

First edition
1999-12-15

**Information technology — Message
Handling Systems (MHS): MHS Routing —
Guide for messaging systems managers**

*Technologies de l'information — Systèmes de messagerie (MHS): Routage
MHS — Guide pour responsables de systèmes de messagerie*

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC TR 10021-11:1999

Reference number
ISO/IEC TR 10021-11:1999(E)



PDF disclaimer

This PDF file may contain embedded typefaces. In accordance with Adobe's licensing policy, this file may be printed or viewed but shall not be edited unless the typefaces which are embedded are licensed to and installed on the computer performing the editing. In downloading this file, parties accept therein the responsibility of not infringing Adobe's licensing policy. The ISO Central Secretariat accepts no liability in this area.

Adobe is a trademark of Adobe Systems Incorporated.

Details of the software products used to create this PDF file can be found in the General Info relative to the file; the PDF-creation parameters were optimized for printing. Every care has been taken to ensure that the file is suitable for use by ISO member bodies. In the unlikely event that a problem relating to it is found, please inform the Central Secretariat at the address given below.

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC TR 10021-11:1999

© ISO/IEC 1999

All rights reserved. Unless otherwise specified, no part of this publication may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying and microfilm, without permission in writing from either ISO at the address below or ISO's member body in the country of the requester.

ISO copyright office
Case postale 56 • CH-1211 Geneva 20
Tel. + 41 22 749 01 11
Fax + 41 22 749 09 47
E-mail copyright@iso.ch
Web www.iso.ch

Published by ISO in 2000

Printed in Switzerland

CONTENTS

	<i>Page</i>
1 Scope	1
2 Normative References	1
2.1 Identical Recommendations International Standards.....	1
2.2 Paired Recommendations International Standards equivalent in technical	1
3 Definitions	1
3.1 MHS definitions	2
3.2 ASN.1 definitions.....	2
3.3 Directory definitions.....	2
3.4 Presentation Service definitions	2
3.5 MHS-routing definitions	2
3.6 MHS Routing Methodology Definitions	3
4 Abbreviations	3
5 Overview	3
5.1 The role of MHS Routing.....	3
5.2 Administrative Roles.....	4
5.3 The Role of the MHS Routing Standard	5
5.4 The advantages of using MHS Routing	6
6 MHS Routing Concepts.....	6
6.1 Introduction	6
6.2 MHS	7
6.3 The Message Routing Problem	7
6.4 A Directory Solution to Message Routing.....	8
6.5 General Directory Services and Functions	8
6.6 MHS-Routing use of the Directory	9
6.7 Scenario.....	9
6.8 Routing Collectives and Connection Groups	10
6.9 Routing Collective Directory Representation.....	12
6.10 Connection Groups.....	12
6.11 Connection Group Directory Representation	13
6.12 OR-address Plan.....	13
6.13 OR-address plan directory representation in OR-address subtrees	13
6.14 MTA definitions	14
6.15 The MTA's Message Routing Process	14
6.16 MHS Routing Administrative Roles	15
6.17 MHS Routing Administrative Tasks	15
6.17.1 Organizational Requirements Analysis	16
6.17.2 Design Tasks	16
6.17.3 Configuring the Directory	17
6.17.4 MTA Configuration.....	20
6.17.5 MTA Initialization.....	20
6.18 Prerequisites for MHS Routing	21
6.19 Organizational Requirements Analysis	21
6.20 The user community	22
6.21 Cultural requirements and constraints	22
6.22 Organizational messaging administrative structures	22
6.23 Existing, non-standard and other Messaging Systems	22
6.24 Available Communications Infrastructure.....	23
6.25 Messaging Traffic Patterns and Volumes.....	23

