INTERNATIONAL STANDARD

ISO/IEC 10021-10

First edition
1998-09-01

Information technology — Message Handling Systems (MHS) —

Part 10:
MHS routing

Technologies de l'information — Systèmes de messagerie (MHS) —
Partie 10: Routage MHS

Reference number
ISO/IEC 10021-10:1998(E)

Contents

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. Draft International Standards adopted by the joint technical committee are circulated to national bodies for voting. Publication as an International Standard requires approval by at least 75% of the national bodies casting a vote.

International Standard ISO/IEC 10021-10 was prepared by Joint Technical Committee ISO/IEC JTC 1, *Information technology*, Subcommittee SC 18, *Document processing and related communication*.

ISO/IEC 10021 consists of the following parts, under the general title *Information technology - Message Handling Systems (MHS):*

- *Part 1: System and service overview*
- *Part 2: Overall architecture*
- *Part 3: Abstract Service Definition Conventions*
- *Part 4: Message Transfer System: Abstract service definition and procedures*
- *Part 5: Message Store: Abstract service definition*
- *Part 6: Protocol specifications*
- *Part 7: Interpersonal messaging system*
- *Part 8: Electronic data interchange messaging service*
- *Part 9: Electronic data interchange messaging system*
- *Part 10: MHS routing*

Annexes A to D form an integral part of this part of ISO/IEC 10021. Annexes E to H are for information only.

Introduction

This part of ISO/IEC 10021 is one of a number of parts of ISO/IEC 10021 defining Message Handling in a distributed open systems environment.

Message Handling provides for the exchange of messages between users on a store-and-forward basis. A message submitted by one user (the originator) is transferred through the message-transfer-system (MTS) and delivered to one or more other users (the recipients).

This part of ISO/IEC 10021 defines a method for routing messages through the Message Handling System (MHS).

Information technology – Message Handling Systems (MHS) – Part 10: MHS routing

1 Scope

This part of ISO/IEC 10021 specifies the means by which messages are routed through the MHS, and supplements the procedures defined in 14.3 of ISO/IEC 10021-4.

Other parts of ISO/IEC 10021 define other aspects of the MHS. ITU-T Rec. F.400/X.400 | ISO/IEC 10021-1 defines the user-oriented services provided by the MHS. ITU-T Rec. X.402 | ISO/IEC 10021-2 provides an architectural overview of the MHS. ITU-T Rec. X.411 | ISO/IEC 10021-4 defines the abstract-service of the Message Transfer System.

2 Normative references

The following standards contain provisions which, through reference in this text, constitute provisions of this part of ISO/IEC 10021. At the time of publication, the editions indicated were valid. All standards are subject to revision, and parties to agreements based on this part of ISO/IEC 10021 are encouraged to investigate the possibility of applying the most recent editions of the standards indicated below. Members of IEC and ISO maintain registers of currently valid International Standards.

2.1 Presentation references

This part of ISO/IEC 10021 cites the following Presentation specifications:

— ITU-T Recommendation X.680 (1994) | ISO/IEC 8824-1: 1995, *Information technology - Abstract Syntax Notation One (ASN.1): Specification of basic notation.*

— ITU-T Recommendation X.681 (1994) | ISO/IEC 8824-2: 1995, *Information technology - Abstract Syntax Notation One (ASN.1): Information object specification.*

2.2 Directory references

This part of ISO/IEC 10021 cites the following Directory specifications:

— ITU-T Recommendation X.500 (1997) | ISO/IEC 9594-1: —[1], *Information technology —Open Systems Interconnection —The Directory: Overview of Concepts, Models, and Services.*

— ITU-T Recommendation X.501 (1997) | ISO/IEC 9594-2: —[1], *Information technology —Open Systems Interconnection —The Directory: Models.*

— ITU-T Recommendation X.520 (1997) | ISO/IEC 9594-6: —[1], *Information technology —Open Systems Interconnection —The Directory: Selected attribute types.*

— ITU-T Recommendation X.521 (1997) | ISO/IEC 9594-7: —[1], *Information technology —Open Systems Interconnection —The Directory: Selected object classes.*

2.3 Message Handling references

This part of ISO/IEC 10021 cites the following Message Handling System specifications:

— ITU-T Recommendation F.400/X.400 (1996), *Message handling services: Message handling system and service overview.*

ISO/IEC 10021-1: —[1], *Information technology — Message Handling Systems (MHS) — Part 1: System and service overview.*

