can be expressed in terms of aggregates (RDCs), rather than individual components (RDs). It also allows simplification of route selection policies, since these policies can be expressed in terms of aggregates (RDCs) rather than individual components (RDs).

Aggregation and abstraction of Network Layer Reachability Information (NLRI) is supported by the "route aggregation" mechanism of IDRP. This mechanism is complementary to the Routing Domain Confederations mechanism. Both mechanisms are intended to provide scalable routing via information reduction/abstraction. However, the two mechanisms are used for different purposes: route aggregation for aggregation and abstraction of routes (i.e., Network Layer Reachability Information), Routing Domain Confederations for aggregation and abstraction of topology and/or policy information. To provide maximum benefits, both mechanisms can be used together. This implies that address assignment that will facilitate route aggregation does not conflict with the ability to form RDCs, and vice versa; formation

of RDCs should be done in a manner consistent with the address assignment needed for route aggregation.

3.4. Requirements of IS-IS and IDRP on NSAPs

The preferred NSAP format for IS-IS is shown in Figure 1. A number of points should be noted from IS-IS:

- * The IDP is as specified in ISO 8348, the OSI network layer service specification [11];
- * The high-order portion of the DSP (HO-DSP) is that portion of the DSP whose assignment, structure, and meaning are not constrained by IS-IS;
- * The area address (i.e., the concatenation of the IDP and the HO-DSP) must be globally unique. If the area address of an NSAP matches one of the area addresses of a router, it is in the router's area and is routed to by level 1 routing;
- * Level 2 routing acts on address prefixes, using the longest address prefix that matches the destination address;
- * Level 1 routing acts on the ID field. 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. The ID field is assumed to be flat. The method presented in RFC 1526 [18] may optionally be used to assure globally unique IDs;
- * The one-octet NSAP Selector, SEL, determines the entity to receive the CLNP packet within the system identified by the rest of the NSAP (i.e., a transport entity) and is always the last octet of the NSAP; and,
- * A system shall be able to generate and forward data packets containing addresses in any of the formats specified by ISO 8348. However, within a routing domain that conforms to IS-IS, 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 IS-IS routing. End systems with addresses which do not conform may require additional manual configuration and be subject to inferior routing performance.

For purposes of efficient operation of the IS-IS routing protocol, several observations may be made. First, although the IS-IS protocol specifies an algorithm for routing within a single routing domain, the routing algorithm must efficiently route both: (i) Packets whose final destination is in the domain (these must, of course, be routed

to the correct destination end system in the domain); and (ii) Packets whose final destination is outside of the domain (these must be routed to an appropriate "border" router, from which they will exit the domain).

For those destinations which are in the domain, level 2 routing treats the entire area address (i.e., all of the NSAP address except the ID and SEL fields) as if it were a flat field. Thus, the efficiency of level 2 routing to destinations within the domain is affected only by the number of areas in the domain, and the number of area addresses assigned to each area.

For those destinations which are outside of the domain, level 2 routing routes according to address prefixes. In this case, there is considerable potential advantage (in terms of reducing the amount of routing information that is required) if the number of address prefixes required to describe any particular set of external destinations can be minimized. Efficient routing with IDRP similarly also requires minimization of the number of address prefixes needed to describe specific destinations. In other words, addresses need to be assigned with topological significance. This requirement is described in more detail in the following sections.

4. NSAPs and Routing

4.1. Routing Data Abstraction

When determining an administrative policy for NSAP assignment, it is important to understand the technical consequences. The objective behind the use of hierarchical routing is to achieve some level of routing data abstraction, or summarization, to reduce the processing time, memory requirements, and transmission bandwidth consumed in support of routing. This implies that address assignment must serve the needs of routing, in order for routing to scale to very large networks.

While the notion of routing data abstraction may be applied to various types of routing information, this and the following sections primarily emphasize one particular type, namely reachability information. Reachability information describes the set of reachable destinations.

Abstraction of reachability information dictates that NSAPs be assigned according to topological routing structures. However, administrative assignment falls along organizational or political boundaries. These may not be congruent to topological boundaries, and therefore the requirements of the two may collide. A balance between these two needs is necessary.

Routing data abstraction occurs at the boundary between hierarchically arranged topological routing structures. An element lower in the hierarchy reports summary routing information to its parent(s). Within the current OSI routing framework [13] and routing protocols, the lowest boundary at which this can occur is the boundary between an area and the level 2 subdomain within a IS-IS routing domain. Data abstraction is designed into IS-IS at this boundary, since level 1 ISs are constrained to reporting only area addresses.

Level 2 routing is based upon address prefixes. Level 2 routers (ISs) distribute, throughout the level 2 subdomain, the area addresses of the level 1 areas to which they are attached (and any manually configured reachable address prefixes). Level 2 routers compute next-hop forwarding information to all advertised address prefixes. Level 2 routing is determined by the longest advertised address prefix that matches the destination address.

At routing domain boundaries, address prefix information is exchanged with other routing domains via IDRP. If area addresses within a routing domain are all drawn from distinct NSAP assignment authorities (allowing no abstraction), then the boundary prefix information consists of an enumerated list of all area addresses.

Alternatively, should the routing domain "own" an address prefix and assign area addresses based upon it, boundary routing information can be summarized into the single prefix. This can allow substantial data reduction and, therefore, will allow much better scaling (as compared to the uncoordinated area addresses discussed in the previous paragraph).

If routing domains are interconnected in a more-or-less random (non-hierarchical) scheme, it is quite likely that no further abstraction of routing data can occur. Since routing domains would have no defined hierarchical relationship, administrators would not be able to assign area addresses out of some common prefix for the purpose of data abstraction. The result would be flat inter-domain routing; all routing domains would need explicit knowledge of all other routing domains that they route to. This can work well in small- and medium-sized internets, up to a size somewhat larger than the current IP Internet. However, this does not scale to very large internets. For example, we expect growth in the future to an international Internet which has tens or hundreds of thousands of routing domains in the U.S. alone. Even larger numbers of routing domains are possible when each home, or each small company, becomes its own routing domain. This requires a greater degree of data abstraction beyond that which can be achieved at the "routing domain" level.

In the Internet, however, it should be possible to exploit the existing hierarchical routing structure interconnections, as discussed in Section 5. Thus, there is the opportunity for a group of subscribers each to be assigned an address prefix from a shorter prefix assigned to their provider. Each subscriber now "owns" its (somewhat longer) prefix, from which it assigns its area addresses.

The most straightforward case of this occurs when there is a set of subscribers whose routing domains are all attached only to a single service provider, and which use that provider for all external (inter-domain) traffic. A short address prefix may be assigned to the provider, which then assigns slightly longer prefixes (based on the provider's prefix) to each of the subscribers. This allows the provider, when informing other providers of the addresses that it can reach, to abbreviate the reachability information for a large number of routing domains as a single prefix. This approach therefore can allow a great deal of hierarchical abbreviation of routing information, and thereby can greatly improve the scalability of inter-domain routing.

Clearly, this approach is recursive and can be carried through several iterations. Routing domains at any "level" in the hierarchy may use their prefix as the basis for subsequent suballocations, assuming that the NSAP addresses remain within the overall length and structure constraints. The flexibility of NSAP addresses facilitates this form of hierarchical address assignment and routing. As one example of how NSAPs may be used, the GOSIP Version 2 NSAP structure is discussed later in this section.

At this point, we observe that the number of nodes at each lower level of a hierarchy tends to grow exponentially. Thus the greatest gains in data abstraction occur at the leaves and the gains drop significantly at each higher level. Therefore, the law of diminishing returns suggests that at some point data abstraction ceases to produce significant benefits. Determination of the point at which data abstraction ceases to be of benefit requires a careful consideration of the number of routing domains that are expected to occur at each level of the hierarchy (over a given period of time), compared to the number of routing domains and address prefixes that can conveniently and efficiently be handled via dynamic inter-domain routing protocols. As the Internet grows, further levels of hierarchy may become necessary. Again, this requires considerable flexibility in the addressing scheme, such as is provided by NSAP addresses.

4.2. NSAP Administration and Efficiency

There is a balance that must be sought between the requirements on NSAPs for efficient routing and the need for decentralized NSAP administration. The NSAP structure from Version 2 of GOSIP (Figure 2) offers one example of how these two needs might be met. The AFI, IDI, DSP Format Identifier (DFI), and Administrative Authority (AA) fields provide for administrative decentralization. The AFI/IDI pair of values 47.0005 identify the U.S. Government as the authority responsible for defining the DSP structure and allocating values within it (see the Appendix for more information on NSAP structure).

