Network Working Group
Request for Comments: 1237

Richard Colella (NIST)
Ella Gardner (Mitre)
Ross Callon (DEC)

July 1991

Guidelines for OSI NSAP Allocation in the Internet

Status of This Memo

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

Abstract

The Internet is moving towards a multi-protocol environment that includes OSI. To support OSI in the Internet, an OSI lower layers infrastructure is required. This infrastructure comprises the connectionless network protocol (CLNP) 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 NSAPs in the Internet.

This document provides our current best judgment for the allocation of NSAP addresses in the Internet. This is intended to guide initial deployment of OSI 8473 (Connectionless Network Layer Protocol) in the Internet, as well as to solicit comments. It is expected that these guidelines may be further refined and this document updated as a result of experience gained during this initial deployment.

Contents

1 Introduction

The Internet is moving towards a multi-protocol environment that includes OSI. To support OSI in the Internet, an OSI lower layers infrastructure is required. This infrastructure comprises the connectionless network protocol (CLNP) [12] (see also RFC 994 [8]) 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 NSAPs in the Internet (NSAP and NSAP address are used interchangeably throughout this paper in referring to NSAP addresses).

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 NSAPs.

Section 4 addresses the specific relationship between NSAPs 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 allocation in the Internet.

Appendix A 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.

2 Scope

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 NSAPs; 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 ES-IS (end-system to intermediate system) routing protocol applicable for point-to-point links and LANs [13] (see also RFC 995 [7]) and the emerging intra-domain IS-IS protocol [17]. Also expected is the deployment of an inter-domain routing protocol similar to Border Gateway Protocol (BGP) [18].

The guidelines provided in this paper are intended for immediate deployment as CLNP is made available in the Internet. This paper specifically does not address long-term research issues, such as complex policy-based routing requirements.

In the current Internet many routing domains (such as corporate and campus networks) attach to transit networks (such as NSFNET regionals) in only one or a small number of carefully controlled access points. Addressing solutions which require substantial changes or constraints on the current topology are not considered.

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 DIS10589 IS-IS routing protocol;
- Benefits of some topological information in NSAPs to reduce routing protocol overhead;
- The anticipated need for additional levels of hierarchy in Internet addressing to support network growth;
- The recommended mapping between Internet topological entities (i.e., backbone networks, regional networks, and site networks) and OSI addressing and routing components;
- The recommended division of NSAP address assignment authority among backbones, regionals (also called mid-levels), and sites;
- 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 leaf routing domains that are connected to more than one regional or backbone.

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.

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 protocol in sufficient detail to understand the issues involved in NSAP allocation. Finally, the specific constraints that the intra-domain protocol places on NSAPs are listed.

3.1 OSI Routing Standards

OSI partitions the routing problem into three parts:

- routing exchanges between end systems and intermediate systems (ES-IS),
- routing exchanges between ISs in the same routing domain (intra-domain IS-IS), and,
- routing among routing domains (inter-domain IS-IS).

ES-IS, international standard ISO9542 [13] approved in 1987, is available in vendor products and is planned for the next release of Berkeley UNIX (UNIX is a trademark of AT&T). It is also cited in GOSIP Version 2 [4], which became effective in April 1991 for all applicable federal procurements, and mandatory beginning eighteen months later in 1992.

Intra-domain IS-IS advanced to draft international standard (DIS) status within ISO in November, 1990 as DIS10589 [17]. It is reasonable to expect that final text for the intra-domain IS-IS standard will be available by mid-1991.

There are two candidate proposals which address OSI inter-domain routing, ECMA TR/50 [3] and Border Router Protocol (BRP) [19], a direct derivative of the IETF Border Gateway Protocol [18]. ECMA TR/50 has been proposed as base text in the ISO/IEC JTC1 SC6/WG2 committee, which is responsible for the Network layer of the ISO Reference Model [11]. X3S3.3, the ANSI counterpart to WG2, has incorporated features of TR/50 into BRP and submitted this as alternate base text at the WG2 meeting in October, 1990. Currently, it is out for ISO Member Body comment. The proposed protocol is referred to as the Inter-domain Routing Protocol (IDRP) [20].

