

CICNet OSI ROUTING ARCHITECTURE

July 23, 1991

Linda Winkler (lwinkler@anl.gov)

Tom Easterday (tom@nisca.acs.ohio-state.edu)

This is an informational document specifying guidelines for the implementation of connectionless network protocol (CLNP) service within CICNet. This document describes a plan for coordinating OSI addresses, coordinating the supporting routing protocols, assigning routing domains within CICNet, establishing routing between the CICNet backbone and its connected sites, and coordinating routing between CICNet and other networks. It is expected that this plan will be able to coexist with the DECnet Phase IV to Phase V transition. This document specifies a starting point for the implementation of CLNP within CICNet; however, this plan may require modification during the implementation phases. This document does not attempt to set guidelines for OSI naming conventions and name registration within CICNet. A subsequent document will address naming and name registration issues.

Background

CICNet is a mid-level regional network spanning seven states in the Midwest. The CICNet backbone is formed by T1 links interconnecting twelve sites using Cisco routers. Currently CICNet implements IP and DECnet routing on the backbone and with each of its connected sites. The CICNet backbone peers with two state networks (Ohio Academic Resources Network and WISCnet) and two national backbones (NSFNET and ESnet). Additionally, some CICNet sites participate in HEPnet, a wide-area DECnet-based network.

OSI Addressing

An OSI address consists of a Network Service Access Point (NSAP) plus the selectors required to specify the destination entity within the system identified by the rest of the NSAP. An NSAP address is analogous to an Internet (IP) address. Selectors are similar to TCP port numbers in the Internet world. ISO8348/Addendum 2 defines the semantics of the NSAP [1]. The principle concept is based on hierarchical Addressing Domains. An Addressing Domain is a set of addresses administered by a common authority that specifies the structure and ensures the uniqueness of identifiers within the domain. An Addressing Domain also shares common intra-domain routing protocols. The NSAP is made up of the initial domain part (IDP) and the domain specific part (DSP). The IDP identifies the authority responsible for the structure and assignment of the remainder of the address, the DSP (see Figure 1). The IDP consists of an authority and format identifier (AFI) and an initial domain identifier (IDI). The AFI specifies the format of the IDI, the authority responsible for allocation of IDI values, and the syntax of the DSP. The recommended DSP format for use in the U.S. portion of the Internet is under the IDP format administered by the General Services Administration (GSA) for the National Institute of Standards and Technology (NIST) [2] under Section 5.1.1 of GOSIP Version 2 [3]. NIST has been designated as the authority to administer the Addressing Domain identified by IDI=0005 under AFI=47. The AFI value 47 specifies that the IDI part is interpreted as a four decimal digit International Code Designator (ICD), and the DSP has a binary abstract syntax. The Addressing Domain under IDI=0005 is available for use by all of the U.S. Federal Government. NIST will specify the structure of the DSP under IDI=0005 and delegate the task of assigning DSP values to GSA.

Within the DSP, the DSP Format Identifier (DFI) specifies the structure, semantics and administration requirements associated with the remainder of the DSP. Currently, only one DSP format (DFI=80) is defined under IDI=0005. The Administrative Authority identifies the entity responsible for the organization

of Routing Domains, establishes policies that govern the information collected and disseminated internally and externally, establishes subdomains (Areas), and delegates the corresponding responsibilities. A Routing Domain is a set of Intermediate Systems (routers) and End Systems (hosts) which operate according to the same routing procedures and is wholly contained within a single Administrative Domain. An Administrative Domain is a collection of End Systems and Intermediate Systems administered by a single authority [4].

	IDP		HO-DSP						
ISO NSAP	AFI	IDI	DSP						
US Government IDI	47	0005	DFI	AA	Res	RD	Area	System ID	NSEL
octets	1	2	1	3	2	2	2	6	1

Figure 1. GOSIP V2 NSAP Structure

- IDP - Initial Domain Part
- HO-DSP - High-Order Domain Specific Part
- AFI - Authority and Format Identifier
- IDI - Initial Domain Identifier
- DSP - Domain Specific Part
- DFI - DSP Format Identifier
- AA - Administrative Authority
- Res - reserved for future use
- RD - Routing Domain
- Area - Area Identifier
- System ID - System Identifier
- NSEL - NSAP Selector

CICNet Addressing

CICNet has been assigned the Administrative Authority (AA) value=FFF400 (hex), within the U.S. Government IDI. The CICNet backbone will be assigned one Routing Domain Identifier (RD) out of the CICNet AA. Each CICNet site (backbone site or connected site for which CICNet is the primary network provider) will be assigned one or more RDs (although requests for multiple RDs will be discouraged) from the CICNet AA (for details see the NSAP Registration section at the end of this paper). Each site has the authority to assign Area Identifiers and System Identifiers within its RD to best serve its topology and its users. CICNet backbone routers will only carry RD information about sites within the CICNet AA, making a site's local addressing (Area Identifier, System ID and NSEL) and topology transparent to the backbone. CICNet sites with connections to networks belonging to other administrative authorities, or whose primary path is a network other than CICNet, may want to obtain an address from that network's addressing authority. These cases should be decided on an individual basis, based on current and expected topologies. For further information on these special cases of multihomed sites, refer to the April 9, 1991 Internet Draft "Guidelines for OSI NSAP Allocation in the Internet" by Colella, Gardner and Callon [5].