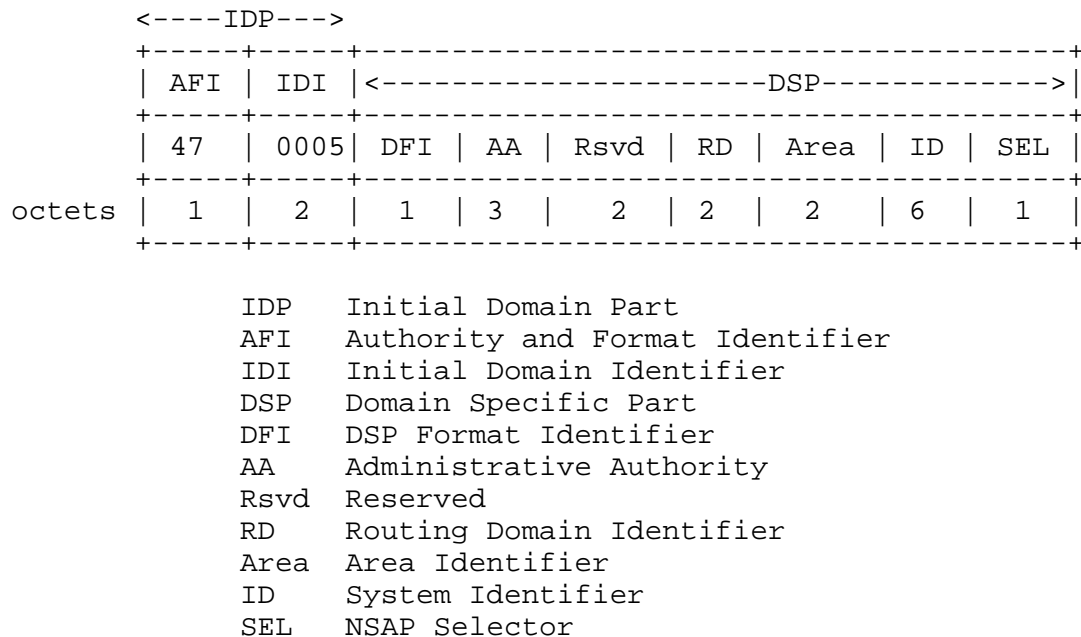


Figure 2: GOSIP Version 2 NSAP structure.

[Note: We are using U.S. GOSIP version 2 addresses only as an example. It is not necessary that NSAPs be allocated from the GOSIP Version 2 authority under 47.0005. The ANSI format under the Data Country Code for the U.S. (DCC=840) and formats assigned to other countries and ISO members or liaison organizations are also being used, and work equally well. For parts of the Internet outside of the U.S. there may in some cases be strong reasons to prefer a country- or area-specific format rather than the U.S. GOSIP format. However, GOSIP addresses are used in most cases in the examples in this paper because:

- * The DSP format has been defined and allows hierarchical allocation; and,

- * An operational registration authority for suballocation of AA values under the GOSIP address space has already been established at GSA.]

GOSIP Version 2 defines the DSP structure as shown (under DFI=80h) and provides for the allocation of AA values to administrations. Thus, the fields from the AFI to the AA, inclusive, represent a unique address prefix assigned to an administration.

American National Standard X3.216-1992 [1] specifies the structure of the DSP for NSAP addresses that use an Authority and Format Identifier (AFI) value of (decimal) 39, which identifies the "ISO-DCC" (data country code) format, in which the value of the Initial Domain Identifier (IDI) is (decimal) 840, which identifies the U.S. National Body (ANSI). This DSP structure is identical to the structure that is specified by GOSIP Version 2. The AA field is called "org" for organization identifier in the ANSI standard, and the ID field is called "system". The ANSI format, therefore, differs from the GOSIP format illustrated above only in that the AFI and IDI specify the "ISO-DCC" format rather than the "ISO 6523-ICD" format used by GOSIP, and the "AA" field is administered by an ANSI registration authority rather than by the GSA. Organization identifiers may be obtained from ANSI. The technical considerations applicable to NSAP administration are independent of whether a GOSIP Version 2 or an ANSI value is used for the NSAP assignment.

Similarly, although other countries make use of different NSAP formats, the principles of NSAP assignment and use are the same. The NSAP formats recommended by RARE WG4 for use in Europe are discussed in Section 6.2.

In the low-order part of the GOSIP Version 2 NSAP format, two fields are defined in addition to those required by IS-IS. These fields, RD and Area, are defined to allow allocation of NSAPs along topological boundaries in support of increased data abstraction. Administrations assign RD identifiers underneath their unique address prefix (the reserved field is left to accommodate future growth and to provide additional flexibility for inter-domain routing). Routing domains allocate Area identifiers from their unique prefix. The result is:

- * AFI+IDI+DFI+AA = administration prefix,
- * administration prefix(+Rsvd)+RD = routing domain prefix, and,
- * routing domain prefix+Area = area address.

This provides for summarization of all area addresses within a routing domain into one prefix. If the AA identifier is accorded topological significance (in addition to administrative significance), an additional level of data abstraction can be obtained, as is discussed in the next section.

5. NSAP Administration and Routing in the Internet

Basic Internet routing components are service providers and service subscribers. A natural mapping from these components to OSI routing components is that each provider and subscriber operates as a routing domain.

Alternatively, a subscriber may choose to operate as a part of a provider domain; that is, as an area within the provider's routing domain. However, in such a case the discussion in Section 5.1 applies.

We assume that most subscribers will prefer to operate a routing domain separate from their provider's. Such subscribers can exchange routing information with their provider via interior routing protocol route leaking or via IDRP; for the purposes of this discussion, the choice is not significant. The subscriber is still allocated a prefix from the provider's address space, and the provider advertises its own prefix into inter-domain routing.

Given such a mapping, where should address administration and allocation be performed to satisfy both administrative decentralization and data abstraction? Three possibilities are considered:

1. at the area,
2. at the subscriber routing domain, and,
3. at the provider routing domain.

Subscriber routing domains correspond to end-user sites, where the primary purpose is to provide intra-domain routing services. Provider routing domains are deployed to carry transit (i.e., inter-domain) traffic.

The greatest burden in transmitting and operating on routing information is at the top of the routing hierarchy, where routing information tends to accumulate. In the Internet, for example, each provider must manage the set of network numbers for all networks reachable through the provider.

For traffic destined for other networks, the provider will route based on inter-domain routing information obtained from other providers or, in some cases, to a default provider.

In general, higher levels of the routing hierarchy will benefit the most from the abstraction of routing information at a lower level of the routing hierarchy. There is relatively little direct benefit to the administration that performs the abstraction, since it must maintain routing information individually on each attached topological routing structure.

For example, suppose that a given subscriber is trying to decide whether to obtain an NSAP address prefix based on an AA value from GSA (implying that the first four octets of the address would be those assigned out of the GOSIP space), or based on an RD value from its provider (implying that the first seven octets of the address are those obtained by that provider). If considering only their own self-interest, the subscriber and its local provider have little reason to choose one approach or the other. The subscriber must use one prefix or another; the source of the prefix has little effect on routing efficiency within the subscriber's routing domain. The provider must maintain information about each attached subscriber in order to route, regardless of any commonality in the prefixes of its subscribers.

However, there is a difference when the local provider distributes routing information to other providers. In the first case, the provider cannot aggregate the subscriber's address into its own prefix; the address must be explicitly listed in routing exchanges, resulting in an additional burden to other providers which must exchange and maintain this information.

In the second case, each other provider sees a single address prefix for the local provider which encompasses the new subscriber. This avoids the exchange of additional routing information to identify the new subscriber's address prefix. Thus, the advantages primarily benefit other providers which maintain routing information about this provider (and its subscribers).

Clearly, a symmetric application of these principles is in the interest of all providers, enabling them to more efficiently support CLNP routing to their customers. The guidelines discussed below describe reasonable ways of managing the OSI address space that benefit the entire community.

5.1. Administration at the Area

If areas take their area addresses from a myriad of unrelated NSAP allocation authorities, there will be effectively no data abstraction beyond what is built into IS-IS. For example, assume that within a routing domain three areas take their area addresses, respectively, out of:

- * the GOSIP Version 2 authority assigned to the Department of Commerce, with an AA of nnn:

AFI=47, IDI=0005, DFI=80h, AA=nnn, ... ;

- * the GOSIP Version 2 authority assigned to the Department of the Interior, with an AA of mmm:

AFI=47, IDI=0005, DFI=80h, AA=mmm, ... ; and,

- * the ANSI authority under the U.S. Data Country Code (DCC)

(Section A.2) for organization XYZ with ORG identifier = xxx:

AFI=39, IDI=840, DFI=dd, ORG=xxx,

As described in Section 3.3, from the point of view of any particular routing domain, there is no harm in having the different areas in the routing domain use addresses obtained from a wide variety of administrations. For routing within the domain, the area addresses are treated as a flat field.

