

Network Working Group
Request for Comments: 1070

R. Hagens
U of Wisconsin - Madison
N. Hall
U of Wisconsin - Madison
M. Rose
The Wollongong Group
February 1989

Use of the Internet as a Subnetwork for
Experimentation with the OSI Network Layer

Status of this Memo

This RFC proposes a scenario for experimentation with the International Organization for Standardization (ISO) Open Systems Interconnection (OSI) network layer protocols over the Internet and requests discussion and suggestions for improvements to this scenario. This RFC also proposes the creation of an experimental OSI internet. To participate in the experimental OSI internet, a system must abide by the agreements set forth in this RFC. Distribution of this memo is unlimited.

WARNING

The methods proposed in this RFC are suitable ONLY for experimental use on a limited scale. These methods are not suitable for use in an operational environment.

Introduction

Since the International Organization for Standardization (ISO) Open Systems Interconnection (OSI) network layer protocols are in their infancy, both interest in their development and concern for their potential impact on internetworking are widespread. This interest has grown substantially with the introduction of the US Government OSI Profile (GOSIP), which mandates, for the US Government, the use of OSI products in the near future. The OSI network layer protocols have not yet received significant experimentation and testing. The status of the protocols in the OSI network layer varies from ISO International Standard to "contribution" (not yet a Draft Proposal). We believe that thorough testing of the protocols and implementations of the protocols should take place concurrently with the progression of the protocols to ISO standards. For this reason, the creation of an environment for experimentation with these protocols is timely.

Thorough testing of network and transport layer protocols for

internetworking requires a large, varied, and complex environment. While an implementor of the OSI protocols may of course test an implementation locally, few implementors have the resources to create a sufficiently large dynamic topology in which to test the protocols and implementations well.

One way to create such an environment is to implement the OSI network layer protocols in the existing routers in an existing internetwork. This solution is likely to be disruptive due to the immature state of the OSI network layer protocols and implementations, coupled with the fact that a large set of routers would have to implement the OSI network layer in order to do realistic testing.

This memo suggests a scenario that will make it easy for implementors to test with other implementors, exploiting the existing connectivity of the Internet without disturbing existing gateways.

The method suggested is to treat the Internet as a subnetwork, hereinafter called the "IP subnet." We do this by encapsulating OSI connectionless network layer protocol (ISO 8473) packets in IP datagrams, where IP refers to the Internet network layer protocol, RFC 791. This encapsulation occurs only with packets travelling over the IP subnet to sites not reachable over a local area network. The intent is for implementations to use OSI network layer protocols directly over links locally, and to use the IP subnet as a link only when necessary to reach a site that is separated from the source by an IP gateway. While it is true that almost any system at a participating site may be reachable with IP, it is expected that experimenters will configure their systems so that a subset of their systems will consider themselves to be directly connected to the IP subnet for the purpose of testing the OSI network layer protocols or their implementations. The proposed scheme permits systems to change their topological relationship to the IP subnet at any time, also to change their behavior as an end system (ES), intermediate system (IS), or both at any time. This flexibility is necessary to test the dynamic adaptive properties of the routing exchange protocols.

A variant of this scheme is proposed for implementors who do not have direct access to the IP layer in their systems. This variation uses the User Datagram Protocol over IP (UDP/IP) as the subnetwork.

In this memo we will call the experiment based on the IP subnet EON, an acronym for "Experimental OSI-based Network". We will call the experiment based on the UDP/IP subnet EON-UDP.

It is assumed that the reader is familiar with the OSI connectionless network layer and, in particular, with the following documents:

RFC 768

User Datagram Protocol.

RFC 791

Internet Protocol.

ISO 8473

Protocol for Providing the Connectionless mode Network Service.

ISO DP 9542

End System to Intermediate System Routing Exchange Protocol for Use in Conjunction with the Protocol for the Provision of the Connectionless-mode Network Service (ISO 8473).

