

TECHNICAL REPORT

ISO/IEC
TR 9575

Second edition
1995-10-15

Information technology — Telecommunications and information exchange between systems — OSI Routeing Framework

*Technologies de l'information — Communication de données et échange
d'information entre systèmes — Cadre général de routage OSI*



Reference number
ISO/IEC TR 9575:1995(E)

Contents	Page
Foreword.....	iii
Introduction	iv
1 Scope.....	1
2 References.....	1
3 Definitions	1
4 Symbols and Abbreviations.....	2
5 Routeing Concepts	2
6 Environment for OSI Routeing	6
7 Goals for OSI Routeing	6
8 Structure of Global OSI Routeing.....	8

Foreword

ISO (the International Organization for Standardization) and IEC (the International Electrotechnical Commission) form the specialized system for worldwide standardization. National bodies that are members of ISO or IEC participate in the development of International Standards through technical committees established by the respective organization to deal with particular fields of technical activity. ISO and IEC technical committees collaborate in fields of mutual interest. Other international organizations, governmental and non-governmental, in liaison with ISO and IEC, also take part in the work.

In the field of information technology, ISO and IEC have established a joint technical committee, ISO/IEC JTC 1.

The main task of technical committees is to prepare International Standards, but in exceptional circumstances a technical committee may propose the publication of a Technical Report of one of the following types:

- type 1, when the required support cannot be obtained for the publication of an International Standard, despite repeated efforts;
- type 2, when the subject is still under technical development or where for any other reason there is the future but not immediate possibility of an agreement on an International Standard;
- type 3, when a technical committee has collected data of a different kind from that which is normally published as an International Standard ("state of the art", for example).

Technical Reports of types 1 and 2 are subject to review within three years of publication, to decide whether they can be transformed into International Standards. Technical Reports of type 3 do not necessarily have to be reviewed until the data they provide are considered to be no longer valid or useful.

ISO/IEC TR 9575, which is a Technical Report of type 3, was prepared by Joint Technical Committee ISO/IEC JTC 1, *Information technology*, Subcommittee SC 6, *Telecommunications and information exchange between systems*.

This second edition cancels and replaces the first edition (ISO/IEC TR 9575:1990), which has been technically revised.

Introduction

In the OSI environment (OSIE), the possibility exists for any End System (ES) to communicate with any other ES. The physical path (or paths) over which this communication takes place may

- include multiple Intermediate Systems (IS);
- include multiple subnetwork types; and
- traverse multiple, independent organisations.

Furthermore, one instance of communications may follow a different path from another instance of communications.

Within the Network Layer, the Internal Organisation of the Network Layer (ISO 8648) identifies two functions, Routeing and Relaying, as being central to the ability for End Systems to communicate through an arbitrary concatenation of subnetworks and Intermediate Systems.

Part of the overall function of routeing and relaying is to allow ESs and ISs to find an appropriate path between two or more ESs for a given instance of communications.

Relaying is concerned primarily with the actual transformation and manipulation of Network Protocol Data Units (NPDUs) as they transit Intermediate Systems. Routeing, on the other hand, is primarily concerned with the maintenance and selection of paths through multiple subnetworks and Intermediate Systems which allow NPDUs to flow smoothly between End Systems.

There are four important aspects to routeing, i.e.:

- a) the information required by ESs and ISs (5.1.1),
- b) the techniques used by ESs and ISs to collect that information (5.1.2),
- c) the techniques used by ESs and ISs to distribute that information (5.1.3), and
- d) the functions executed by ESs and ISs on that information to determine the paths over which NPDUs flow between pairs of NSAPs (5.1.4).

This Technical Report discusses these aspects of routeing, and describes how various protocols may be employed to effect the OSI routeing functions. It does not discuss relaying, except where relaying functions are closely allied with routeing functions.

This second edition of ISO/IEC TR 9575 adds the option of interconnecting Routeing Domains using multicast subnetworks, interconnecting Administrative Domains using multicast subnetworks, and providing connectivity within Routeing Domains using multicast subnetworks.

Information technology — Telecommunications and information exchange between systems — OSI Routeing Framework

1 Scope

This Technical Report provides a framework in which OSI protocols for routeing may be developed and to expedite the progression of routeing protocols through the standardisation process. At the time of publication, this report reflected the current state of OSI Routeing, and does not preclude future extensions and developments.

2 References

The following International Standards | ITU-T Recommendations contain provisions which, through reference in this text, constitute provisions of this Technical Report. At the time of publication, the editions indicated were valid. All Standards | Recommendations are subject to revision, and parties to agreements based on this Technical Report are encouraged to investigate the possibility of applying the most recent editions of the Standards | Recommendations listed below. Members of IEC and ISO maintain registers of currently valid International Standards. The Telecommunication Standardization Bureau of the ITU maintains a list of currently valid ITU-T Recommendations.