---

[1] To be published

— ITU-T Recommendation X.402 (1995) | ISO/IEC 10021-2: 1996, *Information technology — Message Handling Systems (MHS): Overall architecture.*

— ITU-T Recommendation X.411 (1995) | ISO/IEC 10021-4: 1997, *Information technology — Message Handling Systems (MHS): Message transfer system: Abstract service definition and procedures.*

2.4 Country Code references

This part of ISO/IEC 10021 cites the following Country Code specifications:

— ISO 3166-1: 1997, *Codes for the representation of names of countries and their subdivisions — Part 1: Country codes.*

— CCITT Recommendation X.121 (1996), *International number plan for public data networks.*

2.5 Additional references

This part of ISO/IEC 10021 cites the following specification:

— ISO/IEC 9945-2: 1993, *Information technology — Portable Operating System Interface (POSIX) — Part 2: Shell and Utilities.*

3 Definitions

For the purposes of this part of ISO/IEC 10021, the following definitions apply.

3.1 MHS-routing definitions

The following terms are defined in clauses 6 and 7 of this part of ISO/IEC 10021:

— connection-group

— entry-connection-group

— enumerated connection-group

— indirect-exit-connection-group

— key-routing-collective

— local-exit-connection-group

— local-use-tables

— MHS-routing

— next-MTA

— OR-address-element

— OR-address-subtree

— routing-advice

— routing-collective

— routing-collective-subtree

— routing-MTA

— transit-exit-connection-group

— unenumerated connection-group

A glossary of these terms appears in Annex H.

3.2 MHS definitions

The following terms are defined in ITU-T Rec. X.402 | ISO/IEC 10021-2:

— Administration Management Domain (ADMD)

— Management Domain (MD)

— Message Handling System (MHS)

— Message Transfer Agent (MTA)

— Message Transfer System (MTS)

— Originator/recipient Address (OR-address)

— Private Management Domain (PRMD)

— Reliable Transfer Service Element (RTSE)

For the purposes of this part of ISO/IEC 10021, the term *message*, where used unqualified, refers generically to a *message*, *probe*, or *report*.

3.3 Directory definitions

The following terms are defined in ITU-T Rec. X.501 | ISO/IEC 9594-2:

— Directory Information Tree (DIT)

— Directory System Agent (DSA)

— Directory User Agent (DUA)

— Relative Distinguished Name (RDN)

4 Abbreviations

The abbreviations used in this part of ISO/IEC 10021 are defined in ITU-T Rec. X.402 | ISO/IEC 10021-2, and ITU-T Rec. X.501 | ISO/IEC 9594-2 (see 3.2 and 3.3), except for the following.

— ACSE Association Control Service Element (ITU-T Rec. X.217 | ISO/IEC 8649)

— APS Asynchronous Protocol Specification (ITU-T Rec.X.445)

— IP Internet Protocol

— LAN Local Area Network

— PSAP Presentation Service Access Point (ISO/IEC 7498-3)

— WAN Wide Area Network

— X.25 A packet-switched network conforming to ITU-T Rec. X.25

5 Conventions

This part of ISO/IEC 10021 uses the descriptive conventions listed below.

5.1 Conventions for routing model specification

This part of ISO/IEC 10021 uses the following ASN.1-based descriptive conventions for the indicated purposes:

a) To define the data types and values for MHS-routing, ASN.1 itself.

b) To define the Directory entries for MHS-routing, the OBJECT-CLASS, ATTRIBUTE, NAME-FORM, STRUCTURE-RULE, and MATCHING-RULE information object classes of ITU-T Rec. X.501 | ISO/IEC 9594-2.

Whenever this part of ISO/IEC 10021 describes a class of data structure having components, each component is categorized as one of the following **grades**:

a) **Mandatory** (M): A mandatory component shall be present in every instance of the class.

b) **Optional** (O): An optional component may be present in an instance of the class at the discretion of the object (e.g. user) supplying that instance.

c) **Conditional** (C): A conditional component shall be present in an instance of the class as specified by this part of ISO/IEC 10021.

5.2 General font conventions

Throughout this part of ISO/IEC 10021, terms are rendered in **bold** when defined. Terms that are proper nouns are capitalized, generic terms are not. Multi-word generic terms are hyphenated. Italic font is used for ASN.1 identifiers defined in other International Standards.

5.3 Font conventions for ASN.1 definitions

Throughout this part of ISO/IEC 10021, ASN.1 definitions appear in fixed-pitch font (`of this Type`) to highlight the difference between normal text and ASN.1 definitions. The font used for ASN.1 definitions is smaller than that used for normal text.

