

# TECHNICAL REPORT

**ISO/IEC**  
**TR 9575**

First edition  
1990-06-01

---

---

## **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 : 1990 (E)

## Contents

|   |                                        |     |
|---|----------------------------------------|-----|
|   | Foreword.....                          | iii |
|   | Introduction .....                     | iv  |
| 1 | Scope .....                            | 1   |
| 2 | Normative References .....             | 1   |
| 3 | Definitions .....                      | 1   |
| 4 | Symbols and abbreviations .....        | 2   |
| 5 | Routeing Concepts .....                | 2   |
| 6 | Environment for OSI routeing .....     | 5   |
| 7 | Goals for OSI Routeing .....           | 6   |
| 8 | Structure of global OSI Routeing ..... | 8   |

© ISO/IEC 1990

All rights reserved. No part of this publication may be reproduced or utilized in any form by any means, electronic or mechanical, including photocopying and microfilm, without permission in writing from the publisher

ISO/IEC Copyright Office, P.O. Box 56, 1211 Geneva 20, Switzerland  
Printed in Switzerland.

## Foreword

ISO (the International Organization for Standardization) and IEC (the International Electrotechnical Commission) together form a system for worldwide standardization as a whole. 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 a technical committee is to prepare International Standards but in exceptional circumstances, the publication of a Technical Report of one of the following types may be proposed:

- 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 ISO/IEC JTC1, *Information technology*.

## 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.

# 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.

## 2 Normative References

The following standards 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 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 listed below. Members of IEC and ISO maintain registers of currently valid International Standards.

ISO 7498:1984, *Information processing systems — Open Systems Interconnection — Basic Reference Model*.

ISO 7498/Add.1:1984, *Information processing systems — Open Systems Interconnection — Basic Reference Model — Addendum 1: Connectionless-mode Transmission*.

ISO 7498/Add.3:1989, *Information processing systems — Open Systems Interconnection — Basic Reference Model — Addendum 3: Naming, including Addressing*.

ISO 7498/Add.4:1989, *Information processing systems — Open Systems Interconnection — Basic Reference Model — Addendum 4: Management Framework*.

ISO 8348:1987, *Information processing systems — Data communications — Network Service Definition*.

ISO 8348/Add.1:1987, *Information processing systems — Data communications — Network Service Definition — Addendum 1: Connectionless-mode transmission*.

ISO 8348/Add.2:1988, *Information processing systems — Data communications — Network Service Definition — Addendum 2: Network layer addressing*.

ISO 8473:1988, *Information processing systems — Data communications — Protocol for providing the connectionless-mode network service*.

ISO 8648:1988, *Information processing systems — Open Systems Interconnection — Internal organisation of the Network Layer*.

ISO 9542:1988, *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 8473)*.

## 3 Definitions

### 3.1 Reference Model definitions

This Technical Report makes use of the following terms defined in ISO 7498:

- 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 ISO 8348/Add.2:

- 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 CCITT. A CCITT 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 8.2.1 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.

## 4 Symbols and abbreviations

|       |                                          |
|-------|------------------------------------------|
| 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           |
| OSIE  | Open Systems Interconnection Environment |
| WAN   | Wide Area Network                        |

## 5 Routing Concepts

### 5.1 Functional decomposition of routing

OSI Routing can be decomposed into four different but inter-related 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 SNPAs) 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




Figure 1 - Decomposition of the routing function

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*.

### 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

- the called or destination NSAP address;
- the calling or source NSAP address;
- a source route (optional). A source route is a sequence of network entity titles which identify Network relay systems. See, for example, the source routing function of ISO 8473. 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 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.
- Quality of service (QoS) parameters (optional);
- the Forwarding Information Base.

For each NPDU that is routed, F1 determines

- The Network entity title of a Network relay system on the path to the destination NSAP or else,
- The title of the destination Network entity, if no relay function is necessary to reach the destination. The 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 1, 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 Network Layer



### Figure 2 - Routing Exchange using Network layer protocols



(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.
- capabilities existing in the lower layers but not available in the upper layers, can be used (e.g. 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 1 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 sat-




Figure 1 - Routing exchange using system management

isfy. 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, 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 with 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 8.2.1 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 SNPA's. 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 reimplementation 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 optimisations 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 on the same subnetwork 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 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 Single Subnetwork routing

Across a single subnetwork, an ES-IS protocol (e.g. ISO 9542) 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 lines which associate each ES with at least one IS on the local subnetwork.

Often, the ES-IS protocol is closely associated with the particular Network layer protocol used to carry data across that subnetwork. For example, ISO 9542 is designed to operate in close conjunction with ISO 8473 and exploit the connectionless environment. It makes extensive use of the broadcast connectivity of LANs to collect and distribute routing information, and takes advantage of the independence of individual PDUs to offer new routes to End systems as PDUs are being forwarded.

Conversely, in connection-oriented subnetworks, an approach to ES-IS routing might instead utilise a specialised entity on the subnetwork to disseminate routing information to End systems. Such a function is referred to as a *Subnetwork Address Resolution Entity*, or SNARE. The SNARE function exploits the ability of some Connection-mode protocols to deflect calls from one system on a subnetwork to another, and also may make use of the connection context to allow End systems to reliably query for and obtain information about available Intermediate systems.

#### 8.1.2 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 ISs 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.

In figure 4, two types of routing within an administrative domain are shown. The first is routing within a Routing Domain, illustrated by the curved dotted lines connecting the ISs within each routing domain. The second is routing between Routing Domains, illustrated by the solid curved line between the two routing domains in the upper left Administrative Domain. 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.

#### 8.1.3 Inter-Administrative Domain routing

Routing between Administrative Domains is concerned with managing and controlling the exchange of information between ISs 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