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 DIS10589

The IS-IS intra-domain routing protocol, DIS10589, developed in ISO, provides routing for OSI environments. In particular, DIS10589 is designed to work in conjunction with CLNP and ES-IS. This section briefly describes the manner in which DIS10589 operates.

In DIS10589, 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. Typically, a routing domain may consist of a corporate network, a university campus network, a regional network, 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 DIS10589 routing messages are sent on that link.

Currently, ISO does not have a standard for inter-domain routing (i.e., for routing between separate autonomous routing domains). In the interim, DIS10589 uses manual configuration. An inter-domain link is statically configured with the set of *address prefixes* reachable via that link, and with the method by which they can be reached (such as the DTE address to be dialed to reach that address, or the fact that the DTE address should be extracted from the OSI NSAP address).

DIS10589 routing makes use of two-level hierarchical routing. A routing domain is subdivided into areas (also known as level 1 subdomains). Level 1 ISs know the topology in their area, including all ISs and ESs in their area. However, level 1 ISs do not know the identity of ISs or destinations outside of their area. Level 1 ISs forward all traffic for destinations outside of their area to a level 2 IS within their area.

Similarly, level 2 ISs know the level 2 topology and know which addresses are reachable via each level 2 IS. The set of all level 2 ISs 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 ISs do not need to know the topology within any level 1 area, except to the extent that a level 2 IS may also be a level 1 IS within a single area. Only level 2 ISs can exchange data packets or routing information directly with external ISs located outside of their routing domain.

As illustrated in Figure 1, ISO addresses are subdivided into the Initial Domain Part (IDP) and the Domain Specific Part (DSP), as specified in ISO8348/Addendum 2, the OSI network layer addressing standard [14] (also RFC 941 [6]). 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). The DSP is further subdivided, by DIS10589, into a High Order Part of DSP (HO-DSP), a system identifier (ID), and an NSAP selector (SEL). The HO-DSP may use any format desired by the authority which is identified by the IDP. Together, the combination of [IDP,HO-DSP] identify an area within a routing domain and, implicitly, the routing domain containing the area. The combination of [IDP,HO-DSP] is therefore referred to as the *area address*.

IDP		DSP		
AFI	IDI	HO-DSP	ID	SEL

IDP	Initial Domain Part
AFI	Authority and Format Identifier
IDI	Initial Domain Identifier
DSP	Domain Specific Part
HO-DSP	High-order DSP
ID	System Identifier
SEL	NSAP Selector

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).

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 from having address A to having address 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 ESs 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 ESs (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 ISs to identify packets going to destinations outside of their area, which need to be forwarded to level 2 ISs. Thus, in DIS10589 the two types of ISs route as follows:

- Level 1 intermediate systems – these nodes route based on the ID portion of the ISO address. They route within an area. Level 1 ISs 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 IS.
- Level 2 intermediate systems – these nodes 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 ISs on the routing domain boundary that have advertised external address prefixes into the level 2 subdomain. A level 2 IS may also be operating as a level 1 IS in one area.

A level 1 IS will have the area portion of its address manually configured. It will refuse to become a neighbor with an IS whose area addresses do not overlap its own area addresses. However, if a level 1 IS 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 IS will accept another level 2 IS as a neighbor, regardless of area address. However, if the area addresses do not overlap, the link would be considered by both ISs 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 ISs in different routing domains.

DIS10589 provides an optional partition repair function. In the unlikely case that a level 1 area becomes partitioned, this function, if implemented, allows the partition to be repaired via use of level 2 routes.

DIS10589 requires that the set of level 2 ISs 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.