2.1 Identical International Standards | Recommendations

- ITU-T Recommendation X.200 (1994) | ISO/IEC 7498-1:1994, *Information technology - Open Systems Interconnection - Basic Reference Model: The Basic Model*.
- CCITT Recommendation X.213 (1992) | ISO/IEC 8348:1993, *Information technology - Open Systems Interconnection - Network service definition*.
- ITU-T Recommendation X.233 (1993) | ISO/IEC 8473-1:1994, *Information technology - Protocol for providing the connectionless-mode network service: Protocol specification*.

2.2 Paired International Standards | Recommendations

- ITU-T Recommendation X.223 (1993), *Use of X.25 to Provide the OSI Connection-mode Network Service for ITU-T Applications*.
- ISO/IEC 8878:1992, *Information technology - Telecommunications and information exchange between systems - Use of X.25 to provide the OSI Connection-mode Network Service*.

2.3 Additional references

- ISO 8648:1988, *Information processing systems - Open Systems Interconnection - Internal organization of the Network Layer*.
- ISO 9542:1988¹, *Information processing systems - Telecommunications and information exchange between systems - End system to intermediate system routeing exchange protocol for use in conjunction with the protocol for providing the connectionless mode network service*.
- ISO/IEC 10030:1995, *Information technology - Telecommunications and information exchange between systems - End System Routeing Information Exchange Protocol for use in conjunction with ISO/IEC 8878*.
- ISO/IEC 10589:1992, *Information technology - Telecommunications and information exchange between systems - Intermediate system to intermediate system intra-domain-routeing routine information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode Network Service (ISO 8473)*.
- ISO/IEC 10747:1994, *Information technology - Telecommunications and information exchange between systems - Protocol for exchange of inter-domain routeing information among intermediate systems to support forwarding of ISO 8473 PDUs*.
- RFC 1629, *Guidelines for OSI NSAP allocation in the Internet*.

¹ Currently under revision.

3 Definitions

3.1 Reference Model Definitions

This Technical Report makes use of the following terms defined in ITU-T Rec. X.200 | ISO 7498-1:

- a) Network Layer
- b) Network Service Access Point
- c) Network Service Access Point address
- d) Network entity
- e) Routeing

- f) Network protocol
- g) Network relay
- h) Network protocol data unit
- i) System management
- j) Layer management

3.2 Network Layer Architecture Definitions

This Technical Report makes use of the following terms defined in ISO 8648:

- a) Subnetwork
- b) End system
- c) Intermediate system
- d) Subnetwork service

3.3 Network Layer Addressing Definitions

This Technical Report makes use of the following terms defined in CCITT Rec. X.213 | ISO/IEC 8348.

- a) Subnetwork address
- b) Subnetwork Point of Attachment

3.4 Routing Framework Definitions

For the purpose of this Technical Report the following definitions apply.

3.4.1 Administrative Domain: A collection of End systems, Intermediate systems, and subnetworks operated by a single organisation or administrative authority.

The components which make up the domain are assumed to interoperate with a significant degree of mutual trust among themselves, but interoperate with other Administrative Domains in a mutually suspicious manner.

NOTE: The term *Administrative Domain* is not intended to have any particular relationship to an *Administration* as defined by the ITU-T. A ITU Administration may in fact operate an Administrative Domain, but this would be no different from an Administrative Domain operated by any organisation from the point of view of this Routing Framework.

3.4.2 Routing Domain: A set of End Systems and Intermediate Systems which operate according to the same routing procedures and which is wholly contained within a single Administrative Domain.

See 8.1.2.1 for a precise formal definition of this concept.

3.4.3 common domain: An Administrative Domain which is not a member of a higher level domain.

A common domain is the highest level in the routing hierarchy. There is no single domain above the common domain. In this sense, the routing hierarchy is in fact

multiple hierarchies, with the common domain as the highest element of each hierarchy.

Where there are multiple common domains, they co-operate as peers to make it possible to route to any NSAP in the OSIE.

3.4.4 hop: The traversal of a single subnetwork by a PDU.

3.4.5 black hole: A situation in which an Intermediate System, due to a breakdown of the routing procedures, malicious intent, or lack of information, discards or otherwise refuses to forward all traffic directed to it.

A black hole may also be formed on a connectionless subnetwork when the intended recipient of traffic is unavailable.

3.4.6 Subnetwork Address Resolution Entity: A network layer entity available on a subnetwork which acts as a repository for, and source of, routing information for that subnetwork.