	<i>Page</i>	
6.26	Security requirements.....	23
6.27	Specific message routing requirements or policies	23
6.28	Alternative routing capability.....	23
6.29	Directory systems availability	23
6.30	Distribution lists	24
6.31	MTS Topology Design.....	24
6.32	General Guidance.....	24
6.33	Input information.....	24
6.34	Results of the Topology design process	25
6.35	Identifying Connection Groups for MTAs	26
6.36	Results	26
7	OR-address Plan Design.....	27
7.1	General	27
7.2	Information Input	27
7.3	OR-address plan specification.....	27
7.4	Special Cases.....	28
7.5	OR-address Registration Authority Roles	28
7.6	Results of the OR-address plan design.....	28
8	Routing Collective Design and Configuration	29
8.1	Introduction	29
8.2	Directory Information Base Preparation	30
8.3	Directory Information Base Configuration	31
8.4	Connection Groups.....	32
8.4.1	Defining Connection Groups	32
8.4.2	Adding a Connection Group to a Routing Collective	33
8.4.3	Deleting a Connection Group Definition.....	33
8.4.4	Routing Collectives	34
8.4.5	Identifying Routing Collectives	34
8.4.6	Routing Collective Definition	35
8.4.7	Creating the top level Routing Collective	35
8.4.8	Adding a Subordinate Routing Collective.....	35
8.4.9	Deleting a Routing Collective	35
8.4.10	Adding Proxy routing collectives.....	35
8.5	Configuring Routes for MTAs in OR-address subtrees	36
8.5.1	OR-address subtree types	36
8.5.2	OR-address subtrees Model.....	36
8.5.3	Routing Information.....	37
8.5.4	Specifying OR-address subtree bases.....	37
8.5.5	Building OR-address subtrees	38
8.5.6	Establishing External Routes to destinations outside the routing collective	40
8.5.7	Distributing Access to External routes through a routing collective	41
8.5.8	Establishing Routes to non-Routing MTAs or proprietary messaging systems	42
8.5.9	Default Routes.....	42
8.6	OR-address subtree Entry Routing Information Configuration	42
8.6.1	OR-address attribute registration.....	42
8.6.2	The Target Routing Collective Instruction.....	43
8.6.3	MHS User Instruction	43
8.6.4	Aliases	44
8.6.5	Aliasing Techniques	44
8.6.6	Alias OR-address Instructions for Personal Names	44
8.6.7	The Alias Redirection Instruction	44
8.6.8	The Non-delivery Instruction	45
8.6.9	The Distribution List Instruction.....	46
8.6.10	The Recipient MD Assigned Alternate Recipient Instruction.....	46
8.6.11	The Double Enveloping Instruction	47
8.6.12	The Expression Matches Instruction	48
8.6.13	Truncating an OR-address subtree	48

	<i>Page</i>
8.7 Organizing an MTA's OR-address subtrees.....	49
8.7.1 The MTA's OR-address subtree sequence.....	49
8.8 Publishing Routing Capabilities.....	49
8.9 Configuring an MTA.....	49
8.9.1 Routing MTA Entry.....	50
8.9.2 mHSMMessageTransferAgent Entry.....	50
8.10 MTA Initialization.....	50
8.11 MTA Cache information.....	51
9 Directory Information Base Guide.....	51
9.1 Directory Information Structure.....	51
9.2 Routing collective subtree components.....	51
9.2.1 The Routing Collective Object Class.....	51
9.2.2 Routing MTA Object Class.....	51
9.3 Connection Group.....	52
9.3.1 The Connection Group Object Class.....	52
9.4 MTA Components.....	52
9.4.1 MTA Information Object Class.....	52
9.5 OR-address subtree Components.....	53
9.5.1 The OR-address Element object class.....	53
10 Provision of the MHS Routing Directory Service.....	53
Annex A – Scenarios.....	55
A.1 Single MTA MD connected only to an ADMD.....	55
A.2 A small MD under a single management.....	55
A.3 Large MD with autonomous management.....	55
A.4 The open access connection group case.....	56
A.5 Collection of MDs.....	56
A.6 Secret OR-addresses.....	56
Annex B – MHS Routing DUA Specification.....	57

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.

International Standards are drafted in accordance with the rules given in the ISO/IEC Directives, Part 3.

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.

In exceptional circumstances, 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), it may decide by a simple majority vote of its participating members to publish a Technical Report. A Technical Report is entirely informative in nature and does not have to be reviewed until the data it provides are considered to be no longer valid or useful.

In exceptional circumstances, 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), it may decide by a simple majority vote of its participating members to publish a Technical Report. A Technical Report is entirely informative in nature and does not have to be reviewed until the data it provides are considered to be no longer valid or useful.

Attention is drawn to the possibility that some of the elements of this part of ISO/IEC TR 10021 may be the subject of patent rights. ISO and IEC shall not be held responsible for identifying any or all such patent rights.

ISO/IEC TR 10021-11 was prepared by Joint Technical Committee ISO/IEC JTC 1, *Information technology*, Subcommittee SC 6, *Telecommunications and information exchange between systems*, in collaboration with ITU-T. The identical text is published as ITU-T Recommendation X.404.

ISO/IEC TR 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*
- *Part 11: MHS Routing — Guide for messaging systems managers*

Annexes A and B of this part of ISO/IEC 10021 are for information only.

Introduction

This Recommendation | Technical Report is one of a set of Recommendations | number of parts of ISO/IEC 10021 defining Message Handling in a distributed open systems environment.

ITU-T Rec. X.412 | ISO/IEC 10021-10 defines a method for routing messages through the Message Handling System (MHS). This Recommendation | Technical Report provides guidance for Configuring MTS Routing using the Directory, and suggests the characteristics of a Directory User Agent for managing that process. It allows OR-address plans, MTA interconnection topology and the management structures applied to MHS to be dealt with independently of each other whilst remaining within a co-ordinated framework.