ISO TC 97/SC 6/N xxxx

Intermediate System to Intermediate System Intra-Domain Routing Exchange Protocol.

PD TR 97/SC 6/N 9575

OSI Routing Framework.

Definitions

EON

An acronym for Experimental OSI Network, a name for the proposed experimental OSI-based internetwork that uses the IP over the Internet as a subnetwork.

EON-UDP

A name for the proposed experimental OSI-based internetwork that uses the UDP/IP over the Internet as a subnetwork.

ES

End system as defined by OSI: an OSI network layer entity that provides the OSI network layer service to a transport layer.

IANA

The Internet Assigned Numbers Authority. Contact Joyce K. Reynolds (JKREY@ISI.EDU).

IS

An OSI network layer entity that provides the routing and forwarding functions of the OSI connectionless network layer.

OSI CLNL

OSI connectionless network layer.

NSDU

Network Service Data Unit.

PDU

Protocol Data Unit, or packet.

NPDU

Network Protocol Data Unit.

ISO-gram

An NPDU for any protocol in the OSI CLNL, including ISO 8473 (CLNP), ISO DP 9542 (ES-IS), and ISO TC 97/SC 6/N xxxx (IS-IS).

Participating system

An ES or IS that is running a subset of the OSI CLNL protocols and is reachable through the application of these protocols and the agreements set forth in this memo.

Core system

An ES or IS that considers itself directly connected to the IP subnet for the purpose of participating in EON.

NSAP-address

Network Service Access Point address, or an address at which the OSI network services are available to a transport entity.

SNPA-address

SubNetwork Point of Attachment address, or an address at which the subnetwork service is available to the network entity.

Issues to be Addressed by this Memo

In order to make the experimental OSI internet work, participating experimenters must agree upon:

- how ISO-grams will be encapsulated in IP or UDP packets,
- the format of NSAP-addresses to be used,
- how NSAP-addresses will be mapped to SNPA-addresses on the IP subnet,
- how multicasting, which is assumed by some OSI CLNL protocols, will be satisfied, and
- how topology information and host names will be disseminated.

This memo contains proposals for each of these issues.

Design Considerations

The goals of this memo are:

- to facilitate the testing of the OSI network layer protocols among different implementations,
- to do this as soon as possible, exploiting existing connectivity,
- to do this without requiring any changes to existing IP gateways,
- to create a logical topology that can be changed easily, for the purpose of testing the dynamic adaptive properties of the protocols, and
- to minimize the administrative requirements of this experimental internetwork.

The following are not goals of this memo:

- to permit the use of arbitrary ISO-style NSAP-addresses,
- to require that participants have working implementations of all of the OSI routing protocols before they can participate in any capacity,
- to permit or encourage the use of existing IP routing methods and algorithms for the routing of ISO-grams among participating ESSs and ISSs,
- to create a production-like environment accommodating a very large number of systems (ESSs, ISSs or both), and
- to provide or to encourage IP-to-CLNP gatewaying.

Encapsulating ISO-grams in IP datagrams

The entire OSI network layer PDU, whether it be an ISO 8473 PDU, an ISO DP 9542 PDU, or an IS-IS PDU, will be placed in the data portion of an IP datagram at the source. The ISO 8473 entity may fragment an NSDU into several NPDUs, in which case each NPDU will be encapsulated in an IP datagram. The intent is for the OSI CLNL to fragment rather than to have IP fragment, for the purpose of testing the OSI CLNL. Of course, there is no guarantee that fragmentation will not occur within the IP subnet, so reassembly must be supported at the IP level in the destination participating system.

SNPA-addresses (Internet addresses) will be algorithmically derived from the NSAP-addresses as described below. The "protocol" field of the IP datagram will take the value 80 (decimal), which has been assigned for this purpose.

NSAP-Address Format

The OSI internetwork described here will form one routing domain, with one form of NSAP address recognized by all level 1 routers in this domain. Other address formats may be agreed upon in later editions of this memo.