However, this does have a negative effect on inter-domain routing, particularly on those other domains which need to maintain routes to this domain. There is no common prefix that can be used to represent these NSAPs and therefore no summarization can take place at the routing domain boundary. When addresses are advertised by this routing domain to other routing domains, an enumerated list must be used consisting of the three area addresses.

This situation is roughly analogous to the dissemination of routing information in the TCP/IP Internet prior to the introduction of CIDR. Areas correspond roughly to networks and area addresses to network numbers. The result of allowing areas within a routing domain to take their NSAPs from unrelated authorities is flat routing at the area address level. The number of address prefixes that subscriber routing domains would advertise is on the order of the number of attached areas; the number of prefixes a provider routing domain would advertise is approximately the number of areas attached to all

its subscriber routing domains. For "default-less" providers (i.e., those that don't use default routes) the size of the routing tables would be on the order of the number of area addresses globally. As the CLNP internet grows this would quickly become intractable. A greater degree of hierarchical information reduction is necessary to allow greater growth.

5.2. Administration at the Subscriber Routing Domain

As mentioned previously, the greatest degree of data abstraction comes at the lowest levels of the hierarchy. Providing each subscriber routing domain (that is, site) with a unique prefix results in the biggest single increase in abstraction, with each subscriber domain assigning area addresses from its prefix. From outside the subscriber routing domain, the set of all addresses reachable in the domain can then be represented by a single prefix.

As an example, assume a government agency has been assigned the AA value of zzz under ICD=0005. The agency then assigns a routing domain identifier to a routing domain under its administrative authority identifier, rrr. The resulting prefix for the routing domain is:

AFI=47, IDI=0005, DFI=80h, AA=zzz, (Rsvd=0), RD=rrr.

All areas within this routing domain would have area addresses comprising this prefix followed by an Area identifier. The prefix represents the summary of reachable addresses within the routing domain.

There is a close relationship between areas and routing domains implicit in the fact that they operate a common routing protocol and are under the control of a single administration. The routing domain administration subdivides the domain into areas and structures a level 2 subdomain (i.e., a level 2 backbone) which provides connectivity among the areas. The routing domain represents the only path between an area and the rest of the internetwork. It is reasonable that this relationship also extend to include a common NSAP addressing authority. Thus, the areas within the subscriber RD should take their NSAPs from the prefix assigned to the subscriber RD.

5.3. Administration at the Provider Routing Domain

Two kinds of provider routing domains are considered, direct providers and indirect providers. Most of the subscribers of a direct provider are domains that act solely as service subscribers (i.e., they carry no transit traffic). Most of the "subscribers" of

an indirect provider are, themselves, service providers. In present terminology a backbone is an indirect provider, while a regional is a direct provider. Each case is discussed separately below.

5.3.1. Direct Service Providers

It is interesting to consider whether direct service providers' routing domains should be the common authority for assigning NSAPs from a unique prefix to the subscriber routing domains that they serve. In the long term the number of routing domains in the Internet will grow to the point that it will be infeasible to route on the basis of a flat field of routing domains. It will therefore be essential to provide a greater degree of information abstraction.

Direct providers may assign prefixes to subscriber domains, based on a single (shorter length) address prefix assigned to the provider. For example, given the GOSIP Version 2 address structure, an AA value may be assigned to each direct provider, and routing domain values may be assigned by the provider to each attached subscriber routing domain. A similar hierarchical address assignment based on a prefix assigned to each provider may be used for other NSAP formats. This results in direct providers advertising to other providers (both direct and indirect) a small fraction of the number of address prefixes that would be necessary if they enumerated the individual prefixes of the subscriber routing domains. This represents a significant savings given the expected scale of global internetworking.

Are subscriber routing domains willing to accept prefixes derived from the direct providers? In the supplier/consumer model, the direct provider is offering connectivity as the service, priced according to its costs of operation. This includes the "price" of obtaining service from one or more indirect providers and exchanging routing information with other direct providers. In general, providers will want to handle as few address prefixes as possible to keep costs low. In the Internet environment, subscriber routing domains must be sensitive to the resource constraints of the providers (both direct and indirect). The efficiencies gained in routing clearly warrant the adoption of NSAP administration by the direct providers.

The mechanics of this scenario are straightforward. Each direct provider is assigned a unique prefix, from which it allocates slightly longer routing domain prefixes for its attached subscriber routing domains. For GOSIP NSAPs, this means that a direct provider would be assigned an AA identifier. Attached subscriber routing domains would be assigned RD identifiers under the direct provider's unique prefix. For example, assume that NIST is a subscriber routing domain whose sole inter-domain link is via SURANet. If SURANet is

assigned an AA identifier kkk, NIST could be assigned an RD of jjj, resulting in a unique prefix for SURANet of:

AFI=47, IDI=0005, DFI=80h, AA=kkk

and a unique prefix for NIST of

AFI=47, IDI=0005, DFI=80h, AA=kkk, (Rsvd=0), RD=jjj.

A similar scheme can be established using NSAPs allocated under DCC=840. In this case, a direct provider applies for an ORG identifier from ANSI, which serves the same purpose as the AA identifier in GOSIP.

5.3.2. Indirect Providers

There does not appear to be a strong case for direct service providers to take their address spaces from the NSAP space of an indirect provider (e.g. backbone in today's terms). The benefit in routing data abstraction is relatively small. The number of direct providers today is in the tens and an order of magnitude increase would not cause an undue burden on the indirect providers. Also, it may be expected that as time goes by there will be increased direct inter-connection of the direct providers, subscriber routing domains directly attached to the "indirect" providers, and international links directly attached to the providers. Under these circumstances, the distinction between direct and indirect providers would become blurred.

An additional factor that discourages allocation of NSAPs from an indirect provider's prefix is that the indirect providers and their attached direct providers are perceived as being independent. Direct providers may take their indirect provider service from one or more providers, or may switch indirect providers should a more cost-effective service be available elsewhere (essentially, indirect providers can be thought of the same way as long-distance telephone carriers). Having NSAPs derived from the indirect providers is inconsistent with the nature of the relationship.

5.4. Multi-homed Routing Domains

The discussions in Section 5.3 suggest methods for allocating NSAP addresses based on service provider connectivity. This allows a great deal of information reduction to be achieved for those routing domains which are attached to a single provider. In particular, such routing domains may select their NSAP addresses from a space allocated to them by their direct service provider. This allows the provider, when announcing the addresses that it can reach to other

providers, to use a single address prefix to describe a large number of NSAP addresses corresponding to multiple routing domains.

However, there are additional considerations for routing domains which are attached to multiple providers. Such "multi-homed" routing domains may, for example, consist of single-site campuses and companies which are attached to multiple providers, large organizations which are attached to different providers at different locations in the same country, or multi-national organizations which are attached to providers in a variety of countries worldwide. There are a number of possible ways to deal with these multi-homed routing domains.

One possible solution is to assign addresses to each multi-homed organization independently from the providers to which it is attached. This allows each multi-homed organization to base its NSAP assignments on a single prefix, and to thereby summarize the set of all NSAPs reachable within that organization via a single prefix. The disadvantage of this approach is that since the NSAP address for that organization has no relationship to the addresses of any particular provider, the providers to which this organization is attached will need to advertise the prefix for this organization to other providers. Other providers (potentially worldwide) will need to maintain an explicit entry for that organization in their routing tables. If other providers do not maintain a separate route for this organization, then packets destined to this organization will be lost.

For example, suppose that a very large U.S.-wide company "Mega Big International Incorporated" (MBII) has a fully interconnected internal network and is assigned a single AA value under the U.S. GOSIP Version 2 address space. It is likely that outside of the U.S., a single entry may be maintained in routing tables for all U.S. GOSIP addresses. However, within the U.S., every "default-less" provider will need to maintain a separate address entry for MBII. If MBII is in fact an international corporation, then it may be necessary for every "default-less" provider worldwide to maintain a separate entry for MBII (including providers to which MBII is not attached). Clearly this may be acceptable if there are a small number of such multihomed routing domains, but would place an unacceptable load on routers within providers if all organizations were to choose such address assignments. This solution may not scale to internets where there are many hundreds of thousands of multi-homed organizations.

A second possible approach would be for multi-homed organizations to be assigned a separate NSAP space for each connection to a provider, and to assign a single address prefix to each area within its routing

domain(s) based on the closest interconnection point. For example, if MBII had connections to two providers in the U.S. (one east coast, and one west coast), as well as three connections to national providers in Europe, and one in the far east, then MBII may make use of six different address prefixes. Each area within MBII would be assigned a single address prefix based on the nearest connection.