3.4.7 multicast subnetwork: a subnetwork in which a single data unit transmitted by a source is received by multiple destinations.

3.4.8 multicast communication: the use of a multicast subnetwork for data transmission.

4 Symbols and Abbreviations

BIS	Border Intermediate System
ES	End System
IS	Intermediate System
LAN	Local Area Network
NPDU	Network Protocol Data Unit
NSAP	Network Service Access Point
OSIE	Open System Interconnection Environment
PDU	Protocol Data Unit
QoS	Quality of Service
SN	Subnetwork
SNARE	Subnetwork Address Resolution Entity
SNPA	Subnetwork Point of Attachment
WAN	Wide Area Network

5 Routing Concepts

5.1 Functional Decomposition of Routing

OSI Routing can be decomposed into four different but interrelated aspects. The purposes of this division are to:

- conceptually clarify the functions of routing;

- simplify the design of routing protocols by breaking routing into its component parts; and
- make the routing functions as flexible as is practical by allowing for degrees of freedom in each aspect.

The four aspects are described in the following clauses. Figure 1 below illustrates the relationship among these four aspects of routing.

5.1.1 Routing Information Base

The Routing Information Base comprises the complete information required by a particular ES or IS to accomplish routing. Such information might include:

- Next hop routing tables. These are tables which relate destination NSAPs to the potential next subnetwork hops (e.g. local and remote SNPs) which might be used to forward the PDU closer to its destination.
- Lists of neighbour ESs and ISs. These lists enable an ES or IS to ascertain the local topology.
- Measured QoS characteristics of a datalink or subnetwork path. These measurements allow the routing functions to adapt to QoS changes.

- Network maps. These are complete topological graphs of a portion of the global network. Such maps can be used to compute shortest paths to destination NSAPs using any of a number of routing metrics.

5.1.2 Information Collection

ESs and ISs build up their routing information bases by collecting information from their local environment and from other systems. Some example sources of information are: measurement protocols, policy input from System Management, directory lookup functions, and routing protocols. The information collection function is illustrated in figure 1 by the box labelled *Update Receive*.

5.1.3 Information Distribution

Systems may inform other systems of pertinent information in their local routing information base by distributing this information. Some examples of information distribution techniques include: routing protocols and interactions through the management information bases. The information distribution function is illustrated in figure 1 by the box labelled *Update Send*.




Figure 1 - Decomposition of the Routing Function

5.1.4 Route Calculation and Maintenance

These are the internal functions executed by ESs and ISs on the routing information base to accomplish routing. The major function in this category is the generation of the forwarding information base which is used to actually relay NPDUs. This function is illustrated in figure 1 by the box labelled *Decision*. Other examples of these internal functions include: timing functions such as ageing old routing information base entries, and the functions F1 and F2 described below.

5.1.4.1 Functions F1 and F2

The functions F1 and F2 are two functions required by every ES and IS to route an NPDU. The inputs to F1 are

- a) the called or destination NSAP address;
- b) the calling or source NSAP address;
- c) a source route (optional). A source route is a

sequence of network entity titles or network entity title prefixes which identify Network relay systems. See, for example, the source routing function of ITU-T Rec. X.233 | ISO/IEC 8473-1. In a complete source route the next network entity title in the sequence is the output of F1. In a partial source route, the next network entity title or network entity title prefix in the sequence is used to determine the network entity title of a Network relay system used to reach the Network relay identified by the source route.

- e) Quality of service (QoS) parameters (optional);
- f) the Forwarding Information Base.

For each NPDU that is routed, F1 determines

- f) The Network entity title of a Network relay system on the path to the destination NSAP or else,
- g) The title of the destination Network entity, if no relay function is necessary to reach the destination. The




Figure 2 - Routeing Exchange using Network Layer Protocols

title may be the same as the destination NSAP address.

The inputs to F2 are:

- h) The network entity title of the Network relay or destination End system determined by F1.
- i) QoS.
- j) the Forwarding information base.

This function is performed after F1 to determine which subnetwork point of attachment (SNPA) to use when sending an NPDU to the Network relay or destination network entity. The information yielded by this function is:

- k) identification of the selected SNPA.
- l) values of parameters which are input to the subnetwork service provider associated with that SNPA.

5.2 Relationship of Routing to OSI Management

Operation of the Network Layer, in fulfilment of the role assigned to it in the OSI Reference Model, requires shared knowledge concerning the location of NSAPs and routes through the available subnetworks.

As shown in figures 2 and 3, the routing function intersects with OSI management through information stored in, and retrieved from, the management information base (the boxes labelled *Local Information Base* in the figures). Routing information is placed in the management information base either through interaction with Network Layer entities or through interaction with System Management (the box labelled *OSI Management Application*).