The address format to be used in this experiment is that specified in RFC 1069. According to RFC 1069, the low-order portion of the Domain Specific Part of the NSAP address may vary depending on the conventions of the particular routing domain. For the purposes of this experiment, we shall use the following address format:

Address Format for EON		
Octet	Value	Meaning
1	47	Authority and Format Identifier
2,3	00, 06	International Code Designator
4	3	Version Number
5,6	0	Global Area Number, see RFC 1069
7,8	RDN	Routing Domain Number, assigned by IANA
9-11	0	Pad
12,13	0	LOC-AREA, see below
14,15	0	unused
16-19	A.B.C.D	Internet address
20		NSAP Selector, assigned IANA

Note: It is our desire that the address format used by EON be consistent with RFC 1069. To that end, the address format proposed by this RFC may change as future editions of RFC 1069 become available.

The mapping between NSAP-addresses and SNPA-addresses (Internet addresses) on the proposed IP subnet is straightforward. The SNPA-address is embedded in the NSAP-address.

There are several ways in which the LOC-AREA field could be used.

- (1) Assign local areas, administered by the Internet Assigned Numbers Authority (IANA). The advantage of this is that it permits experimentation with area routing. The disadvantage is that it will require an additional directory service to map host names to NSAP-addresses. We would like to use the existing domain name servers to derive Internet addresses from names, and we would like NSAP-addresses to be derivable from the Internet addresses alone.
- (2) Have one local area in the EON, with LOC-AREA value 0. This would eliminate the problem of name-to-NSAP-address binding, but would not permit experimentation with area routing. It would not, however preclude the use of areas later, for example, when OSI directory services are widely available.
- (3) Make the local area a simple function of the Internet address. The advantage of this is that it would permit experimentation with area addressing without requiring additional directory services, but the areas derived would not be under the control of the experimenters and may not correspond to anything useful or meaningful for the purposes of this experiment.

We believe that initially, the preferred alternative is to use only

zero-valued local areas. Later editions of this memo may contain proposals for use of the local area field, when the IS-IS proposal is more mature and perhaps when OSI directory services are in use among experimenters.

The value of the high-order portion of the DSP will be set in accordance with RFC 1069.

Other NSAP-Address Formats

PDUs carrying NSAP-addresses of other formats can be routed through this domain. This is the job of the level 2 routers, described in the IS-IS document.

Multicast Addresses on the IP Subnet

The ES-IS and IS-IS routing exchange protocols assume that broadcast subnetworks support two multicast addresses: one for all ESs and the other for all ISs. While one could obviate this issue by treating the IP subnet as a general topology subnetwork or as a set of point-to-point links, it is also desirable to treat the IP subnet as a broadcast subnetwork for the purpose of testing those parts of an implementation that run over broadcast subnets. A participating implementor not having access to several local machines running the OSI CLNL may test the protocols over the IP subnet as if the IP subnet were a broadcast subnet.

The multicasting assumed by the OSI CLNL can be simulated by a small sublayer lying between the OSI CLNL and the IP subnet layer. For the purpose of this discussion, call this sublayer an SNACP, a SubNetwork Access Protocol, in OSI argot. In each system the SNACP caches a table of the Internet addresses of systems that it considers to be reachable in one ISO 8473-hop over the IP subnet. These are called "core" systems. In this sense, the use of the cache simulates a set of links over which a system will send ISO configuration messages (ES Hello, IS Hello, etc.) when it comes up. As a local matter, the table of core systems may or may not expand during the system's lifetime, in response to configuration messages from other core systems.

On the outgoing path, the SNACP accepts an ISO-gram and a parameter indicating the intended use of this ISO-gram: send to a single system, to all ESs, to all ISs, or to all systems. If the intended destination is a single system, the ISO-gram is sent only to its destination. Otherwise, the SNACP makes a copy of the ISO-gram for each of the SNPA-addresses in the cache, effecting a broadcast to all participating systems. Before passing an ISO-gram to the IP subnet layer, the SNACP prepends an SNACP header to each outgoing ISO-gram.