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

An ES may autoconfigure the area portion of its address by extracting the area portion of a neighboring IS's address. If this is the case, then an ES will always accept an IS as a neighbor. Since the standard does not specify that the end system *must* autoconfigure its area address, an end system may be pre-configured with an area address. In this case the end system would ignore IS neighbors with non-matching area addresses.

3.3 Requirements of DIS10589 on NSAPs

The preferred NSAP format for DIS10589 is shown in Figure 1. A number of points should be noted from DIS10589:

- The IDP is as specified in ISO 8348/Addendum 2, the OSI network layer addressing standard [14];
- The high-order portion of the DSP (HO-DSP) is that portion of the DSP whose assignment, structure, and meaning are not constrained by DIS10589;
- The concatenation of the IDP and the HO-DSP, the area address, must be globally unique (if the area address of an NSAP matches one of the area addresses of a system, it is in the system'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 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/Addendum 2. However, within a routing domain that conforms to DIS10589, 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 DIS10589 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 a correct "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 (which can range from one up to a maximum of three).

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 destinations can be minimized.

4 NSAPs and Routing

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 cpu, memory, and transmission bandwidth consumed in support of routing. This 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. It is necessary to find a balance between these two needs.

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 [16] and routing protocols, the lowest boundary at which this can occur is the boundary between an area and the level 2 subdomain within a DIS10589 routing domain. Data abstraction is designed into DIS10589 at this boundary, since level 1 ISs are constrained to reporting only area addresses, and a maximum number of three area addresses are allowed in one area (This is an architectural constant in DIS10589. See [17], Clause 7.2.11 and Table 2 of Clause 7.5.1).

Level 2 routing is based upon address prefixes. Level 2 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 ISs 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 (statically or dynamically) with other routing domains. 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. 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 routing domains each to be assigned an address prefix from a shorter prefix assigned to another routing domain whose function is to interconnect the group of routing domains. Each member of the group of routing domains 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 routing domains which are all attached only to a single regional (or backbone) domain, and which use that regional for all external (inter-domain) traffic. A small address prefix may be assigned to the regional, which then assigns slightly longer prefixes (based on the regional's prefix) to each of the routing domains that it interconnects. This allows the regional, when informing other routing domains 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 GOSIP Version 2 NSAP structure, discussed later in this section, allows for multiple levels of routing hierarchy.

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.

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 an example of how these two needs might be met. The AFI, IDI, DFI, and 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 Appendix A for more information on NSAP structure).

> [Note: It is not important 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 expected to be used, and will work equally well. For parts of the Internet outside of the U.S. there may in some cases be strong reasons to prefer a local format rather than the 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.

	←— IDP —→									
	AFI	IDI	←— DSP —→							
	47	0005	DFI	AA	Rsvd	RD	Area	ID	Sel
octets	1	2	1	3	2	2	2	6	1

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

Figure 2: GOSIP Version 2 NSAP structure.

Currently, a proposal is being progressed in ANSI for an American National Standard (ANS) for the DSP of the NSAP address space administered by ANSI. This will provide an identical DSP structure to that provided by GOSIP Version 2. The ANSI format, therefore, differs from that illustrated above only in that the IDP is based on an ISO DCC assignment, and in that the AA will be administered by a different organization (ANSI secretariat instead of GSA). 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 may make use of slightly different NSAP formats, the principles of NSAP assignment and use are the same.

In the low-order part of the GOSIP Version 2 NSAP format, two fields are defined in addition to those required by DIS10589. 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

Internet routing components—backbones, regionals, and sites or campuses—are arranged hierarchically for the most part. A natural mapping from these components to OSI routing components is that backbones, regionals, and sites act as routing domains. (Alternatively, a site may choose to operate as an area within a regional. However, in such a case the area is part of the regional's routing domain and the discussion in Section 5.1 applies. We assume that some, if not most, sites will prefer to operate as routing domains. By operating as a routing domain, a site operates a level 2 subdomain as well as one or more level 1 areas.)

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 leaf routing domain, and,
3. at the transit routing domain (TRD).