It may be desirable to collect and distribute routing information automatically through the operation of OSI Routing protocols; these protocols may be located at the




Figure 3 - Routing Exchange using System Management

Network Layer (Layer management) or the application layer (System management).

The use of a "network layer routing information exchange protocol" has (among others) the following advantages:

- it confines the generation, exchange, and synchronisation of routing information within the Network layer. This keeps routing a "closed system" and avoids difficult issues in cross-layer co-ordination.
- it permits the efficient and direct use of subnetwork capabilities which may be available, such as inherent multicast.

Figure 2 illustrates the use of a layer management protocol to exchange routing information. Use of an "application layer routing information exchange protocol" has (among others) the following advantage:

- context negotiation and the establishment of management associations over a reliable end-to-end transport service is possible.

Figure 3 illustrates the use of a system management protocol to exchange routing information.

In general, it is likely that a complete and realistic solution to the global routing problem in the OSI environment will require a combination of techniques, involving both Network layer management protocols and System management protocols.

6 Environment for OSI Routing

OSI Routing shall be capable of operating effectively in a variety of environments, which when considered together result in a number of difficult goals for any Routing scheme to satisfy. This clause discusses some of the environments envisioned for OSI routing and points out some of their salient features from the point of view of routing.

6.1 Interconnection of LANs

LANs may be connected either locally through an Interworking Unit, or across larger distances via point-to-point subnetworks, multicast subnetworks, private leased-line networks, or public data networks. In all cases routing functions are needed to determine paths through the WAN that meet connectivity and/or QoS requirements. These routing functions may be arbitrarily complex, depending on whether organisational boundaries are being crossed, the need for optimal routes, resilience from failure, etc. In addition, the routing functions for LAN-WAN interconnection need to take account of the wide disparity of transmission bandwidth between the two environments.

6.2 Public/Private Network Interconnection

Many organisations already operate private data networks. In order to communicate with other organisations their

private networks may be interconnected using public network facilities. Routing functions shall be capable of efficient routing within both the private and public domains, while providing the organisational isolation necessary for the separate management of these domains. Further, the routing methods, metrics, and policies may be very different in the public domain than in the private domain. Routing functions shall be capable of successfully dealing with the limited control and data flow across these sorts of boundaries.

6.3 Factory and Campus Networks

The use of networks in factories and campus environments such as universities, corporate headquarters, government ministries, and research establishments is growing rapidly. These environments are characterised by large numbers of systems (sometimes in the thousands) connected by rich topologies. In these environments, the configuration of the network tends to change rapidly and the exercise of centralised control over the installation and operation of systems minimal. Routing schemes for this kind of environment need to be robust against unanticipated configuration changes and be able to adapt to changes in network usage, applications, and traffic patterns without the need for the extensive intervention of a centralised administrative function.

6.4 Multi-vendor Subnetworks

Networks are inherently a multi-vendor environment. It is extremely rare for a consumer to acquire an entire subnetwork (LAN or WAN) from a single vendor, since systems are purchased at different times for different applications. In many cases however, operators of private (and in some cases public) networks are forced to acquire all of their ISs from a single source because of the lack of routing standards. This situation limits the ability of organisations to build cost-effective networks and severely constrains the ways in which their networks can be interconnected with those of other organisations. An effective set of OSI routing standards will enable the construction of practical multi-vendor subnetworks, much as the rest of OSI has enabled systems from multiple vendors to interoperate.

7 Goals for OSI Routing

The environment identified for OSI routing in clause 6 results in some difficult goals for any OSI routing scheme. These goals are discussed in the following subclauses.

7.1 Multiple Subnetwork Types

The routing functions defined within this framework shall be designed to operate without regard to any specific underlying technology or transmission medium, to the extent that they do not rely upon any technology specific service for their correct operation. These functions shall also be designed to operate correctly irrespective of the geographic distribution of ESs and ISs which comprise the

global Routing Domain (i.e. they are not topology dependent).

7.2 Very Large Number of ESs and ISs

The global OSIE in which End System data are to be transferred is assumed to consist of a very large number of NSAPs ($>10^7$) which, in the most general situation, may be logically interconnected by means of paths consisting of concatenated intermediate systems. The total number of intermediate systems is assumed to be one to two orders of magnitude fewer than the number of NSAPs (or End Systems), but very large as well. Any routing scheme adopted for OSI shall be capable of indefinite scaling.

7.2.1 End Systems Should be Kept Simple

A consequence of the preponderance of End-systems over Intermediate systems is the desire to keep the ESs simple, even at the expense of making the ISs more complicated. This makes sense also because Intermediate systems are often dedicated to routing and relaying. End-systems, on the other hand, perform routing as an overhead function necessary to enable them to do their real job of executing applications.