This header will take the form:

SNACP Header Format		
Octet	Value	Meaning
1	01	Version number
2		Semantics of address:
	00	Not a multicast address
	01	All ESs
	02	All ISs
	03	Broadcast
3,4		OSI checksum as defined in ISO 8473

The SNACP header has three fields, a version number field, a semantics field, and a checksum field. The version number will take the value 01. The checksum field will take the two octet ISO (Fletcher) checksum of the SNACP header. The checksum algorithm is described in ISO 8473.

The semantics field will take one of 4 values, indicating "all ESs", "all ISs", "broadcast", or "not a multicast address". The value of the semantics field is determined by a parameter passed to the SNACP by the calling OSI network entity. A participant in the experiment may test the OSI network layer over a set of point-to-point links by choosing not to use the multicast capabilities provided by the SNACP on the outgoing path.

On the incoming path, the SNACP inspects the SNACP header and decides whether or not to accept the ISO-gram. If it accepts the ISO-gram, the SNACP removes the SNACP header and passes the ISO-gram to the OSI CLNL, otherwise, it discards the ISO-gram. The SNACP will always accept ISO-grams with SNACP headers indicating "not a multicast address" (value zero in the semantics field) and "broadcast" (value 03). Whether an SNACP will accept ISO-grams for either of the two multicast groups "all ESs" (value 1) and "all ISs" (value 2) will depend on local configuration information. If the system on which the SNACP resides is configured as an end system, it will accept ISO-grams destined for "all ESs" and if it is configured as an intermediate system, it will accept ISO-grams destined for "all ISs".

A participant who is testing the OSI network layer over a set of point-to-point links will accept ISO-grams according to these rules as well.

Consideration was given to making the SNACP extensible by making the semantics and checksum fields variable-length parameters, in the

manner of ISO 8473. We feel that the presence of a version number provides sufficient extensibility.

Errors on the IP subnet

The IP subnet layer may receive ICMP messages and may pass error indications to the SNACp layer as a result of having received these ICMP messages. It is assumed that in this context, the IP subnet will handle ICMP messages in the same way that it handles them in any other context. For example, upon receipt of an ICMP echo message, the IP subnet will respond with an ICMP echo reply, and the SNACp need not be informed of the receipt of the ICMP echo message. Certain ICMP messages such as source quench are likely to produce an error indication to all layers using the IP subnet. The actions taken by the SNACp for these indications is purely a local matter, however the following actions are suggested.

Suggested SNACp Actions in Response to ICMP-related Error Indications	
ICMP message type	Action taken by the SNACp
Destination unreachable, Parameter problem, Time exceeded	If the remote address is present in the cache of core systems' addresses, mark it unusable. Inform network management.
Source quench	If the remote address is present in the cache of core systems' addresses, mark the remote address as unusable and set a timer for a time after which the address becomes usable again. Inform network management.
All others	Ignored by the SNACp layer.

To "inform network management" may mean to print a message on the system console, to inform a local process, to increment a counter, to write a message in a log file, or it may mean to do nothing.

The effect of marking a cached address as unusable is as follows. When the SNACp attempts to broadcast or multicast an ISO-gram, addresses in the cache that are marked as unusable are ignored. When the SNACp attempts to send a non-multicast ISO-gram to an unusable cached address, the SNACp returns an error indication to the OSI CLNL. In this way, when the OSI CLNL uses the SNACp to simulate a

set of point-to-point links, it is notified when a link fails, but when the OSI CLNL uses the SNACp to simulate a multicast subnet, it is not notified when one system on the subnet goes down.