5.4 Rules for ASN.1 definitions

ASN.1 definitions appear both in the body of this part of ISO/IEC 10021 to aid the exposition, and again, formally, in annexes for reference. If a difference is found between the ASN.1 used in the exposition and that formally defined in the corresponding annex, a specification error is indicated.

6 MHS-routing Overview

The purpose of a Message Handling System (MHS) is to enable users to exchange messages on a store-and-forward basis. A message submitted on behalf of one user, the originator, is conveyed by the Message Transfer System (MTS) and subsequently delivered to the agents of one or more additional users, the recipients. The MTS comprises a collection of Message Transfer Agents (MTAs), which are highly distributed, and a message may traverse a number of MTAs on the journey from its originator to its recipient.

In MHS, the originator does not specify a path through the MTS to reach a recipient, but simply specifies a recipient OR-name (from which the OR-address is obtained). It is the responsibility of each MTA to determine the next MTA to which the message should be transferred to progress its journey to its recipient. Routing is thus the process of selecting, given an OR-address, the MTA to which the message should next be transferred.

The other parts of ISO/IEC 10021 specify the services provided by MHS, and the protocols used to transfer messages within the MTS. The means by which an MTA determines an appropriate route that will convey a message to its recipient is the subject of this part of ISO/IEC 10021. The various mechanisms defined herein which enable MTAs to make this determination constitute MHS-routing.

The path taken between an originator and recipient may vary on different occasions, since there will in general be a number of possible paths between them, and factors such as congestion and availability may influence route selection.

MTAs acquire routing information by accessing Directory entries whose maintenance is the responsibility of an MHS administrator. These entries model the various properties of the MHS that are relevant to routing. However, MHS-routing does not depend on the provision of a fully interconnected Directory.

6.1 Operational characteristics

MHS-routing has the following operational characteristics:

a) In MHS-routing, the OR-address name-space is decoupled from the constraints of hardware organization (e.g. the assignment of users to particular machines, or the ability of groups of machines to interconnect). This is in contrast with routing strategies that use OR-address pattern-matching methods to make routing decisions, which constrain an MHS administrator's choice in the allocation of OR-addresses to users, and require users to change their OR-addresses whenever their MTA changes.

   OR-addresses are intended to reflect the organizational hierarchy, and to use only as many levels of OR-address element as is required to achieve this. However, many organizations have staff distributed over multiple locations (where staff will necessarily use a local MTA), or have multiple messaging systems in use (e.g. mainframe-based, integrated office automation systems, and PC LAN email systems) where staff in any one department are arbitrarily allocated to different systems. Separating the OR-address hierarchy from the physical topology allows for the assignment to users of addresses which are compact and related to their organizational role, regardless of their physical location.

b) MHS-routing supports the range of connection densities possible among MTAs. One extreme is where all MTAs are connected to a common network and any MTA can connect to any other (e.g. a public wide-area network). At the other extreme an administrator specifies precisely which MTA pairs are able to communicate, as if the MTAs were connected by individual cables, regardless of the actual connection method.

   Some MHS installations restrict the topology (e.g. by performing all routing through a central switch MTA), because of the management cost in maintaining routing information, even when all the MTAs concerned are connected to a LAN or WAN and could, in principle, exchange messages directly. However, moving to the fully connected possibility may be undesirable for some organizations, due to concerns of security and cost optimization.

c) MHS-routing permits varying degrees of autonomy and segregation in the management of an MHS system. Many organizations will have the overall MHS managed by a central IT unit, with the management of individual system components devolved to local administrators. Often, the MHS management hierarchy will be different from the organizational hierarchy. Typically, the MHS management reflects geographical locality, while the users are organized into geographically dispersed business units. In MHS-routing, this practice is accommodated by the use of two separate hierarchies for these two aspects of MHS organization.

d) MHS-routing accommodates a range of routing problem sizes, from an organization managing two or three MTAs, up to a PRMD operating hundreds or thousands of MTAs. An organization might organize only its internal connectivity, using ADMDs or specific bilateral agreements for all external connectivity, or it might join a more open community which publishes routing information in the Directory and allows for the open exchange of messages across a common network infrastructure.

e) MHS-routing does not constrain the choice of routing policies available to an organization. It provides a framework for the management of relevant routing information such that a range of routing strategies may be implemented.

f) MHS-routing provides a balanced trade-off in terms of cost/efficiency. Any scheme which relies on access to the Directory is likely to be slower to execute than one based purely on local tables held in the MTA. However, the Directory-based approach should provide the MTA with a better quality of information and so improve overall efficiency by the determination of optimal routes. For example, in a large PRMD it may often be infeasible to provide each MTA with customised tables listing every address in the domain, and consequently messages will take multiple hops towards their destination, passing through central MTAs that hold the tables. When Directory-based MHS-routing is used, the processing required at the originating MTA is greater, but is likely to identify a direct route and so enable transfer of the message in a single step.