7.3 Multiple Organisations

The presence of multiple organisations within the OSI environment will require the following attributes of OSI routing.

7.3.1 Distribution of Control

Global routing shall by necessity be able to operate correctly under the distributed control of multiple organisations. Furthermore, the control of routing within a single organisation may be distributed for reasons of efficiency, economy, performance, etc.

7.3.2 Trust, Firewalls, and Security

The exchange of routing information between ISs has the effects of

- a) allowing one IS to impact the routing decisions made by another IS; and
- b) providing one IS with information about another IS.

Of specific concern here is

- c) the effect of bad routing information exchanged between ISs in different administrative domains, and
- d) the implicit or explicit exchange of private, proprietary, or secret information across administrative domains.

The exchange of routing information across administrative boundaries should maximise the usefulness of that information while minimising the potential adverse effects of that information. Examples of adverse routing effects are routing loops and "black holes" (both of which can severely degrade network performance), incorrect routing,

and illegal routing (especially with regard to national boundaries). Furthermore, the routing information exchanged should provide a minimum of ancillary information while providing a maximum of routing functionality. Examples of ancillary information about the internals of an administrative domain might include

- topology information,
- size,
- level of activity,
- reliability,
- quality and/or type of service, and
- tariff structure.

Organisations operating administrative domains should be able to control the "leakage" of information outside their administrative domain(s). That is, they should be able to control the amount and kind of information which enters or leaves their administrative boundaries while still providing and receiving some minimum global routing capability.

When establishing administrative boundaries and exchanging routing information across those boundaries, strong authentication of Network entities may be required. Authentication of Network entities is necessary to prevent an IS belonging to one administrative domain from claiming to be a different IS belonging to another (possibly the local) administrative domain. Without authentication, administrative domains may be susceptible to a variety of external attacks, including the denial of service to large numbers of systems.

7.3.3 Routing Domains

The routing functions shall be designed to operate across multiple Routing Domains (see 8.1.2 for detailed information on Routing Domains). These Routing Domains may be private, in the sense that a given Routing Domain may make use of non-standard routing functions and protocols internally while supporting standard routing functions and protocols externally. A Routing Domain may, however, make use of standard routing functions and protocols both internally and externally. The routing functions shall be able to accommodate network topologies consisting of both of these Routing Domain types.

7.4 Performance

Performance is characterised by the efficiency and robustness of the Network Service as seen by participating systems. It is important that these systems avoid

- introducing a high degree of overhead (control) traffic; and
- concentrating traffic on a few paths when other paths are relatively free, thus causing unnecessary congestion.

OSI routing is expected to be "fair" to all participating Network entities in the sense of providing equitable access

that is only constrained by the bandwidths of the connected SNPAs. Measures shall be taken to avoid service denial and to allow for a "reasonable" throughput and response time to all systems. Good performance also requires a dynamic adaptive capability in the Network entities to respond to various types of failures. The routing scheme should be able to adapt appropriately to changes in topology, including partitions.

7.5 Existing Network Layer Protocols

Any scheme adopted for OSI Routing should not have an excessive impact on existing Network Layer Protocols. An ideal routing scheme would not require any modifications to current network layer protocol standards nor necessitate their re-implementation to accommodate the routing functions.

7.6 Communication Type Independence

Routing procedures should accommodate both Connection-mode and Connectionless-mode communication. However, Routing procedures may vary for the two types. For instance, routing loops are easier to tolerate for a few connectionless-mode PDUs than for an entire Network Layer connection. Optimisations may be made by taking advantage of the characteristics of either form of communication, however, these optimizations may be at the expense of communication type independence.

7.7 Resilience

The Routing procedures should be able to respond to physical or logical connectivity failures by finding routes around those failures. A tradeoff must be made between speedy and robust recovery from connectivity failures and routing overhead and complexity. A connectivity failure which results in a partition is more difficult to recover from than a failure which does not cause a partition.

7.8 Routing Procedure Diagnosis

Routing procedures should be diagnosable. This diagnosis ranges from solving severe problems in the routing procedures which cause them to fail, to optimising the routing procedures for a particular environment. Creating routing procedures which are diagnosable may involve avoiding certain algorithms, adding error and probing PDUs to a protocol, and maintaining trace and error state in systems.

8 Structure of Global OSI Routing

8.1 Categories of Routing

The development of an architectural framework for routing within the OSI environment proceeds from two basic principles:

- a) The aspects of routing that are concerned with communication between ESs and ISs to collect configuration information and distribute configuration and redirection information are to a great extent

separable from the aspects that are concerned with communication among the Intermediate systems that connect multiple subnetworks.