Use of UDP/IP in Lieu of IP

In addition to using IP directly, for testing purposes it may be useful to support the OSI CLNL over the User Datagram Protocol (UDP). This is because some implementors do not have direct access to IP, but do have access to the UDP. For example, an implementor may have an a binary license for an operating system that provides TCP/IP and UDP/IP, but no direct access to IP. These implementors may participate in a parallel experiment, called EON-UDP, by using UDP/IP as a subnetwork instead of using the IP subnet. In this case, the OSI NPDU (after fragmentation, if applicable) will be placed in the data portion of a UDP datagram. UDP port 147 (decimal) has been assigned for this purpose. These participants will place an SNACp between UDP and ISO 8473 rather than between IP and ISO 8473. In all other respects, the EON-UDP experiment is identical to the EON experiment.

Of course, network layers entities using the UDP/IP subnet will not interoperate directly with network layers entities using the IP subnet. The procedures proposed in this memo do not prevent an implementor from building an EON to EON-UDP gateway, however the issues related to building and routing to such a gateway are not addressed in this memo. This memo simply proposes two distinct parallel experiments for two groups of experimenters having different resources.

The preferred method of experimentation is to use the IP subnet, in other words, EON. The EON-UDP variant is intended for use only by those who cannot participate in EON.

Dissemination of Topological Information and Host Names

The EON experiment simulates a logical topology that is not as connected as the underlying logical topology offered by the Internet. The topology of the IP subnet is, in effect, simulated by the SNACp layer in each of the core systems. Each of the core systems caches a list of the other core systems in the EON. When a system boots, it needs some initial list of the participating core systems. For this reason, a list of core systems will be maintained by the IANA.

In addition, a list of all participating ESs will be maintained by the IANA. This list is not necessary for the functioning of the EON network layer. It is a convenience to the experimenters, and is meant for use by application layer software or human users.

Two pairs of lists are kept, one for the EON and one for EON-UDP.

core.EON This is a list of SNPA-addresses of those systems that will be (logically) reachable via the IP subnet in one ISO 8473-hop from any other core system. This does not mean that systems in this file are gateways or ISSs. They may be ESSs, ISSs or both. A site may participate as a core system before its address is included in this file and distributed to other core systems, but under these circumstances other core systems will not know to send configuration messages (ESHs and ISHs) to the new site when coming up or rebooting. The SNPA-addresses in this file will be ASCII strings of the form A.B.C.D, no more than one per line. White space (tabs, blanks) will be optional before A and after D. A pound-sign (#) will indicate that it and everything following it on that line is a comment. For example:

```
128.105.2.153 # bounty.cs.wisc.edu
```

core.EON-UDP

This is the equivalent of core.EON for use with the UDP/IP subnet. The format is the same that of core.EON.

hosts.EON This is a list of the ASCII host names of all end systems participating in the IP subnet experiment, one host name per line. It is not used by the OSI CLNL.

hosts.EON-UDP

This is a list of the ASCII host names of all end systems participating in the UDP/IP subnet experiment, one host name per line. It is meant for the use of applications. It is not used by the OSI CLNL.

The files will be available from the IANA via anonymous ftp. Sites wishing to join the experimental OSI internet will have to have their host names and core system addresses added to the appropriate files. They may do so by sending requests to Joyce K. Reynolds at the electronic mail address:

```
JKREY@ISI.EDU
```

Hypothetical EON Topology

Figure 1 describes the logical links in a hypothetical topology, in which three university computer sciences departments are participating in the experiment: the University of Wisconsin (U of W), the University of Tudor (U of Tudor), and the University of Fordor (U of Fordor). The U of W has two local area networks (LANs), 128.105.4 and 128.105.2, and four systems that are acting as ESs in the experiment. Two systems are attached to both LANs. Only one of these two systems is forwarding ISO-grams, in other words, acting as an IS. The U of Tudor has only one participating system, and it is acting as an ES. The U of Fordor has two systems that are participating in the experiment, one of which is an IS only, and the other of which is acting as an ES only.