For purposes of external routing of traffic from outside MBII to a destination inside of MBII, this approach works similarly to treating MBII as six separate organizations. For purposes of internal routing, or for routing traffic from inside of MBII to a destination outside of MBII, this approach works the same as the first solution.

If we assume that incoming traffic (coming from outside of MBII, with a destination within MBII) is always to enter via the nearest point to the destination, then each provider which has a connection to MBII needs to announce to other providers the ability to reach only those parts of MBII whose address is taken from its own address space. This implies that no additional routing information needs to be exchanged between providers, resulting in a smaller load on the inter-domain routing tables maintained by providers when compared to the first solution. This solution therefore scales better to extremely large internets containing very large numbers of multi-homed organizations.

One problem with the second solution is that backup routes to multi-homed organizations are not automatically maintained. With the first solution, each provider, in announcing the ability to reach MBII, specifies that it is able to reach all of the NSAPs within MBII. With the second solution, each provider announces that it can reach all of the NSAPs based on its own address prefix, which only includes some of the NSAPs within MBII. If the connection between MBII and one particular provider were severed, then the NSAPs within MBII with addresses based on that provider would become unreachable via inter-domain routing. The impact of this problem can be reduced somewhat by maintenance of additional information within routing tables, but this reduces the scaling advantage of the second approach.

The second solution also requires that when external connectivity changes, internal addresses also change.

Also note that this and the previous approach will tend to cause packets to take different routes. With the first approach, packets from outside of MBII destined for within MBII will tend to enter via the point which is closest to the source (which will therefore tend to maximize the load on the networks internal to MBII). With the second solution, packets from outside destined for within MBII will tend to enter via the point which is closest to the destination

(which will tend to minimize the load on the networks within MBII, and maximize the load on the providers).

These solutions also have different effects on policies. For example, suppose that country "X" has a law that traffic from a source within country X to a destination within country X must at all times stay entirely within the country. With the first solution, it is not possible to determine from the destination address whether or not the destination is within the country. With the second solution, a separate address may be assigned to those NSAPs which are within country X, thereby allowing routing policies to be followed. Similarly, suppose that "Little Small Company" (LSC) has a policy that its packets may never be sent to a destination that is within MBII. With either solution, the routers within LSC may be configured to discard any traffic that has a destination within MBII's address space. However, with the first solution this requires one entry; with the second it requires many entries and may be impossible as a practical matter.

There are other possible solutions as well. A third approach is to assign each multi-homed organization a single address prefix, based on one of its connections to a provider. Other providers to which the multi-homed organization are attached maintain a routing table entry for the organization, but are extremely selective in terms of which indirect providers are told of this route. This approach will produce a single "default" routing entry which all providers will know how to reach the organization (since presumably all providers will maintain routes to each other), while providing more direct routing in those cases where providers agree to maintain additional routing information.

There is at least one situation in which this third approach is particularly appropriate. Suppose that a special interest group of organizations have deployed their own backbone. For example, let's suppose that the U.S. National Widget Manufacturers and Researchers have set up a U.S.-wide backbone, which is used by corporations who manufacture widgets, and certain universities which are known for their widget research efforts. We can expect that the various organizations which are in the widget group will run their internal networks as separate routing domains, and most of them will also be attached to other providers (since most of the organizations involved in widget manufacture and research will also be involved in other activities). We can therefore expect that many or most of the organizations in the widget group are dual-homed, with one attachment for widget-associated communications and the other attachment for other types of communications. Let's also assume that the total number of organizations involved in the widget group is small enough that it is reasonable to maintain a routing table containing one

entry per organization, but that they are distributed throughout a larger internet with many millions of (mostly not widget-associated) routing domains.

With the third approach, each multi-homed organization in the widget group would make use of an address assignment based on its other attachment(s) to providers (the attachments not associated with the widget group). The widget backbone would need to maintain routes to the routing domains associated with the various member organizations. Similarly, all members of the widget group would need to maintain a table of routes to the other members via the widget backbone. However, since the widget backbone does not inform other general world-wide providers of what addresses it can reach (since the backbone is not intended for use by other outside organizations), the relatively large set of routing prefixes needs to be maintained only in a limited number of places. The addresses assigned to the various organizations which are members of the widget group would provide a "default route" via each members other attachments to providers, while allowing communications within the widget group to use the preferred path.

A fourth solution involves assignment of a particular address prefix for routing domains which are attached to two or more specific cooperative public service providers. For example, suppose that there are two providers "SouthNorthNet" and "NorthSouthNet" which have a very large number of customers in common (i.e., there are a large number of routing domains which are attached to both). Rather than getting two address prefixes (such as two AA values assigned under the GOSIP address space) these organizations could obtain three prefixes. Those routing domains which are attached to NorthSouthNet but not attached to SouthNorthNet obtain an address assignment based on one of the prefixes. Those routing domains which are attached to SouthNorthNet but not to NorthSouthNet would obtain an address based on the second prefix. Finally, those routing domains which are multi-homed to both of these networks would obtain an address based on the third prefix. Each of these two providers would then advertise two prefixes to other providers, one prefix for subscriber routing domains attached to it only, and one prefix for subscriber routing domains attached to both.

This fourth solution could become important when use of public data networks becomes more common. In particular, it is likely that at some point in the future a substantial percentage of all routing domains will be attached to public data networks. In this case, nearly all government-sponsored networks (such as some regional networks which receive funding from NSF, as well as government sponsored backbones) may have a set of customers which overlaps substantially with the public networks.

There are therefore a number of possible solutions to the problem of assigning NSAP addresses to multi-homed routing domains. Each of these solutions has very different advantages and disadvantages. Each solution places a different real (i.e., financial) cost on the multi-homed organizations, and on the providers (including those to which the multi-homed organizations are not attached).

In addition, most of the solutions described also highlight the need for each provider to develop policy on whether and under what conditions to accept customers with addresses that are not based on its own address prefix, and how such non-local addresses will be treated. For example, a somewhat conservative policy might be that an attached subscriber RD may use any NSAP address prefix, but that addresses which are not based on the providers own prefix might not be advertised to other providers. In a less conservative policy, a provider might accept customers using such non-local prefixes and agree to exchange them in routing information with a defined set of other providers (this set could be an a priori group of providers that have something in common such as geographical location, or the result of an agreement specific to the requesting subscriber). Various policies involve real costs to providers, which may be reflected in those policies.

5.5. Private Links

The discussion up to this point concentrates on the relationship between NSAP addresses and routing between various routing domains over transit routing domains, where each transit routing domain interconnects a large number of routing domains and offers a more-or-less public service.

However, there may also exist a large number of private point-to-point links which interconnect two private routing domains. In many cases such private point-to-point links may be limited to forwarding packets directly between the two private routing domains.

For example, let's suppose that the XYZ corporation does a lot of business with MBII. In this case, XYZ and MBII may contract with a carrier to provide a private link between the two corporations, where this link may only be used for packets whose source is within one of the two corporations, and whose destination is within the other of the two corporations. Finally, suppose that the point-to-point link is connected between a single router (router X) within XYZ corporation and a single router (router M) within MBII. It is therefore necessary to configure router X to know which addresses can be reached over this link (specifically, all addresses reachable in MBII). Similarly, it is necessary to configure router M to know which addresses can be reached over this link (specifically, all

addresses reachable in XYZ Corporation).

The important observation to be made here is that such private links may be ignored for the purpose of NSAP allocation, and do not pose a problem for routing. This is because the routing information associated with private links is not propagated throughout the internet, and therefore does not need to be collapsed into a provider's prefix.

In our example, let's suppose that the XYZ corporation has a single connection to a service provider, and has therefore received an address allocation from the space administered by that provider. Similarly, let's suppose that MBII, as an international corporation with connections to six different providers, has chosen the second solution from Section 5.4, and therefore has obtained six different address allocations. In this case, all addresses reachable in the XYZ Corporation can be described by a single address prefix (implying that router M only needs to be configured with a single address prefix to represent the addresses reachable over this point-to-point link). All addresses reachable in MBII can be described by six address prefixes (implying that router X needs to be configured with six address prefixes to represent the addresses reachable over the point-to-point link).

In some cases, such private point-to-point links may be permitted to forward traffic for a small number of other routing domains, such as closely affiliated organizations. This will increase the configuration requirements slightly. However, provided that the number of organizations using the link is relatively small, then this still does not represent a significant problem.

Note that the relationship between routing and NSAP addressing described in other sections of this paper is concerned with problems in scaling caused by large, essentially public transit routing domains which interconnect a large number of routing domains. However, for the purpose of NSAP allocation, private point-to-point links which interconnect only a small number of private routing domains do not pose a problem, and may be ignored. For example, this implies that a single subscriber routing domain which has a single connection to a "public" provider, plus a number of private point-to-point links to other subscriber routing domains, can be treated as if it were single-homed to the provider for the purpose of NSAP address allocation.