- b) Establishing communication among Intermediate systems to connect multiple subnetworks presents both technical and administrative challenges. The technical issues have to do with establishing global connectivity and performing global routing functions; the administrative issues have to do with controlling the way in which groups of systems managed by different administrative authorities are permitted to communicate.

These two principles lead to a decomposition of the global routing framework into three categories. See figure 4 which illustrates these categories of routing.

The first principle establishes an initial distinction between local ES-IS operation and global IS-IS operation. There are a number of technical advantages to such a distinction. These include

- a protocol which specifically addresses communication between ESs and ISs can be designed to place the more difficult and complicated procedures in the ISs, thus keeping the End systems simple;
- a specialised ES-IS protocol can be made relatively independent of the routing procedures used among ISs. This allows more than one IS-IS protocol to be used, if necessary, without burdening the End-systems;
- often the topology which connects End-systems to Intermediate systems is much richer and more highly connected than the topology employed to connect ISs together (e.g. LANs, PDNs). The operation of both ESs and ISs can be enhanced by exploiting this difference and designing a separate routing procedure for the two classes of topology.
- the topology which connects Intermediate Systems may support multicast communication. The operation of ISs can be enhanced by exploiting the multicast communication capability in the design of the routing procedures.

The second principle establishes a further distinction between IS-IS operation within the purview of a single (possibly composite) organisation ("Intra-Administrative Domain") and IS-IS operation that spans one or more significant administrative boundaries ("Inter-Administrative Domain").

8.1.1 ES-IS Routing

An ES-IS protocol may operate to establish connectivity and reachability between End systems and Intermediate systems where this is not an inherent service of the subnetwork service provider.

The operation of an ES-IS protocol ensures that each End System knows about at least one IS that is directly reachable

on each subnetwork to which the ES is attached, and each IS knows about every End System reachable on each subnetwork to which the IS is attached.

In figure 4, an ES-IS routing protocol is illustrated by the curved dotted lines which associate each ES with at least one IS on the local subnetwork. The subnetwork paths contained within Routing Domains are illustrated using straight solid lines. These subnetwork paths may support point-to-point and/or multicast communication.

ISO 9542 is an ES-IS protocol designed to operate in close conjunction with ISO/IEC 8473-1 | ITU Rec. X.233. It makes extensive use of LAN based multicast communication to collect and distribute routing information. The use of ISO 9542 when ISO/IEC 8473-1 | ITU-T Rec. X.233 is operated over a WAN is not precluded.

ISO/IEC 10030 is an ES-IS protocol designed to operate in close conjunction with ISO 8208. It makes use of a *Subnetwork Address Resolution Entity* (SNARE), which collects configuration information from ESs and distributes configuration information and redirection information to both ESs and ISs. The use of a ISO/IEC 10030 SNARE for obtaining a *Subnetwork Point of Attachment* (SNPA) for the forwarding of ISO/IEC 8473-1 | ITU-T Rec. X.233 PDUs is not precluded.

The above characteristics of the routing protocols can be further refined to recognize a division between the relay functions performed by an IS and the collection and distribution of reachability information performed by a SNARE. While it might be desirable in some cases to combine these functions within a single open system, this is not required in all cases.

8.1.2 Routing Domains

The global OSIE will, of necessity, be composed of multiple domains of administrative responsibility and multiple domains of routing responsibility. An administrative domain wholly encompasses one or more Routing Domains.

A Routing Domain is a set of ISs and ESs bound by a common routing procedure; namely:

- they use the same set of routing metrics;
- they use compatible metric measurement techniques;
- they use the same information distribution protocol; and
- they use the same path computation algorithm.

8.1.2.1 Formal Definition of a Routing Domain

A Routing Domain D can be defined formally as a couplet (S, R) where S is the set of ISs and ESs in the domain and R is the common routing procedure. It is understood that

- a) every IS within a domain D can determine if a given NSAP is reachable within D . If it is, then the routing procedure is capable of deriving a path to

that NSAP;

- b) an IS α within a domain D has means of ascertaining if another neighbouring IS β participates in D ;
- c) an IS may participate in more than one Routing Domain. In such a case,
 - 1) the IS will fully and completely, but independently, participate in the routing procedures of each domain;
 - 2) routing information from one domain will not be utilised in any way in the routing procedures of the other; and
 - 3) when an IS participating in two domains D_a and D_b receives a PDU from an ES, the IS will have to determine in which domain this PDU will be routed.