For CICNet sites who are assigned an address by another addressing authority, the CICNet backbone

routers will carry the minimum NSAP prefix necessary to route to that site. For example, Argonne is a CICNet site which also has direct connections to the ESnet and NSFNET backbones. Argonne has been assigned an RD from the ESnet AA (see Figure 2). The CICNet routers will carry Argonne's RD prefix as well as the ESnet AA prefix in their routing tables. With this information, traffic from CICNet sites to Argonne will be sent to the Argonne CICNet router, while traffic from CICNet sites to another ESnet site, DOE-HQ for example, will be forwarded to the CICNet/ESnet peering point located at Argonne.

CICNet AA	47	0005	80	FFF400	0000	
ESnet AA	47	0005	80	005700	0000	
Argonne	47	0005	80	005700	0000	0010

Figure 2. NSAP Prefixes

CICNet CLNP Routing Strategy

CICNet backbone routers will carry routing information for RDs within the CICNet AA, other AAs within the U.S. Government IDI, and other IDPs (see Figure 3). For CICNet sites which have addresses outside the CICNet AA, the CICNet backbone will carry the shortest NSAP prefix necessary to route to the site.

In order to delimit Routing Domain boundaries, the CICNet backbone router and the site router should peer on a common firewall network. The CICNet router will be assigned one address out of the site's RD and one address out of the CICNet backbone RD. Eventually, the backbone router and the site router will run the Interdomain Routing Protocol (IDRP) to exchange routing information; however, until the IDRP specification [6] is complete and vendors implement it, CICNet will run Cisco's ISO Interior Gateway Routing Protocol (IGRP) [7] to exchange CLNP routing information with a CICNet site if the site's firewall router happens to be a Cisco. If it is not, static CLNP routing will be necessary until IDRP is widely available from various router vendors.

Routing between Administrative Domains will occur where CICNet peers with the national backbones (and potentially other regionals in the future). At the Administrative Domain boundary, it is important that hierarchical abbreviation of routing information occur (route aggregation) and that a minimum of information be exchanged (address prefix information). Herein lies a problem. CICNet currently peers with NSFNET at two (soon to be three) sites. If CICNet has only one AA and NSFNET only wants to hear AA information from CICNet, then NSFNET will not be able to distinguish which interconnect point to use to send traffic to a CICNet site. To take care of this problem, CICNet will utilize two bits of the RD field to identify peering relationships with NSFNET. CICNet will assign sites a RD based on the nearest (according to the topology) NSFNET peering point (see Figure 4). CICNet will announce unique NSAP prefixes to NSFNET at each of its peering points. In the case that one peer fails, NSFNET will accept CICNet NSAP prefixes from any one of the three peers. Such peering issues remain to be worked out among the national backbones. IDRP promises to solve many of the problems but, until it can be implemented, static routing based on unique NSAP prefixes will be used.

Although the objective of this document is to deal with addressing and routing issues for CLNP within CICNet, there are concerns about the transition from DECnet Phase IV to Phase V. The CICNet backbone currently utilizes Cisco routers which implement CLNP and DECnet Phase IV in parallel protocol stacks. This means that DECnet Phase IV and CLNP can coexist with completely separate routing tables. The

CICNet	47	0005	80	FFF400	0000	
Argonne	47	0005	80	005700	0000	0010
ESnet	47	0005	80	005700	0000	
NSFNET	47	0005	80	FFFF00	0000	
NSN	47	0005	80	005900	0000	
CERN	47	0020	80	005900	0000	
NORDUNET	47	0023				

Figure 3. NSAP Prefixes

<u>RDs</u>	<u>Primary NSFNET peering point</u>
0000 - 3FFF	University of Illinois
4000 - 7FFF	University of Michigan
8000 - BFFF	Argonne
C000 - FFFF	Reserved

Figure 4. CICNet Routing Domains

problems begin when you try to ensure interoperability between DECnet Phase IV nodes and DECnet Phase V nodes. Although the Digital Equipment Corporation and Cisco Systems, Inc. have each provided a solution to the Phase IV/Phase V interoperability problem, neither solution is suitable in the wide-area DECnet community. The existing HEPnet wide-area DECnet network routing scheme relies heavily on hidden addresses (nodes and areas) and encompasses geographically dispersed sites which share a DECnet Phase IV area. These "features" break the GOSIP guidelines for NSAP allocations. Experts within the ESnet and HEPnet communities are working with Digital and Cisco on solutions to these problems. This document will be revised as solutions are identified.

CICNet NSAP Registration

To apply for a Routing Domain within the CICNet AA, retrieve the ISO Routing Domain Assignment form via anonymous ftp from nic.cic.net. The completed form should be returned to nsap-request@cic.net.

References

- [1] 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.
- [2] R. Colella, R. Callon. OSI NSAP Address Format For Use In The Internet. Internet-Draft, Network Working Group, February 12, 1991.
- [3] GOSIP Advanced Requirements Group. Government Open Systems Interconnect Profile (GOSIP) Version 2. Federal Information Processing Standard, U.S. Department of Commerce, National Institute of Standards and Technology, Gaithersburg, MD.
- [4] ISO/IEC, Information technology - Telecommunications and Information exchange between systems - OSI Routeing Framework, ISO/IEC TR 9575, July, 1989.
- [5] Richard Colella, Ella Gardner, Ross Callon. Guidelines for OSI NSAP Allocation in the Internet. Internet-Draft, Network Working Group, April 9, 1991.
- [6] ISO/IEC, Information Processing Systems - Telecommunications and Information Exchange between Systems - Intermediate System to Intermediate System Inter-domain Routeing Information Exchange Protocol, ISO/IEC JTC1/SC6, November 1990.
- [7] Cisco Systems, Inc. Gateway System Manual, November, 1990.