5.6. Zero-Homed Routing Domains

Currently, a very large number of organizations have internal communications networks which are not connected to any external network. Such organizations may, however, have a number of private point-to-point links that they use for communications with other organizations. Such organizations do not participate in global routing, but are satisfied with reachability to those organizations with which they have established private links. These are referred to as zero-homed routing domains.

Zero-homed routing domains can be considered as the degenerate case of routing domains with private links, as discussed in the previous section, and do not pose a problem for inter-domain routing. As above, the routing information exchanged across the private links sees very limited distribution, usually only to the RD at the other end of the link. Thus, there are no address abstraction requirements beyond those inherent in the address prefixes exchanged across the private link.

However, it is important that zero-homed routing domains use valid globally unique NSAP addresses. Suppose that the zero-homed routing domain is connected through a private link to an RD. Further, this RD participates in an internet that subscribes to the global OSI addressing plan (i.e., ISO 8348). This RD must be able to distinguish between the zero-homed routing domain's NSAPs and any other NSAPs that it may need to route to. The only way this can be guaranteed is if the zero-homed routing domain uses globally unique NSAPs.

5.7. Address Transition Issues

Allocation of NSAP addresses based on connectivity to providers is important to allow scaling of inter-domain routing to an internet containing millions of routing domains. However, such address allocation based on topology also implies that a change in topology may result in a change of address.

This need to allow for change in addresses is a natural, inevitable consequence of any method for routing data abstraction. The basic notion of routing data abstraction is that there is some correspondence between the address and where a system (i.e., a routing domain, area, or end system) is located. Thus if the system moves, in some cases the address will have to change. If it were possible to change the connectivity between routing domains without changing the addresses, then it would clearly be necessary to keep track of the location of that routing domain on an individual basis.

Because of the rapid growth and increased commercialization of the Internet, it is possible that the topology may be relatively volatile. This implies that planning for address transition is very important. Fortunately, there are a number of steps which can be taken to help ease the effort required for address transition. A complete description of address transition issues is outside of the scope of this paper. However, a very brief outline of some transition issues is contained in this section.

Also note that the possible requirement to transition addresses based on changes in topology imply that it is valuable to anticipate the future topology changes before finalizing a plan for address allocation. For example, in the case of a routing domain which is initially single-homed, but which is expecting to become multi-homed in the future, it may be advantageous to assign NSAP addresses based on the anticipated future topology.

In general, it will not be practical to transition the NSAP addresses assigned to a routing domain in an instantaneous "change the address at midnight" manner. Instead, a gradual transition is required in which both the old and the new addresses will remain valid for a limited period of time. During the transition period, both the old and new addresses are accepted by the end systems in the routing domain, and both old and new addresses must result in correct routing of packets to the destination.

Provision for transition has already been built into IS-IS. As described in Section 3, IS-IS allows multiple addresses to be assigned to each area specifically for the purpose of easing transition.

Similarly, there are provisions in OSI for the autoconfiguration of area addresses. This allows OSI end systems to find out their area addresses automatically, either by passively observing the ES-IS IS-Hello packets transmitted by routers, or by actively querying the routers for their NSAP address. If the ID portion of the address is assigned in a manner which allows for globally unique IDs [18], then an end system can reconfigure its entire NSAP address automatically without the need for manual intervention. However, routers will still require manual address reconfiguration.

During the transition period, it is important that packets using the old address be forwarded correctly, even when the topology has changed. This is facilitated by the use of "best match" inter-domain routing.

For example, suppose that the XYZ Corporation was previously connected only to the NorthSouthNet provider. The XYZ Corporation

therefore went off to the NorthSouthNet administration and got a routing domain assignment based on the AA value obtained by the NorthSouthNet under the GOSIP address space. However, for a variety of reasons, the XYZ Corporation decided to terminate its association with the North-SouthNet, and instead connect directly to the NewCommercialNet public data network. Thus the XYZ Corporation now has a new address assignment under the ANSI address assigned to the NewCommercialNet. The old address for the XYZ Corporation would seem to imply that traffic for the XYZ Corporation should be routed to the NorthSouthNet, which no longer has any direct connection with XYZ Corporation.

If the old provider (NorthSouthNet) and the new provider (NewCommercialNet) are adjacent and cooperative, then this transition is easy to accomplish. In this case, packets routed to the XYZ Corporation using the old address assignment could be routed to the NorthSouthNet, which would directly forward them to the NewCommercialNet, which would in turn forward them to XYZ Corporation. In this case only NorthSouthNet and NewCommercialNet need be aware of the fact that the old address refers to a destination which is no longer directly attached to NorthSouthNet.

If the old provider and the new provider are not adjacent, then the situation is a bit more complex, but there are still several possible ways to forward traffic correctly.

If the old provider and the new provider are themselves connected by other cooperative providers, then these intermediate domains may agree to forward traffic for XYZ correctly. For example, suppose that NorthSouthNet and NewCommercialNet are not directly connected, but that they are both directly connected to the NSFNET backbone. In this case, all three of NorthSouthNet, NewCommercialNet, and the NSFNET backbone would need to maintain a special entry for XYZ corporation so that traffic to XYZ using the old address allocation would be forwarded via NewCommercialNet. However, other routing domains would not need to be aware of the new location for XYZ Corporation.

Suppose that the old provider and the new provider are separated by a non-cooperative routing domain, or by a long path of routing domains. In this case, the old provider could encapsulate traffic to XYZ Corporation in order to deliver such packets to the correct backbone.

Also, those locations which do a significant amount of business with XYZ Corporation could have a specific entry in their routing tables added to ensure optimal routing of packets to XYZ. For example, suppose that another commercial backbone "OldCommercialNet" has a large number of customers which exchange traffic with XYZ

Corporation, and that this third provider is directly connected to both NorthSouthNet and NewCommercialNet. In this case OldCommercialNet will continue to have a single entry in its routing tables for other traffic destined for NorthSouthNet, but may choose to add one additional (more specific) entry to ensure that packets sent to XYZ Corporation's old address are routed correctly.

Whichever method is used to ease address transition, the goal is that knowledge relating XYZ to its old address that is held throughout the global internet would eventually be replaced with the new information. It is reasonable to expect this to take weeks or months and will be accomplished through the distributed directory system. Discussion of the directory, along with other address transition techniques such as automatically informing the source of a changed address, are outside the scope of this paper.

6. Recommendations

We anticipate that the current exponential growth of the Internet will continue or accelerate for the foreseeable future. In addition, we anticipate a continuation of the rapid internationalization of the Internet. The ability of routing to scale is dependent upon the use of data abstraction based on hierarchical NSAP addresses. As CLNP use increases in the Internet, it is therefore essential to assign NSAP addresses with great care.

It is in the best interests of the internetworking community that the cost of operations be kept to a minimum where possible. In the case of NSAP allocation, this again means that routing data abstraction must be encouraged.

In order for data abstraction to be possible, the assignment of NSAP addresses must be accomplished in a manner which is consistent with the actual physical topology of the Internet. For example, in those cases where organizational and administrative boundaries are not related to actual network topology, address assignment based on such organization boundaries is not recommended.

The intra-domain IS-IS routing protocol allows for information abstraction to be maintained at two levels: systems are grouped into areas, and areas are interconnected to form a routing domain. The inter-domain IDRP routing protocol allows for information abstraction to be maintained at multiple levels by grouping routing domains into Routing Domain Confederations and using route aggregation capabilities.

For zero-homed and single-homed routing domains (which are expected to remain zero-homed or single-homed), we recommend that the NSAP

addresses assigned for OSI use within a single routing domain use a single address prefix assigned to that domain. Specifically, this allows the set of all NSAP addresses reachable within a single domain to be fully described via a single prefix. We recommend that single-homed routing domains use an address prefix based on its connectivity to a public service provider. We recommend that zero-homed routing domains use globally unique addresses.

We anticipate that the total number of routing domains existing on a worldwide OSI Internet to be great enough that additional levels of hierarchical data abstraction beyond the routing domain level will be necessary. To provide the needed data abstraction we recommend to use Routing Domain Confederations and route aggregation capabilities of IDRP.

The general technical requirements for NSAP address guidelines do not vary from country to country. However, details of address administration may vary between countries. Also, in most cases, network topology will have a close relationship with national boundaries. For example, the degree of network connectivity will often be greater within a single country than between countries. It is therefore appropriate to make specific recommendations based on national boundaries, with the understanding that there may be specific situations where these general recommendations need to be modified. Moreover, that suggests that national boundaries may be used to group domains into Routing Domain Confederations.