   The use of MHS-routing also leads to a more scalable architecture, allowing for medium-capacity MTAs to be distributed throughout the network rather than depending on a small number of MTAs to perform central message switching. This scalable approach reduces the risks associated with a single point of failure on the network.

6.2 Components of the model

The different types of information stored in the Directory represent the modelling objects used to represent the MHS. Use of Directory access-control is not required for the operation of MHS-routing, but it may be employed where it is required to limit visibility to browsers of the Directory. Equally, information may be stored in a private DIT that has no connection to the outside world - there is no assumption that a fully interconnected Directory is available. However some routing policies, particularly those of an unrestricted nature, would benefit from a fully interconnected Directory.

6.2.1 Routing-collective

A **routing-collective** is a collection of one or more MTAs, under common management, which has collective responsibility for a portion of the OR-address name-space, and is capable of routing a message to any MTA managed within the collective. Therefore a routing-collective represents the management structure of some part of the MHS in the context of routing. By grouping together MTAs under common management control, the routing-collective provides abstraction (the internal structure of the routing-collective may be concealed from external MTAs) and is useful in limiting the scope of the information that MTAs need to hold.

The smallest instance of a routing-collective is a single MTA, and indeed each MTA is itself defined as a routing-collective. The largest instance of a routing-collective will typically be a Management Domain. While a routing-collective which comprises several PRMDs is not precluded, it is unusual for such PRMDs to be under genuinely common management control; the more typical case of loose confederations of PRMDs may be handled by other mechanisms.

Routing-collectives are structured in a hierarchical manner, with the number of levels chosen to suit the organization in question. Small MDs, or MDs of a very uniform nature under the complete control of a central IT management unit, will require only a two-level hierarchy. The MD is one routing-collective and contains all the MTAs (each of which is itself a routing-collective). More complex MDs will normally require a third level of routing-collective perhaps to group together all of the MTAs at a given location, or all those belonging to each department.

Each complete hierarchy of routing-collectives is represented by a DIT subtree, the **routing-collective-subtree**, which models their hierarchical relationships. All routing-collectives within a routing-collective subtree co-operate with one another in performing MHS-routing.

Each routing-collective has an entry in the Directory that indicates the connection-groups (see 6.2.3) by which a message may be transferred into the routing-collective, and those by which a message may be transferred out of the routing-collective.

A **routing-MTA** is an MTA that participates in MHS-routing as defined in this part of ISO/IEC 10021. By definition, a routing-MTA is the smallest instance of a routing-collective and occupies the lowest level in the routing-collective hierarchy, i.e. it appears as a leaf of the routing-collective subtree.

The routing-MTA is provided, by local configuration, with its routing-MTA Directory name (a routing-collective name) and credentials which enable it to read that Directory entry. All other information necessary for MHS-routing may be acquired from the Directory. As an implementation choice, this information may be retrieved at initialisation, or may be retrieved dynamically as required. For the purposes of exposition, the former approach is assumed. The caching of information acquired from the Directory may be necessary for efficient implementation, but is beyond the scope of this part of ISO/IEC 10021.

NOTE 1 – Implementors should be aware that out-of-date cached information can lead to routing loops and sub-optimal route choices. While procedures are specified to resolve routing loops, cache refreshing should also be performed at regular intervals.

A routing-MTA Directory entry indicates the one or more OR-address-subtrees which the routing-MTA's administrator has chosen to reflect the routing policy of this MTA, and also indicates the Directory name of the MHS Message Transfer Agent entry for this MTA (see A.1.3 of ITU-T Rec. X.402 | ISO/IEC 10021-2).

In order to fulfil its function, a routing-MTA must acquaint itself with knowledge of the other routing-collectives in its routing-collective-subtree. The minimum knowledge required for a routing-MTA to achieve this is discovered by collecting information from a specific subset of the routing-collectives in its routing-collective-subtree. These **key-routing-collectives** are defined as follows:

— the siblings of the routing-collective;

— the siblings of each of the routing-collective's superior routing-collectives.