8.1.2.2 Hierarchical Structure of Routing Domains

There are a number of reasons why it may be useful to organise a Routing Domain or a collection of domains in a hierarchical fashion. As the number of ESs and ISs in a Routing Domain increases, it becomes more difficult to maintain and process all of the information necessary to perform the routing functions. Typically, the size of the routing information base, the exchange of routing update information, and the computation of routes may consume more resources than are allocated to route determination in the domain.

In order to reduce the overhead associated with route determination, it is often useful to divide a Routing Domain into smaller routing subdomains. Each routing subdomain maintains detailed routing information about its own internal composition, and also maintains routing information which allows it to reach other routing subdomains. Because these other routing subdomains in turn maintain detailed routing information about their own internal composition, there is no need to maintain detailed routing information for the Routing Domain in all routing subdomains.

A significant benefit of this technique is that the number of entries in the routing information base maintained by a Network entity in a given routing subdomain may be reduced to the number of Network entities within the routing subdomain plus the number of other routing subdomains in the Routing Domain. This reduction of the routing information base results in a proportional reduction in the exchange of routing update information and in turn reduces the load imposed by the computation of routes in the Routing Domain. Hierarchical decomposition of Routing Domains can be recursively applied in order to achieve further reductions when necessary.

8.1.3 Intra-administrative Domain Routing

Intra-Administrative Domain routing is concerned with communication among Intermediate systems that are all

managed and controlled by a single (possibly composite) administrative authority. The Administrative Domain controls the organisation of the intermediate systems into Routing Domains, the assignment of NSAP and subnetwork addresses, the way in which the costs of operation are determined and recovered, and the policies that govern the flow of information. Within an Administrative Domain, the organisation responsible for each Routing Domain may establish within that domain any number of hierarchically organised subdomains.

8.1.3.1 Intra-domain Routing

In ISO/IEC 10589 (IS-IS), two levels of subdomain are defined. A Level 1 subdomain - termed an Area - consists of a set of ESs and ISs, in which the ISs maintain detailed information for routing among all the ESs in the Area.

When a Routing Domain consists of two or more Areas, a subset of the ISs from each Area are configured to form the Level 2 subdomain. These Level 2 ISs maintain the information necessary for routing between Areas. When an OSIE consists of more than 1 Routing Domain, Level 2 ISs also perform the functions necessary for routing traffic among those routing domains.

In figure 4, two types of intra-domain routing are shown. The first is routing within an IS-IS Area Boundary, illustrated by the thin curved dashed lines connecting the intermediate systems within each IS-IS Area, and corresponds to intra-domain level 1 routing (see ISO/IEC 10589). The second is routing between IS-IS Areas within the same Routing Domain, illustrated by the solid curved line between the two IS-IS Areas, and corresponds to intra-domain level 2 routing (see ISO/IEC 10589). Note that it is by no means required that there be more than one Routing Domain within a single Administrative Domain. They may in fact be congruent, as shown in figure 4 by the domain on the lower right.

In Figure 4, the subnetwork paths contained within IS-IS Area Boundaries and Routing Domains are illustrated using straight solid lines. For illustrative purposes, intermediate systems performing both intra-domain routing and inter-domain routing (see 8.1.3.2 and 8.1.4.1) are labeled as BISs.

Subnetwork paths contained within IS-IS Area Boundaries and Routing Domains may support point-to-point and/or multicast communication.

8.1.3.2 Inter-domain Routing

In Figure 4, routing between Routing Domains within a single Administrative Domain is illustrated by the thick dashed curved line between the two Routing Domains, and corresponds to inter-domain routing (see ISO/IEC 10747). The subnetwork paths between Routing Domains are illustrated using straight solid lines.

Subnetwork paths between Routing Domains may support point-to-point and/or multicast communication.

8.1.4 Inter-Administrative Domain Routing

Routing between Administrative Domains is concerned with managing and controlling the exchange of information between intermediate systems at a level that requires significant formal co-operation between different organisations. The issues of concern in this environment are for the most part administrative: security, access control, national regulations, the legal and political implications of trans-border data flow, etc. The techniques used to accomplish the actual routing function may be the same as those used within an Administrative Domain (in fact, the same protocol may be used); the context in which they are employed is, however, fundamentally different when exchanging routing information between Administrative Domains.

8.1.4.1 Inter-domain Routing

In figure 4, routing between Administrative Domains is illustrated by the thick dashed curved line between the two Administrative Domains, and corresponds to inter-domain routing (see ISO/IEC 10747). The subnetwork path between the two Administrative Domains is illustrated using a straight solid line.

Subnetwork paths between Administrative Domains may support point-to-point and/or multicast communication.

8.2 Relationship between Network Addresses and Routing

The Network Layer addresses of systems deployed in the OSIE are governed by CCITT Rec. X.213 | ISO/IEC 8348. Routing uses these Network Layer addresses in routing tables as values or indexes in order to derive routes for the PDUs to follow.