Each of the country-specific or continent-specific recommendations presented below are consistent with the technical requirements for scaling of addressing and routing presented in this RFC.

6.1. Recommendations Specific to U.S. Parts of the Internet

NSAP addresses for use within the U.S. portion of the Internet are expected to be based primarily on two address prefixes: the ICD=0005 format used by The U.S. Government, and the DCC=840 format defined by ANSI.

We anticipate that, in the U.S., public interconnectivity between private routing domains will be provided by a diverse set of providers, including (but not necessarily limited to) regional providers and commercial Public Data Networks.

These networks are not expected to be interconnected in a strictly hierarchical manner. For example, the regional providers may be directly connected rather than rely on an indirect provider, and all three of these types of networks may have direct international connections.

However, the total number of such providers is expected to remain (for the foreseeable future) small enough to allow addressing of this set of providers via a flat address space. These providers will be used to interconnect a wide variety of routing domains, each of which may comprise a single corporation, part of a corporation, a university campus, a government agency, or other organizational unit.

In addition, some private corporations may be expected to make use of dedicated private providers for communication within their own corporations.

We anticipate that the great majority of routing domains will be attached to only one of the providers. This will permit hierarchical address abbreviation based on provider. We therefore strongly recommend that addresses be assigned hierarchically, based on address prefixes assigned to individual providers.

For the GOSIP address format, this implies that Administrative Authority (AA) identifiers should be obtained by all providers (explicitly including the NSFNET backbone, the NSFNET regionals, and other major government backbones). For those subscriber routing domains which are connected to a single provider, they should be assigned a Routing Domain (RD) value from the space assigned to that provider.

To provide routing information aggregation/abstraction we recommend that each provider together with all of its subscriber domains form a Routing Domain Confederation. That, combined with hierarchical address assignment, would provide significant reduction in the volume of routing information that needs to be handled by IDRP. Note that the presence of multihomed subscriber domains would imply that such Confederations will overlap, which is explicitly supported by IDRP.

We recommend that all providers explicitly be involved in the task of address administration for those subscriber routing domains which are single-homed to them. This offers a valuable service to their customers, and also greatly reduces the resources (including human and network resources) necessary for that provider to take part in inter-domain routing.

Each provider should develop policy on whether and under what conditions to accept customers using addresses that are not based on the provider's own address prefix, and how such non-local addresses will be treated. Policies should reflect the issue of cost associated with implementing such policies.

We recommend that a similar hierarchical model be used for NSAP addresses using the DCC-based address format. The structure for

DCC=840-based NSAPs is provided in Section A.2.

For routing domains which are not attached to any publically-available provider, no urgent need for hierarchical address abbreviation exists. We do not, therefore, make any additional recommendations for such "isolated" routing domains, except to note that there is no technical reason to preclude assignment of GOSIP AA identifier values or ANSI organization identifiers to such domains. Where such domains are connected to other domains by private point-to-point links, and where such links are used solely for routing between the two domains that they interconnect, no additional technical problems relating to address abbreviation is caused by such a link, and no specific additional recommendations are necessary.

6.2. Recommendations Specific to European Parts of the Internet

This section contains additional RARE recommendations for allocating NSAP addresses within each national domain, administered by a National Standardization Organization (NSO) and national research network organizations.

NSAP addresses are expected to be based on the ISO DCC scheme. Organizations which are not associated with a particular country and which have reasons not to use a national prefix based on ISO DCC should follow the recommendations covered in chapters 6.3 and 6.4.

ISO DCC addresses are not associated with any specific subnetwork type and service provider and are thus independent of the type or ownership of the underlying technology.

6.2.1. General NSAP Structure

The general structure of a Network Address defined in ISO 8348 is further divided into:

	IDP		DSP					
	AFI	IDI	CDP		CDSP			
	AFI	IDI	CFI	CDI	RDAA		ID	SEL
octets	1	2	2..4		0..13		1..8	1

IDP Initial Domain Part
 AFI Authority and Format Identifier, two-decimal-digit, 38 for decimal abstract syntax of the DSP or 39 for binary abstract syntax of the DSP
 IDI Initial Domain Identifier, a three-decimal-digit country code, as defined in ISO 3166
 DSP Domain Specific Part
 CDP Country Domain Part, 2..4 octets
 CFI Country Format Identifier, one digit
 CDI Country Domain Identifier, 3 to 7 digits, fills CDP to an octet boundary
 CDSP Country Domain Specific Part
 RDAA Routing Domain and Area Address
 ID System Identifier (1..8 octet)
 SEL NSAP Selector

The total length of an NSAP can vary from 7 to 20 octets.

6.2.2. Structure of the Country Domain Part

The CDP identifies an organization within a country and the CDSP is then available to that organization for further internal structuring as it wishes. Non-ambiguity of addresses is ensured by there being the NSO a single national body that allocates the CDPs.

The CDP is further divided into CFI and CDI, where the CFI identifies the format of the CDI. The importance of this is that it enables several types of CDI to be assigned in parallel, corresponding to organizations with different requirements and giving different amounts of the total address space to them, and that it conveniently enables a substantial amount of address space to be reserved for future allocation.

The possible structures of the CDP are as follows:

CFI = /0	reserved
CFI = /1 CDI = /aaa	very large organizations or trade associations
CFI = /2 CDI = /aaaaa	organizations of intermediate size
CFI = /3 CDI = /aaaaaaa	small organizations and single users
CFI = /4../F	reserved

Note: this uses the hexadecimal reference publication format defined in ISO 8348 of a solidus "/" followed by a string of hexadecimal digits. Each "a" represents a hexadecimal digit.

Organizations are classified into large, medium and small for the purpose of address allocation, and one CFI is made available for each category of organization.

This recommendation for CDP leaves space for the U.S. GOSIP Version 2 NSAP model (Appendix A.1) by the reserved CFI /8, nevertheless it is not recommended for use in the European Internet.

6.2.3. Structure of the Country Domain Specific Part

The CDSP must have a structure (within the decimal digit or binary octet syntax selected by the AFI value 38 or 39) satisfying both the routing requirements (IS-IS) and the logical requirements of the organization identified (CFI + CDI).

6.3. Recommendations Specific to Other Parts of the Internet

For the part of the Internet which is outside of the U.S. and Europe, it is recommended that the DSP format be structured hierarchically similarly to that specified within the U.S. and Europe no matter whether the addresses are based on DCC or ICD format.

Further, in order to allow aggregation of NSAPs at national boundaries into as few prefixes as possible, we further recommend that NSAPs allocated to routing domains should be assigned based on each routing domain's connectivity to a national Internet backbone.

6.4. Recommendations for Multi-Homed Routing Domains

Some routing domains will be attached to multiple providers within the same country, or to providers within multiple countries. We refer to these as "multi-homed" routing domains. Clearly the strict hierarchical model discussed above does not neatly handle such routing domains.

There are several possible ways that these multi-homed routing domains may be handled. Each of these methods vary with respect to the amount of information that must be maintained for inter-domain routing and also with respect to the inter-domain routes. In addition, the organization that will bear the brunt of this cost varies with the possible solutions. For example, the solutions vary with respect to:

- * resources used within routers within the providers;
- * administrative cost on provider personnel; and,
- * difficulty of configuration of policy-based inter-domain routing information within subscriber routing domains.

Also, the solution used may affect the actual routes which packets follow, and may effect the availability of backup routes when the primary route fails.

For these reasons it is not possible to mandate a single solution for all situations. Rather, economic considerations will require a variety of solutions for different subscriber routing domains and providers.

6.5. Recommendations for RDI and RDCI assignment

While RDIs and RDCIs need not be related to the set of addresses within the domains (confederations) they depict, for the sake of simplicity we recommend that RDIs and RDCIs be assigned based on the NSAP prefixes assigned to domains and confederations.

A subscriber RD should use the NSAP prefix assigned to it as its RDI. A multihomed RD should use one of the NSAP prefixes assigned to it as its RDI. If a service provider forms a Routing Domain Confederation with some of its subscribers and the subscribers take their addresses out of the provider, then the NSAP prefix assigned to the provider should be used as the RDCI of the confederation. In this case the provider may use a longer NSAP prefix for its own RDIs. In all other cases a provider should use the address prefix that it uses for assigning addresses to systems within the provider as its RDI.

7. Security Considerations

Security issues are not discussed in this memo (except for the discussion of IS-IS authentication in Section 3.2).