Leaf routing domains correspond to sites, where the primary purpose is to provide intra-domain routing services. Transit routing domains are deployed to carry transit (i.e., inter-domain) traffic; backbones and regionals are TRDs.

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, regionals must manage the set of network numbers for all networks reachable through the regional. Traffic destined for other networks is generally routed to the backbone. The backbones, however, must be cognizant of the network numbers for all attached regionals and their associated networks.

In general, the advantage of abstracting routing information at a given level of the routing hierarchy is greater at the higher levels of the 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 site 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 regional (implying that the first seven octets of the address are those assigned to that regional). If considering only their own self-interest, the site itself, and the attached regional, have little reason to choose one approach or the other. The site must use one prefix or another; the source of the prefix has little effect on routing efficiency within the site. The regional must maintain information about each attached site in order to route, regardless of any commonality in the prefixes of the sites.

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

In the second case, each other regional and backbone sees a single address prefix for the regional, which encompasses the new site. This avoids the exchange of additional routing information to identify the new site's address prefix. Thus, the advantages primarily accrue to other regionals and backbones which maintain routing information about this site and regional.

One might apply a supplier/consumer model to this problem: the higher level (e.g., a backbone) is a supplier of routing services, while the lower level (e.g., an attached regional) is the consumer of these services. The price charged for services is based upon the cost of providing them. The overhead of managing a large table of addresses for routing to an attached topological entity contributes to this cost.

The Internet, however, is not a market economy. Rather, efficient operation is based on cooperation. 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 DIS10589. 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. 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 leaf routing domains would advertise is on the order of the number of attached areas; the number of prefixes a regional routing domain would advertise is approximately the number of areas attached to the client leaf routing domains; and for a backbone this would be summed across all attached regionals. Although this situation is just barely acceptable in the current Internet, as the Internet grows this will quickly become intractable. A greater degree of hierarchical information reduction is necessary to allow continued growth in the Internet.

5.2 Administration at the Leaf Routing Domain

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

As an example, assume NSF has been assigned the AA value of zzz under ICD=0005. NSF 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 attached to 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 leaf RD should take their NSAPs from the prefix assigned to the leaf RD.

5.3 Administration at the Transit Routing Domain

Two kinds of transit routing domains are considered, backbones and regionals. Each is discussed separately below.

5.3.1 Regionals

It is interesting to consider whether regional routing domains should be the common authority for assigning NSAPs from a unique prefix to the leaf routing domains that they serve. The benefits derived from data abstraction are less than in the case of leaf routing domains, and the additional degree of data abstraction provided by this is not necessary in the short term. However, 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.

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

Are leaf routing domains willing to accept prefixes derived from the regional's? In the supplier/consumer model, the regional is offering connectivity as the service, priced according to its costs of operation. This includes the "price" of obtaining service from one or more backbones. In general, backbones will want to handle as few address prefixes as possible to keep costs low. In the Internet environment, which does not operate as a typical marketplace, leaf routing domains must be sensitive to the resource constraints of the regionals and backbones. The efficiencies gained in routing clearly warrant the adoption of NSAP administration by the regionals.

The mechanics of this scenario are straightforward. Each regional is assigned a unique prefix, from which it allocates slightly longer routing domain prefixes for its attached leaf routing domains. For GOSIP NSAPs, this means that a regional would be assigned an AA identifier. Attached leaf routing domains would be assigned RD identifiers under the regional's unique prefix. For example, assume NIST is a leaf 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 regional applies for an ORG identifier from ANSI, which serves the same purpose as the AA identifier in GOSIP. The current direction in ANSI is to standardize on an NSAP structure identical to GOSIP Version 2 (see Section A.2).