NOTE 2 – While knowledge of a routing-collective's key-routing-collectives is required for the operation of MHS-routing, this information is not recorded explicitly in any Directory entry since it differs for every routing-collective, and may readily be discovered by inspection.

The reason for this definition of key-routing-collectives may be understood as follows. By the definition of routing-collective, a routing-MTA must be able to route a message to any of the other MTAs in each routing-collective of which it is a member; the practical effect is that the routing-MTA must be able to act on routing information that instructs it to transfer a message to one of the routing-collectives in the same routing-collective-subtree. Hence a simplistic definition of key-routing-collectives would include all of the routing-collectives in this tree. However, this gives more information than is strictly required. Plainly, the routing-MTA does not need to consider itself as a destination; similarly, an instruction to transfer the message to one of the routing-MTA's superior routing-collectives is meaningless (since the message has already arrived in that routing-collective), so these are not considered as key-routing-collectives. One of the purposes of having more than one level of routing-collective hierarchy is that a routing-MTA does not need to be aware of the internal structure of routing-collectives in distant parts of the routing-collective-subtree. Hence if one routing-collective is considered to be a key-routing-collective, its subordinates do not need to be. Excluding these unnecessary elements leaves the definition which has been adopted.

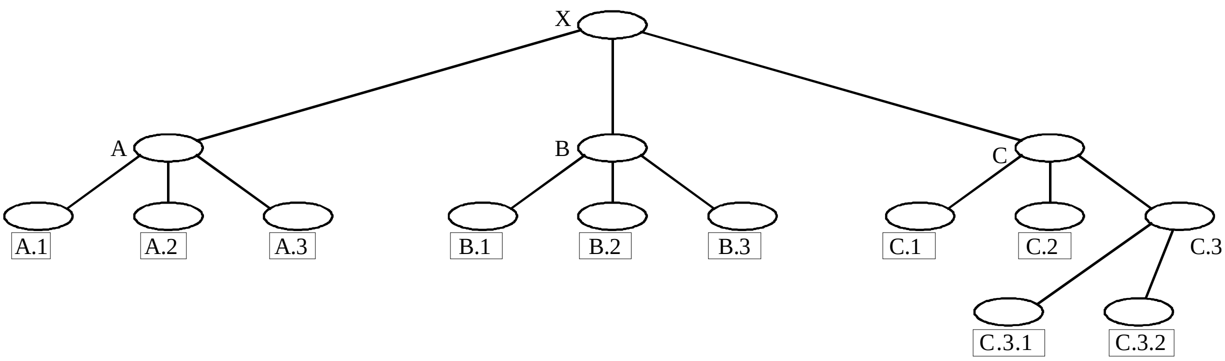

Figure 1 - Example of a routing-collective-subtree

Figure 1 illustrates the hierarchical relationships in a routing-collective-subtree. The routing-collective at the base of the subtree, X, has three immediate subordinate routing-collectives, A, B, and C. In turn, A, B, and C each contain three immediate subordinate routing-collectives. In the case of A and B, these routing-collectives are also routing-MTAs. In the case of C, C.1 and C.2 are routing-MTAs, while C.3 is a routing-collective containing two routing-MTAs. The names of the routing-MTAs are enclosed in boxes. For the routing-MTA B.3, the key-routing-collectives are B.1, B.2, A, and C.

6.2.3 Connection-group

A **connection-group** is a group of connections over which messages may be directly exchanged between members of a set of MTAs, using a specific MHS transfer protocol over a common network. It therefore represents the topology of the MHS, i.e., how the MTAs are physically interconnected. All MTAs in a connection-group have a network technology and a message transfer Application Context in common, and have administrative approval to interwork.

A connection-group is represented by specifying its member MTAs. A connection may be created between every pair of these MTAs. The Directory entry that represents a connection-group may also indicate the specific message transfer protocol used.

The simplest type of connection-group has just two members: this represents a conventional pair of MTAs that have been configured to communicate with one another. Larger sets may be constructed for convenience. If a number of MTAs are connected to a LAN or WAN, such that connections are permitted between any pair of them, they may all be assigned to a single connection-group, rather than to multiple groups that represent all the individual connections possible.

Connection-groups are of two types, enumerated and unenumerated. An **enumerated connection-group** is identified by Directory name, and has an associated list of member MTAs, as described above. An **unenumerated connection-group** is also identified by Directory name, and is typically associated with a network infrastructure such as a public or private wide-area network (e.g. X.25 or IP). Its membership is defined by self-declaration. The administrator who makes a local decision to configure an MTA into such a group is implicitly declaring that the MTA will accept connections from any other MTA on the same network, without prior agreement between MTA administrators.