The contents of the core.EON and hosts.EON files for this topology are shown below.

```
#
# core.EON for hypothetical EON topology
#
128.105.2.153    # IS/ES in cs.wisc.edu
26.5.0.73      # ES in cs.tudor.edu
192.5.2.1      # IS in cs.fordor.edu

#
# hosts.EON hypothetical EON topology
#
128.105.4.150  # ES in cs.wisc.edu
128.105.2.150  # same as above : multihomed ES
128.105.4.154  # ES in cs.wisc.edu
128.105.4.151  # ES in cs.wisc.edu
128.105.2.153  # IS/ES in cs.wisc.edu
26.5.0.73      # ES in cs.tudor.edu
192.5.2.2      # ES in cs.fordor.edu
```


core.EON file is out-of-date for a while. The effect will be that other core systems will continue to send configuration messages to 192.5.2.1, which will respond as an ES, not as an IS, and it will appear that 192.5.2.2 is not reachable from the rest of the topology because the other core systems will not know to send configuration information to it. However, when 192.5.2.2 is booted, it will send configuration messages to all core systems informing them of its existence via the IS-IS protocol. Those core systems that are acting as ISS will respond with their configuration messages, update their core system caches, thereby establishing a set of logical links between 192.5.2.2 and the rest of the core systems.

Relationship of this Memo to other RFCs

RFCs 1006 and 983

ISO Transport Services on top of the TCP. Whereas RFCs 1006 and 983 offer a means of running the OSI session layer protocol and higher OSI layers over TCP/IP, this memo provides a means of running the OSI network and transport layers on an IP internetwork.

RFC 1069

Guidelines for the use of Internet-IP addresses in the ISO Connectionless-Mode Network Protocol. RFC 1069 suggests a method to use the existing Internet routing and addressing in a gateway that forwards ISO connectionless network layer protocol datagrams. In contrast, this memo suggests a method to use the ISO routing and addressing in a gateway that forwards ISO connectionless network layer protocol datagrams.

RFC 982

ANSI Working Document X3S3.3/85-258. This is a set of guidelines for specifying the structure of the DSP part of an ISO address. The addresses described in this memo meet the guidelines set forth in RFC 982.

References

Plummer, D., "An Ethernet Address Resolution Protocol - or - Converting Network Protocol Addresses to 48.bit Ethernet Address for Transmission on Ethernet Hardware", RFC 826, MIT, November 1982.

Finlayson, R., T. Mann, J. Mogul, and M. Theimer, "A Reverse Address Resolution Protocol", RFC 903, Stanford, June 1984.

Postel, J., "Internet Protocol - DARPA Internet Program Protocol Specification", RFC 791, DARPA, September 1981.

Postel, J., "Internet Control Message Protocol - DARPA Internet Program Protocol Specification", RFC 792, ISI, September 1981.

Postel, J., "User Datagram Protocol", RFC 768, ISI, August 1980.

ISO, "Protocol For Providing the Connectionless Mode Network Service", (ISO 8473), March 1986. (This is also published as RFC 994.)

ISO, "End System to Intermediate System Routing Exchange Protocol for Use in Conjunction with the Protocol for the Provision of the Connectionless-mode Network Service (ISO 8473)", (ISO DP 9542). (This is also published as RFC 995.)

ISO, "Intermediate System to Intermediate System Intra-Domain Routing Exchange Protocol", (ISO TC 97/SC 6/N xxxx).

OSI, "OSI Routing Framework", (PD TR 97/SC 6/N 9575).

Authors' Addresses

Robert A. Hagens
Computer Sciences Department
University of Wisconsin - Madison
1210 West Dayton Street
Madison, WI 53706
608/ 262-1017

EMail: hagens@cs.wisc.edu

Nancy E. Hall
Computer Sciences Department
University of Wisconsin - Madison
1210 West Dayton Street
Madison, WI 53706
608/ 262-5945

EMail: nhall@cs.wisc.edu

Marshall T. Rose
The Wollongong Group
San Antonio Blvd.
Palo Alto, California
415/ 962-7100

Email: mrose@twg.com

Comments and Suggestions

Please direct comments, suggestions, and indications of desire to participate to the authors.