8. Authors' Addresses

Richard P. Colella
National Institute of Standards & Technology
Building 225/Room B217
Gaithersburg, MD 20899

Phone: (301) 975-3627
EMail: colella@nist.gov

Ross Callon
c/o Wellfleet Communications, Inc
2 Federal Street
Billerica, MA 01821

Phone: (508) 436-3936
EMail: callon@wellfleet.com

Ella P. Gardner
The MITRE Corporation
7525 Colshire Drive
McLean, VA 22102-3481

Phone: (703) 883-5826
EMail: ep@gateway.mitre.org

Yakov Rekhter
T.J. Watson Research Center, IBM Corporation
P.O. Box 218
Yorktown Heights, NY 10598

Phone: (914) 945-3896
EMail: yakov@watson.ibm.com

9. Acknowledgments

The authors would like to thank the members of the IETF OSI-NSAP Working Group and of RARE WG4 for the helpful suggestions made during the writing of this paper. We would also like to thank Radia Perlman of Novell, Marcel Wiget of SWITCH, and Cathy Wittbrodt of BARRnet for their ideas and help.

10. References

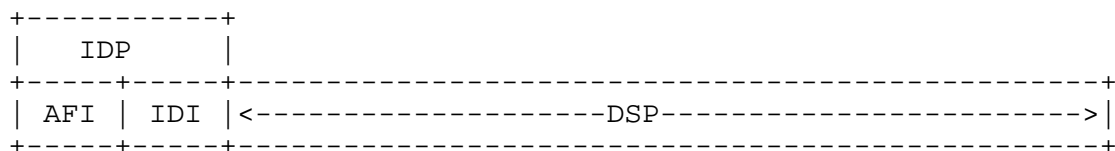
- [1] ANSI, "American National Standard for the Structure and Semantics of the Domain-Specific Part (DSP) of the OSI Network Service Access Point (NSAP) Address", American National Standard X3.216-1992.
- [2] Boland, T., "Government Open Systems Interconnection Profile Users' Guide Version 2 [DRAFT]", NIST Special Publication, National Institute of Standards and Technology, Computer Systems Laboratory, Gaithersburg, MD, June 1991.
- [3] GOSIP Advanced Requirements Group, "Government Open Systems Interconnection Profile (GOSIP) Version 2", Federal Information Processing Standard 146-1, U.S. Department of Commerce, National Institute of Standards and Technology, Gaithersburg, MD, April 1991.
- [4] Hemrick, C., "The OSI Network Layer Addressing Scheme, Its Implications, and Considerations for Implementation", NTIA Report 85186, U.S. Department of Commerce, National Telecommunications and Information Administration, 1985.
- [5] ISO, "Addendum to the Network Service Definition Covering Network Layer Addressing," RFC 941, ISO, April 1985.
- [6] ISO/IEC, "Codes for the Representation of Names of Countries", International Standard 3166, ISO/IEC JTC 1, Switzerland, 1984.
- [7] ISO/IEC, "Data Interchange - Structures for the Identification of Organization", International Standard 6523, ISO/IEC JTC 1, Switzerland, 1984.
- [8] ISO/IEC, "Information Processing Systems - Open Systems Interconnection -- Basic Reference Model", International Standard 7498, ISO/IEC JTC 1, Switzerland, 1984.
- [9] ISO/IEC, "Protocol for Providing the Connectionless-mode Network Service", International Standard 8473, ISO/IEC JTC 1, Switzerland, 1986.
- [10] ISO/IEC, "End System to Intermediate System Routing Exchange Protocol for use in Conjunction with the Protocol for the Provision of the Connectionless-mode Network Service", International Standard 9542, ISO/IEC JTC 1, Switzerland, 1987.

- [11] ISO/IEC, "Information Processing Systems -- Data Communications -- Network Service Definition", International Standard 8348, 1992.
- [12] ISO/IEC, "Information Processing Systems - OSI Reference Model - Part3: Naming and Addressing", Draft International Standard 7498-3, ISO/IEC JTC 1, Switzerland, March 1989.
- [13] ISO/IEC, "Information Technology - Telecommunications and Information Exchange Between Systems - OSI Routeing Framework", Technical Report 9575, ISO/IEC JTC 1, Switzerland, 1989.
- [14] ISO/IEC, "Intermediate System to Intermediate System Intra-Domain Routeing Exchange Protocol for use in Conjunction with the Protocol for Providing the Connectionless-Mode Network Service (ISO 8473)", International Standard ISO/IEC 10589, 1992.
- [15] Loughheed, K., and Y. Rekhter, "A Border Gateway Protocol 3 (BGP-3)" RFC 1267, cisco Systems, T.J. Watson Research Center, IBM Corp., October 1991.
- [16] ISO/IEC, "Protocol for Exchange of Inter-Domain Routeing Information among Intermediate Systems to support Forwarding of ISO 8473 PDUs", International Standard 10747, ISO/IEC JTC 1, Switzerland 1993.
- [17] Callon, R., "TCP and UDP with Bigger Addresses (TUBA), A Simple Proposal for Internet Addressing and Routing", RFC 1347, DEC, June 1992.
- [18] Piscitello, D., "Assignment of System Identifiers for TUBA/CLNP Hosts", RFC 1526, Bellcore, September 1993.
- [19] Fuller, V., Li, T., Yu, J., and K. Varadhan, "Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy", RFC 1519, BARRNet, cisco, OARnet, September 1993.
- [20] ISO/IEC JTC1/SC6, "Addendum to ISO 9542 Covering Address Administration", N6273, March 1991.

A. Administration of NSAPs

NSAPs represent the endpoints of communication through the Network Layer and must be globally unique [4]. ISO 8348 defines the semantics of the NSAP and the abstract syntaxes in which the semantics of the Network address can be expressed [11].

The NSAP consists of the initial domain part (IDP) and the domain specific part (DSP). The initial domain part of the NSAP consists of an authority and format identifier (AFI) and an initial domain identifier (IDI). The AFI specifies the format of the IDI, the network addressing authority responsible for allocating values of the IDI, and the abstract syntax of the DSP. The IDI specifies the addressing subdomain from which values of the DSP are allocated and the network addressing authority responsible for allocating values of the DSP from that domain. The structure and semantics of the DSP are determined by the authority identified by the IDI. Figure 3 shows the NSAP address structure.



IDP Initial Domain Part
 AFI Authority and Format Identifier
 IDI Initial Domain Identifier
 DSP Domain Specific Part

Figure 3: NSAP address structure.

The global network addressing domain consists of all the NSAP addresses in the OSI environment. Within that environment, seven second-level addressing domains and corresponding IDI formats are described in ISO 8348:

- * X.121 for public data networks
- * F.69 for telex
- * E.163 for the public switched telephone network numbers
- * E.164 for ISDN numbers
- * ISO Data Country Code (DCC), allocated according to ISO 3166 [6]

- * ISO International Code Designator (ICD), allocated according to ISO 6523 [7]
- * Local to accommodate the coexistence of OSI and non-OSI network addressing schemes.

For OSI networks in the U.S., portions of the ICD subdomain are available for use through the U.S. Government, and the DCC subdomain is available for use through The American National Standards Institute (ANSI). The British Standards Institute is the registration authority for the ICD subdomain, and has registered four IDIs for the U.S. Government: those used for GOSIP, DoD, OSINET, and the OSI Implementors Workshop. ANSI, as the U.S. ISO Member Body, is the registration authority for the DCC domain in the United States.

A.1 GOSIP Version 2 NSAPs

GOSIP Version 2 makes available for government use an NSAP addressing subdomain with a corresponding address format as illustrated in Figure 2 in Section 4.2. The "47" signifies that it is based on the ICD format and uses a binary syntax for the DSP. The 0005 is an IDI value which has been assigned to the U.S. Government. Although GOSIP Version 2 NSAPs are intended primarily for U.S. Government use, requests from non-government and non-U.S. organizations will be considered on a case-by-case basis.

The format for the DSP under ICD=0005 has been established by the National Institute of Standards and Technology (NIST), the authority for the ICD=0005 domain, in GOSIP Version 2 [3] (see Figure 2, Section 4.2). NIST has delegated the authority to register AA identifiers for GOSIP Version 2 NSAPs to the General Services Administration (GSA).

ISO 8348 allows a maximum length of 20 octets for the NSAP address. The AFI of 47 occupies one octet, and the IDI of 0005 occupies two octets. The DSP is encoded as binary as indicated by the AFI of 47. One octet is allocated for a DSP Format Identifier, three octets for an Administrative Authority identifier, two octets for Routing Domain, two octets for Area, six octets for the System Identifier, and one octet for the NSAP selector. Note that two octets have been reserved to accommodate future growth and to provide additional flexibility for inter-domain routing. The last seven octets of the GOSIP NSAP format are structured in accordance with IS-IS [14], the intra-domain IS-IS routing protocol. The DSP Format Identifier (DFI) identifies the format of the remaining DSP structure and may be used in the future to identify additional DSP formats; the value 80h in the DFI identifies the GOSIP Version 2 NSAP structure.