Each routing-collective classifies the connection-groups available to it as follows:

a) an **entry-connection-group** is one that may be used to transfer messages into the routing-collective.

b) an **exit-connection-group** is one that may be used to transfer messages out of the routing-collective.

   Two types of exit-connection-groups are distinguished: a **transit-exit-connection-group** is one that may be used to transfer a message out of the routing-collective, regardless of origin; a **local-exit-connection-group** is one that may be used to transfer a message out of the routing-collective, but is available only for messages originated, or redirected, or DL-expanded within the routing-collective.

A given connection-group may be classified as both an entry- and exit-connection-group.

An **indirect-exit-connection-group** is an exit-connection-group which is not directly available to a routing-collective, but is available through one of its key-routing-collectives.

The connection-groups available to a routing-collective which is not a routing-MTA (i.e. does not occupy a leaf node in a routing-collective-subtree), correspond to connection-groups available to one or more of its subordinate routing-collectives. However, they are not necessarily the complete union of all such subordinate routing-collectives, since, as a matter of policy, a subordinate routing-collective may confine access to one of its connection-groups to itself and its subordinates.

Optionally, a security-context may be defined for a connection-group to govern the interactions permitted between members of that connection-group.

Figure 2 gives an alternative representation of the routing-collectives illustrated in Figure 1, and shows their connection-groups. The connection-groups CG1, 2, 4 and 5 are shown by indicating every individual connection within routing-collective X. Additional members of CG3 and CG6 exist within neighbouring routing-collectives.

6.2.4 OR-address-subtree

An **OR-address-subtree** is a Directory subtree which models a part of the OR-address name-space and associates routing-advice (see 6.2.5) with names within that name-space.

Directory attributes and object classes are defined corresponding to each element of the OR-address, such that a Directory name may be constructed by treating each OR-address element as an RDN, and assembling these RDNs in a prescribed order to form the Directory name. If entries were created in the Directory corresponding to all of the valid OR-addresses, this would create a subtree of the DIT corresponding to the whole OR-address name-space.

Since it is infeasible to place all routing information in a single Directory subtree (this would assume a fully interconnected Directory, a fully interconnected MHS, and that all users of Directory-based routing are prepared to reveal their complete addressing information to the world), it is usual to prefix the RDNs derived from OR-address elements with some arbitrary Directory name belonging to the administrator preparing the information. This allows multiple subtrees to be created in different parts of the global DIT, potentially containing information for every valid OR-address. A subtree may be incomplete, in that two or more subtrees are required to model a complete branch of the OR-address name-space.