From the perspective of routing, it is important that the size of the tables be minimized, and that the processing time consumed for routing be minimized, and yet to permit as flexible a deployment of systems as possible. In the present context, flexible means that the systems are deployed without having to take (too much) into consideration about what the Network Addresses of the systems are.

This is achieved by several means:

- Paramount is the fact that routing is organized hierarchically, as indicated in 8.1.
- In complement, at each of the routing hierarchy levels, the routing protocols relevant to that level may set those guidelines and/or rules that the Network Addresses or NETs of the systems deployed at that level should follow.
- In general those rules will include provisions for exceptions, that is systems whose addresses or NETs don't meet the rules will still be able to participate in the OSI routing at that level.




Figure 4 - Categories of OSI Routing

- Sometimes, systems that do not abide by the recommended deployment rules will get under-optimized routing performance, or they may impact other systems or the memory usage in Intermediate Systems. In certain cases the routing protocol will not operate correctly if the rules are not met.

It behooves the Routing Administrative Authority (that is the person/organization in charge of deploying systems) at any given level of the routing hierarchy to make sure that guidelines are followed and rules are met such as those in this Technical Report and in the specifications of the Routing protocols, since failure to do so may result in poor performances or even in loss of connectivity (Routing protocols may detect some mis-deployment and not be aware of others).

RFC 1629, "Guidelines for OSI NSAP address allocation in the Internet" also provides further guidance.

8.3 Routing Procedures

OSI may adopt different routing procedures for the different categories of routing identified in 8.1. The reasons for this are many:

- across a single subnetwork, it is desirable to simplify the operation of End systems by off-loading routing functions to the Intermediate systems. This is the case because
 - there are likely to be many more End systems than Intermediate systems;
 - Intermediate systems can be essentially dedicated to the functions of Routing and Relaying; and
 - End systems are less likely to be attached to multiple subnetworks than Intermediate systems and hence have fewer routing choices to make.
- within an Administrative Domain, a single organisation may obtain significant benefits from a sophisticated, dynamic routing scheme which produces optimal routes with acceptable overhead. These organisations may wish to experiment with non-standard routing procedures within their domain(s) while remaining connected to the global OSIE through a standard routing procedure.
- between Administrative Domains, the need to maintain an "arms-length" relationship will restrict the type and detail of routing information available. More stringent procedures for authenticating and propagating routing information may be needed as well.

In order to analyse the strengths and weaknesses of various routing procedures, it is useful to have a taxonomy which can be used to select among the techniques. Routing algorithms may be classified according to how they accomplish the four aspects of routing defined in Clause 5.1, and what types of information are used to select routes.

The following subclauses provide a taxonomy of routing procedures. Included for each class of routing procedure is an indication of its ability to react to a variety of dynamic changes in the global network, such as

- a) changes in topology, due to systems and subnetwork paths coming up or going down,
- b) changes in configuration, such as the addition of new systems/subnetworks or their removal,
- c) changes in the patterns of traffic in the network, and
- d) changes in the Quality of Service available on certain subnetworks or using certain paths.

In addition to the subclauses which introduce each technique, table 1 summarises their operational characteristics according to the four aspects of routing defined in 5.1.

8.3.1 Static Routing

In *static routing* all routing information known to a system is loaded into the Routing Information Base by System management. This information is generally in "pre-computed" form, in that only the paths actually to be used are made available rather than all possible paths. In essence, Static Routing performs the *Decision* function from figure 1 in an off-line fashion and uses System Management protocols to communicate the resulting routing tables to each system.

Static routing has the advantage of permitting extremely sophisticated off-line optimisation algorithms to be executed, since the route computation need not be done in real time while PDUs are being relayed. It has the disadvantages of not being capable of "bootstrapping" the NS-providers since there is no information collection or dissemination by the Network entities themselves. Further, static routing is not capable of reacting to configuration, topology, traffic pattern, or QoS changes in an adaptive fashion since all paths are pre-computed.

8.3.2 Quasi-static Routing

Quasi-static routing is similar to static routing in that paths are computed off-line and loaded into the Routing Information Base through System management. Rather than storing a single, highly optimised path for each routing metric, however, quasi-static routing allows for alternate paths to be stored. This reduces the impact of failures by allowing the forwarding function to select a backup path if the best path is unavailable.

Quasi-static routing has similar advantages and disadvantages to static routing. It can, however, adapt to topological changes in a limited way, at the expense of an increase in the amount of information stored concerning backup paths. Quasi-static routing cannot adapt to configuration, traffic pattern, or QoS changes.