The Administrative Authority identifier names the administrative authority which is responsible for registration within its domain. The administrative authority may delegate the responsibility for registering areas to the routing domains, and the routing domains may delegate the authority to register System Identifiers to the areas. The main responsibility of a registration authority at any level of the addressing hierarchy is to assure that names of entities are unambiguous, i.e., no two entities have the same name. The registration authority is also responsible for advertising the names.

A routing domain is a set of end systems and intermediate systems which operate according to the same routing procedures and is wholly contained within a single administrative domain. An area uniquely identifies a subdomain of the routing domain. The system identifier names a unique system within an area. The value of the system field may be a physical address (SNPA) or a logical value. Address resolution between the NSAP and the SNPA may be accomplished by an ES-IS protocol [10], locally administered tables, or mapping functions. The NSAP selector field identifies the end user of the network layer service, i.e., a transport layer entity.

A.1.1 Application for Administrative Authority Identifiers

The steps required for an agency to acquire an NSAP Administrative Authority identifier under ICD=0005 from GSA will be provided in the updated GOSIP users' guide for Version 2 [2] and are given below. Requests from non-government and non-U.S. organizations should originate from a senior official, such as a vice-president or chief operating officer.

- * Identify all end systems, intermediate systems, subnetworks, and their topological and administrative relationships.
- * Designate one individual (usually the agency head) within an agency to authorize all registration requests from that agency (NOTE: All agency requests must pass through this individual).
- * Send a letter on agency letterhead and signed by the agency head to GSA:

Telecommunications Customer Requirements Office
U.S. General Services Administration
Information Resource Management Service
Office of Telecommunications Services
18th and F Streets, N.W.
Washington, DC 20405
Fax +1 202 208-5555

The letter should contain the following information:

- Requestor's Name and Title,
 - Organization,
 - Postal Address,
 - Telephone and Fax Numbers,
 - Electronic Mail Address(es), and,
 - Reason Needed (one or two paragraphs explaining the intended use).
- * If accepted, GSA will send a return letter to the agency head indicating the NSAP Administrative Authority identifier assigned, effective date of registration, and any other pertinent information.
 - * If rejected, GSA will send a letter to the agency head explaining the reason for rejection.
 - * Each Authority will administer its own subaddress space in accordance with the procedures set forth by the GSA in Section A.1.2.
 - * The GSA will maintain, publicize, and disseminate the assigned values of Administrative Authority identifiers unless specifically requested by an agency not to do so.

A.1.2 Guidelines for NSAP Assignment

Recommendations which should be followed by an administrative authority in making NSAP assignments are given below.

- * The authority should determine the degree of structure of the DSP under its control. Further delegation of address assignment authority (resulting in additional levels of hierarchy in the NSAP) may be desired.
- * The authority should make sure that portions of NSAPs that it specifies are unique, current, and accurate.
- * The authority should ensure that procedures exist for disseminating NSAPs to routing domains and to areas within each routing domain.
- * The systems administrator must determine whether a logical or a physical address should be used in the System Identifier field (Figure 2, Section 4.2). An example of a physical address is a 48-bit MAC address; a logical address is merely a number that meets the uniqueness requirements for the System Identifier field, but bears no relationship to an address on a physical subnetwork. We recommend that IDs should be assigned to be globally unique, as made possible by the method described in [18].
- * The network address itself contains information that may be used to aid routing, but does not contain a source route [12]. Information that enables next-hop determination based on NSAPs is gathered and maintained by each intermediate system through routing protocol exchanges.
- * GOSIP end systems and intermediate systems in federal agencies must be capable of routing information correctly to and from any subdomain defined by ISO 8348.
- * An agency may request the assignment of more than one Administrative Authority identifier. The particular use of each should be specified.

A.2 Data Country Code NSAPs

NSAPs from the Data Country Code (DCC) subdomain will also be common in the international Internet. ANS X3.216-1992 specifies the DSP structure under DCC=840 [1]. In the ANS, the DSP structure is identical to that specified in GOSIP Version 2, with the

Administrative Authority identifier replaced by the numeric form of the ANSI-registered organization name, as shown in Figure 4.

Referring to Figure 4, when the value of the AFI is 39, the IDI denotes an ISO DCC and the abstract syntax of the DSP is binary octets. The value of the IDI for the U.S. is 840, the three-digit numeric code for the United States under ISO 3166 [6]. The numeric form of organization name is analogous to the Administrative Authority identifier in the GOSIP Version 2 NSAP.

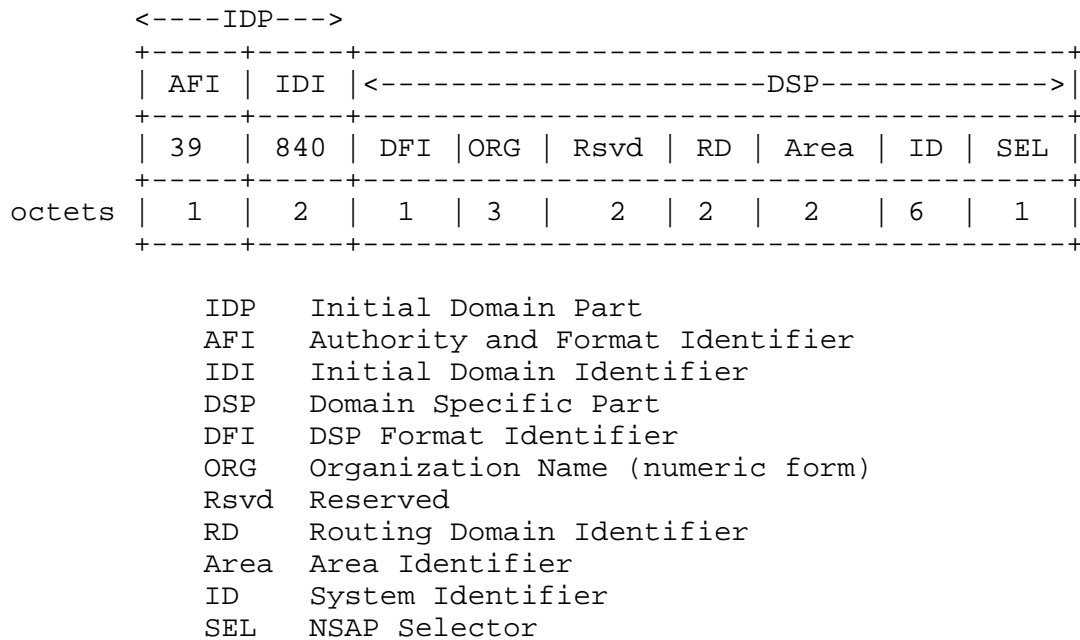


Figure 4: NSAP format for DCC=840 as proposed in ANSI X3S3.3.

A.2.1 Application for Numeric Organization Name

The procedures for registration of numeric organization names in the U.S. have been defined and are operational. To register a numeric organization name, the applicant must submit a request for registration and the \$1,000 (U.S.) fee to the registration authority, the American National Standards Institute (ANSI). ANSI will register a numeric value, along with the information supplied for registration, in the registration database. The registration information will be sent to the applicant within ten working days. The values for numeric organization names are assigned beginning at 113527.

The application form for registering a numeric organization name may be obtained from the ANSI Registration Coordinator at the following address:

Registration Coordinator
American National Standards Institute
11 West 42nd Street
New York, NY 10036
+1 212 642 4884 (tel)
+1 212 398 0023 (fax)
RFC822: mmaas@attmail.com
X.400: G=michelle; S=maas; A=attmail; C=us

Once an organization has registered with ANSI, it becomes a registration authority itself. In turn, it may delegate registration authority to routing domains, and these may make further delegations, for instance, from routing domains to areas. Again, the responsibilities of each Registration Authority are to assure that NSAPs within the domain are unambiguous and to advertise them as applicable.

A.3 Summary of Administrative Requirements

NSAPs must be globally unique, and an organization may assure this uniqueness for OSI addresses in two ways. The organization may apply to GSA for an Administrative Authority identifier. Although registration of Administrative Authority identifiers by GSA primarily serves U.S. Government agencies, requests for non-government and non-U.S. organizations will be considered on a case-by-case basis. Alternatively, the organization may apply to ANSI for a numeric organization name. In either case, the organization becomes the registration authority for its domain and can register NSAPs or delegate the authority to do so.

In the case of GOSIP Version 2 NSAPs, the complete DSP structure is given in GOSIP Version 2. For ANSI DCC-based NSAPs, the DSP structure is specified in ANS X3.216-1992. The DSP structure is identical to that specified in GOSIP Version 2.