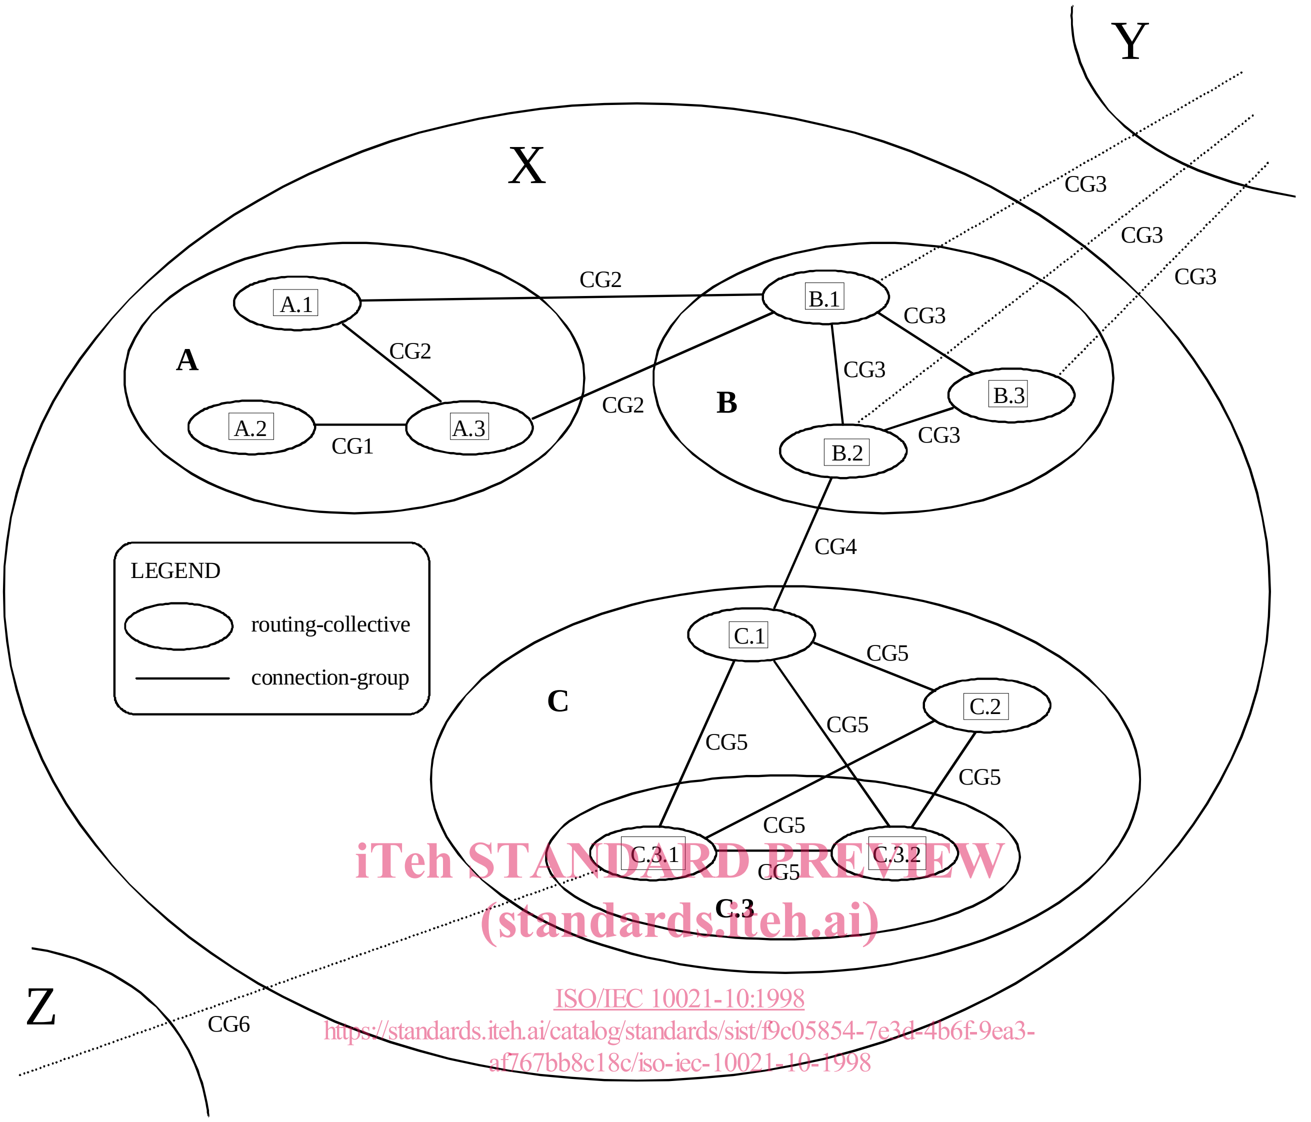

Figure 2 - Example of routing-collectives and their connection-groups

OR-address-subtrees containing different items of routing-advice will be required in various parts of the MHS for the same set of OR-addresses. This is due to the existence of firewalls, sparse connectivity, and the nature of devolved management. Different parts of the MHS will require different location-specific guidance to reach the same OR-address.

An administrator creates OR-address-subtrees to indicate the required routing behaviour for each possible OR-address. The entry corresponding to an OR-address will identify a routing-collective towards which the message should be transferred. In the simplest case, the routing-collective is the delivering-MTA for that recipient. Alternatively, the routing-collective might be larger, and contain the delivering-MTA as a subordinate routing-collective. In this case the message may be transferred to an arbitrary MTA in the routing-collective where more detailed information will be available to identify the delivering-MTA to which it should finally be transferred. Another possibility is that the indicated routing-collective is a relaying MTA that the administrator knows is able to handle all addresses in a particular domain.

Certain OR-addresses have an arbitrary internal structure, and cannot be fully represented in OR-address-subtrees (for example, one containing a domain defined attribute specifying an address in a foreign messaging environment). To represent a branch of the subtree that models OR-addresses of this type, an *expression-matches* attribute may be placed in the subtree to allow for algorithmic matching of these OR-addresses.