5.3.2 Backbones

There does not appear to be a strong case for regionals to take their address spaces from the the NSAP space of a backbone. The benefit in routing data abstraction is relatively small. The number of regionals today is in the tens and an order of magnitude increase would not cause an undue burden on the backbones. Also, it may be expected that as time goes by there will be increased direct interconnection of the regionals, leaf routing domains directly attached to the backbones, and international links directly attached to the regionals. Under these circumstances, the distinction between regionals and backbones may become blurred.

An additional factor that discourages allocation of NSAPs from a backbone prefix is that the backbones and their attached regionals are perceived as being independent. Regionals may take their long-haul service from one or more backbones, or may switch backbones should a more cost-effective service be provided elsewhere (essentially, backbones can be thought of the same way as long-distance telephone carriers). Having NSAPs derived from the backbone 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 regional or backbone connectivity. This allows a great deal of information reduction to be achieved for those routing domains which are attached to a single TRD. In particular, such routing domains may select their NSAP addresses from a space allocated to them by the regional. This allows the regional, when announcing the addresses that it can reach to other regionals and backbones, 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 regionals and backbones. Such "multi-homed" routing domains may, for example, consist of single-site campuses and companies which are attached to multiple backbones, large organizations which are attached to different regionals at different locations in the same country, or multi-national organizations which are attached to backbones 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 regionals and backbones 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 TRD, the TRDs to which this organization is attached will need to advertise the prefix for this organization to other regionals and backbones. Other regionals and backbones (potentially worldwide) will need to maintain an explicit entry for that organization in their routing tables.

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 backbone and regional 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 backbone worldwide to maintain a separate entry for MBII (including backbones to which MBII is *not* attached). Clearly this may be acceptable if there are a small number of such multi-homed routing domains, but would place an unacceptable load on routers within backbones 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 TRD, 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 regionals in the U.S. (one east coast, and one west coast), as well as three connections to national backbones 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 TRD which has a connection to MBII needs to announce to other TRDs 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 TRDs, resulting in a smaller load on the inter-domain routing tables maintained by TRDs 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 TRD, 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 TRD 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 TRD were severed, then the NSAPs within MBII with addresses based on that TRD 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 TRDs).

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 TRD. Other TRDs to which the multi-homed organization are attached maintain a routing table entry for the organization, but are extremely selective in terms of which other TRDs are told of this route. This approach will produce a single "default" routing entry which all TRDs will know how to reach (since presumably all TRDs will maintain routes to each other), while providing more direct routing in some cases.

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, lets 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 TRDs (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 TRDs (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 worldwide TRDs 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 TRDs, 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 precisely two (or more) specific routing domains. For example, suppose that there are two regionals "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 TRDs would then advertise two prefixes to other TRDs, one prefix for leaf routing domains attached to it only, and one prefix for leaf routing domains attached to both.

This fourth solution is likely to be 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 current NSFNET regionals) 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 TRDs (including those to which the multi-homed organizations are not attached).

In addition, most of the solutions described also highlight the need for each TRD to develop policy on whether and under what conditions to accept 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 non-local NSAP prefixes will be accepted from any attached leaf RD, but not advertised to other TRDs. In a less conservative policy, a TRD might accept such non-local prefixes and agree to exchange them with a defined set of other TRDs (this set could be an a priori group of TRDs that have something in common such as geographical location, or the result of an agreement specific to the requesting leaf RD). Various policies involve real costs to TRDs, 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 TRD's prefix.

In our example, lets suppose that the XYZ corporation has a single connection to an NSFNET regional, and has therefore received an address allocation from the space administered by that regional. Similarly, let's suppose that MBII, as an international corporation with connections to six different backbones or regionals, 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 leaf routing domain which has a single connection to a "public" backbone (e.g., the NSFNET), plus a number of private point-to-point links to other leaf routing domains, can be treated as if it were single-homed to the backbone 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., Addendum 2 to ISO8348). 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 Transition Issues

Allocation of NSAP addresses based on connectivity to TRDs 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 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.