As it is impracticable to create entries in each OR-address-subtree for every possible OR-address, an OR-address-subtree may be incomplete in two ways: parts of the OR-address name-space may simply be absent from the tree (e.g. a tree which contains only local users), or whole branches of the tree may be truncated. Portions of the tree that are incomplete are marked as such, so that an MTA reading information from the tree is able to distinguish between an entry which is absent because the OR-address is invalid (causing the MTA to declare the message undeliverable), and an entry which is absent because the administrator has not provided the information (causing the MTA to look elsewhere for more information).

Portions of an OR-address-subtree may be truncated where the behaviour requested of the MTA is the same at all subordinate vertices of a vertex. For example, if all addresses within a particular Organizational Unit require transfer to the same MTA, it is not necessary to create entries for each OR-address in that Organizational Unit. Instead, the routing information is placed at the vertex in the tree corresponding to that Organizational Unit. The procedures used by the MTA for reading information from OR-address-subtrees detect when the desired entry cannot be read because the tree has been truncated, and cause the MTA to consult the leaf node of the truncated tree for routing advice.

It is entirely practicable to have just one OR-address-subtree for a PRMD (or indeed for a number of co-operating PRMDs), even if the PRMD has a nesting of routing-collectives within it. However, there will often be a requirement for multiple OR-address-subtrees, for example, where administrators wish to conceal information about OR-addresses, or for devolved management of OR-address-subtrees, or for private subtrees that identify locally available relaying and ADMD services.

6.2.5 Routing-advice

The **routing-advice** present in an entry in an OR-address-subtree provides advice to the routing-MTA that assists it in arriving at a routing decision for the OR-address corresponding to the entry. It specifies one of the following actions:

a) Transfer the message to the target routing-collective indicated. If this identifies the routing-MTA itself, then local delivery is attempted.

b) Generate a non-delivery report with the reason and diagnostic codes indicated.

c) Redirect the message to a preferred address of the MTS-user. The new OR-address is indicated in the routing-advice.

d) Perform DL-expansion of the OR-address; if this is not possible, transfer the message to the target routing-collective indicated, which is known to be capable of performing the required DL-expansion.

e) Place the message in an inner-envelope content-type, using the information supplied to construct the outer-envelope (see 8.2.1.1.1.34 of ITU-T Rec. X.411 | ISO/IEC 10021-4).

6.2.6 Local use tables

The only configuration information required by a routing-MTA is its routing-collective Directory name, information needed to access the Directory (e.g., its Directory credentials and the address of a local DSA), and knowledge of how to map its exit-connection-groups to appropriate protocol stacks and physical interfaces. On initialization, the routing-MTA reads the entry identified by its routing-collective Directory name, and follows a procedure that provides it with knowledge of its complete routing-collective-subtree, and thereby enables it to participate in MHS-routing (see 9.1.5). This information is recorded by the routing-MTA in its **local-use-tables**, as follows:

a) The Directory name of this MHS Message Transfer Agent.

b) The names of one or more OR-address-subtrees configured for this routing-MTA.

c) A list of the key-routing-collectives of this routing-MTA, each accompanied by the names of one or more next-MTAs for that key-routing-collective.

d) The names of the routing-MTA's transit- and local-exit-connection-groups.

e) The names of the routing-MTA's indirect-exit-connection-groups (i.e. those exit-connection-groups available through its key-routing-collectives), each accompanied by the names of one or more next-MTAs for that connection-group.

f) For each of the routing-MTA's entry-connection-groups that is of type enumerated, authentication information required to identify each possible calling MTA.

6.3 Routing decision overview

When routing a message, the routing-MTA performs the following actions to arrive at a routing decision for each recipient OR-address in the message:

a) The routing-MTA transforms the OR-address into a Directory name, and attempts to read the corresponding entry from the first configured OR-address-subtree. If the read is successful, the routing-advice associated with the entry is obeyed. If the read fails with a Name Error, and the lowest entry that was matched contains routing-advice then that advice is followed. If the lowest matching entry does not contain routing-advice and the next level of the subtree is marked as incomplete, then there may be no information available in this subtree, and the routing-MTA repeats the procedure using the next OR-address-subtree (typically, a routing-MTA is configured with several). However, if the next level of the subtree is marked as complete, then the OR-address is invalid and a non-delivery report is generated.

b) If routing-advice is obtained, this may identify a target routing-collective; alternatively, it may indicate that the OR-address is invalid, or that redirection of the message is required, or that DL-expansion is required.

c) If a target routing-collective is identified, the routing-MTA checks its local-use-tables (generated when the routing-MTA was initialized) to determine whether this is a key-routing-collective. If so, the next-MTA for that routing-collective is known, and the message is transferred to it.