In the short term, due to 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 DIS10589. As described in Section 3, DIS10589 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 by observing the ISO9542 IS-Hello packets transmitted by routers. If the ID portion of the address is assigned by using IEEE style "stamped in PROM at birth" identifiers, then an end system can reconfigure its entire NSAP address automatically without the need for manual intervention. However, routers will still need 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 NSFNET regional. The XYZ Corporation therefore went off to the NorthSouthNet administration and got a routing domain assignment based on the AA value assigned to the NorthSouthNet regional under the GOSIP address space. However, for a variety of reasons, the XYZ Corporation decided to terminate its association with the NorthSouthNet, 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 TRD (NorthSouthNet) and the new TRD (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 TRD and the new TRD 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 TRD and the new TRD are themselves connected by other cooperative transit routing domains, 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 TRD and the new TRD are separated by a non-cooperative routing domain, or by a long path of routing domains. In this case, the old TRD 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 TRD 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 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 OSI is introduced in the Internet, it is therefore essential to choose a hierarchical structure for 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. 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 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.

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.

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 IDP format used by NIST for GOSIP Version 2, 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 TRDs, including (but not necessarily limited to):

- the NSFNET backbone;
- a number of NSFNET regional networks; and,
- a number of commercial Public Data Networks.

It is also expected that these networks will *not* be interconnected in a strictly hierarchical manner (for example, there is expected to be direct connectivity between NSFNET regionals, and all three of these types of networks may have direct international connections). However, the total number of such TRDs is expected to remain (for the foreseeable future) small enough to allow addressing of this set of TRDs via a flat address space. These TRDs 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 TRDs for communication within their own corporation.

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

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

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

Each TRD should develop policy on whether and under what conditions to accept addresses that are not based on its 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 TRD, there is not the same urgent need for hierarchical address abbreviation. 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, again 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 Non-U.S. Parts of the Internet

For the part of the Internet which is outside of the U.S., it is recommended that the DSP format be structured similarly to that specified within GOSIP Version 2 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.

6.3 Recommendations for Multi-Homed Routing Domains

Some routing domains will be attached to multiple TRDs within the same country, or to TRDs within multiple different 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 TRDs;
- administrative cost on TRD personnel; and,
- difficulty of configuration of policy-based inter-domain routing information within leaf 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 routing domains, regionals, and backbones.

7 Security Considerations

Security issues are not discussed in this memo.

8 Authors' Addresses

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

Phone: (301) 975-3627
EMail: colella@osi3.ncsl.nist.gov

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

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

Ross Callon
c/o Digital Equipment Corporation, LKG1-2/A19
550 King Street
Littleton, MA 01460-1289

Phone: (508) 486-5009
Email: Callon@bigfut.enet.dec.com

9 Acknowledgments

The authors would like to thank the members of the IETF OSI-NSAP Working Group for the helpful suggestions made during the writing of this paper.

A Administration of NSAPs

NSAPs represent the endpoints of communication through the Network Layer and must be globally unique [5]. Addendum 2 to ISO8348 defines the semantics of the NSAP and the abstract syntaxes in which the semantics of the Network address can be expressed [14].

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.

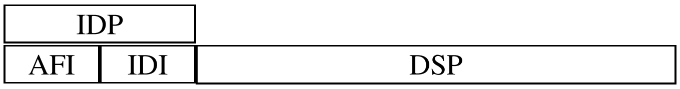

IDP Initial Domain Part
AFI Authority and Format Identifier
IDI Initial Domain Identifier
DSP Domain Specific Part

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 ISO8348/Addendum 2:

- 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 ISO3166 [9]
- ISO International Code Designator (ICD), allocated according to ISO6523 [10]
- 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. (The U.S. Government is registered as an organization by ANSI under the DCC, and in turn, will register object identifiers and X.400 names under this authority.)

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 on page 13. 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 [4] (see Figure 2, page 13). NIST has delegated the authority to register AA identifiers for GOSIP Version 2 NSAPs to the General Services Administration (GSA).

Addendum 2 to ISO8348 allows a maximum length of 20 octets for the NSAP. 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 DIS10589 [17], 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 [13], 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 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, page 13). 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.
- The network address itself contains no routing information [15]. 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 ISO8348/Addendum 2.
- 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. Currently, there is a draft proposed American National Standard (dpANS) in the U.S. for the DSP structure under DCC=840 [1]. Subsequent to an upcoming ANSI X3 Committee ballot, the dpANS will be distributed for public comment.

In the dpANS, 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 ISO3166 [9]. The numeric form of organization name is analogous to the Administrative Authority identifier in the GOSIP Version 2 NSAP.

	⟵ IDP ⟶								
	AFI	IDI	⟵ DSP ⟶						
	39	840	DFI	ORG	Rsvd	RD	Area	ID	Sel
octets	1	2	1	3	2	2	2	6	1

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

Figure 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 4976 (tel)
+1 212 398 0023 (fax)
```

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, there is a draft proposed American National Standard that specifies the DSP structure under DCC=840. The dpANS specifies a DSP structure that is identical to that specified in GOSIP Version 2.

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*. Draft Proposed American National Standard, 1991 (pending final approval by ANSI).

[2] Tim Boland. *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] ECMA. *Inter-Domain Routeing*. Technical Report 50, ISO/IEC JTC 1, Switzerland, 1989.

[4] 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.

[5] Christine Hemrick. *The OSI Network Layer Addressing Scheme, Its Implications, and Considerations for Implementation*. NTIA Report 85-186, U.S. Department of Commerce, National Telecommunications and Information Administration, 1985.

[6] ISO. *Addendum to the Network Service Definition Covering Network Layer Addressing*. RFC 941, Network Working Group, April 1985.

[7] ISO. *End System to Intermediate System Routing Exchange Protocol for use in conjunction with ISO 8473.* RFC 995, Network Working Group, April 1986.

[8] ISO. *Final Text of DIS 8473, Protocol for Providing the Connectionless-mode Network Service*. RFC 994, Network Working Group, March 1986.

[9] ISO/IEC. *Codes for the Representation of Names of Countries*. International Standard 3166, ISO/IEC JTC 1, Switzerland, 1984.

[10] ISO/IEC. *Data Interchange - Structures for the Identification of Organization*. International Standard 6523, ISO/IEC JTC 1, Switzerland, 1984.

[11] ISO/IEC. *Information Processing Systems - Open Systems Interconnection - Basic Reference Model*. International Standard 7498, ISO/IEC JTC 1, Switzerland, 1984.

[12] ISO/IEC. *Protocol for Providing the Connectionless-mode Network Service*. International Standard 8473, ISO/IEC JTC 1, Switzerland, 1986.

[13] 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.

[14] ISO/IEC. *Information Processing Systems – Data Communications – Network Service Definition Addendum 2: Network Layer Addressing*. International Standard 8348/Addendum 2, ISO/IEC JTC 1, Switzerland, 1988.

[15] ISO/IEC. *Information Processing Systems - OSI Reference Model - Part 3: Naming and Addressing*. Draft International Standard 7498-3, ISO/IEC JTC 1, Switzerland, March 1989.

[16] ISO/IEC. *Information Technology - Telecommunications and Information Exchange Between Systems - OSI Routeing Framework*. Technical Report 9575, ISO/IEC JTC 1, Switzerland, 1989.

[17] 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).* Draft International Standard 10589, ISO/IEC JTC 1, Switzerland, November 1990.

[18] K. Loughheed and Y. Rekhter. *A Border Gateway Protocol (BGP)*. RFC 1105, Network Working Group, 1989.

[19] K. Loughheed and Y. Rekhter. *A Border Router Protocol (BRP)*. Draft, Network Working Group, February 1990.

[20] ASC X3S3.3. *Intermediate System to Intermediate System Inter-Domain Routeing Exchange Protocol*. Working Document 90-216, ANSI, New York, July 1990.