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC TR 10021-11:1999

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC TR 10021-11:1999

TECHNICAL REPORT

ITU-T RECOMMENDATION

INFORMATION TECHNOLOGY – MESSAGE HANDLING SYSTEMS (MHS): MHS ROUTING – GUIDE FOR MESSAGING SYSTEM MANAGERS

1 Scope

This Recommendation | Technical Report specifies the means by which the administrator of various aspects of an MHS system may configure information into the directory for MTAs to use to determine the routing of messages.

ITU-T Rec. X.412 | ISO/IEC 10021-10 provides a set of directory structures that may be configured in many different ways to support a particular MHS routing strategy. In order to illustrate the use of these directory structures, this document contains advice on how an MHS Administrator might organize the configuration of directory trees and entries in the directory. In particular, it contains suggestions on the following:

- The types, construction and location of different OR-address subtrees that may be needed;
- The location of routing collective and MTA entries in the directory.

Other ways of using the routing capabilities specified in ITU-T Rec. X.412 | ISO/IEC 10021-10 are also valid.

Other Recommendations | International Standards 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. ITU-T Rec. X.412 | ISO/IEC 10021-10 defines MHS Routing using the directory.

2 Normative References

The following Recommendations and International Standards contain provisions which, through reference in this text, constitute provisions of this Recommendation | Technical Report. At the time of publication, the editions indicated were valid. All Recommendations and Standards are subject to revision, and parties to agreements based on this Recommendation | Technical Report are encouraged to investigate the possibility of applying the most recent edition of the Recommendations and Standards 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 Recommendations | International Standards

- ITU-T Recommendation X.216 (1994) | ISO/IEC 8822:1994, *Information technology – Open Systems Interconnection – Presentation service definition.*
- ITU-T Recommendation X.402 (1995) | ISO/IEC 10021-2:1996, *Information technology – Message Handling Systems (MHS): Overall architecture.*
- ITU-T Recommendation X.412 (1999) | ISO/IEC 10021-10:1998, *Information technology – Message Handling Systems (MHS): MHS routing.*
- ITU-T Recommendation X.500 (1997) | ISO/IEC 9594-1:1998, *Information technology – Open Systems Interconnection – The Directory: Overview of concepts, models and services.*

2.2 Paired Recommendations | International Standards equivalent in technical content

- CCITT Recommendation X.208 (1998), *Specification of Abstract Syntax Notation One (ASN.1).*
ISO/IEC 8824:1990, *Information technology – Open Systems Interconnection – Specification of Abstract Syntax Notation One (ASN.1).*

- ITU-T Recommendation F.400/X.400 (1999), *Message handling services: Message handling system and service overview*.
- ISO/IEC 10021-1:1990, *Information technology – Text Communication – Message Oriented Text Interchange Systems (MOTIS) – Part 1: System and service overview*.

3 Definitions

For the purposes of this Recommendation | Technical Report, the following definitions apply.

3.1 MHS definitions

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

- OR-address;
- MTA;
- MTS;
- MHS;
- Message Store;
- User Agent;
- P7;
- P3;
- MD;
- ADMD;
- PRMD;
- Security Context.

3.2 ASN.1 definitions

The following term is formally defined in ITU-T Rec. X.208 | ISO/IEC 8824:

- ASN.1.

3.3 Directory definitions

The following terms are formally defined in ITU-T Rec. X.500 | ISO/IEC 9594:

- Directory Name;
- Relative Distinguished Name.

3.4 Presentation Service definitions

The following term is formally defined in ITU-T Rec. X.216 | ISO/IEC 8822:

- Presentation Service Access Point.

3.5 MHS-routing definitions

The following terms are formally defined in ITU-T Rec. X.412 | ISO/IEC 10021-10:

- Connection Group;
- Enumerated Connection Group;
- Next Level Complete;
- Routing collective;
- Routing-MTA;
- Target MTA;
- Un-enumerated Connection Group.

3.6 MHS Routing Methodology Definitions

The following terms are defined in clauses 6 to 10 of this Recommendation | Technical Report:

- 3.6.1 adjacent MTA:** An MTA that is directly connected (i.e. through some connection group) to the current MTA.
- 3.6.2 administrator:** A person or role which manages a particular routing collective in the MHS.
- 3.6.3 current MTA:** The MTA taking a routing decision for a message.
- 3.6.4 exit MTA:** A routing MTA within the routing collective which has access to connection groups allowing it to transfer messages to MTAs outside a routing collective. In an extreme case, each MTA in a routing collective might be an Exit MTA.
- 3.6.5 external OR-address subtree:** An OR-address subtree which holds routing information to parts of the MTS lying outside the routing collective under construction.
- 3.6.6 external Route:** A route from a routing collective Exit MTA to another MTA outside the routing collective.
- 3.6.7 internal Route:** A route between two MTAs within a routing collective.
- 3.6.8 mailbox:** term used to indicate the delivery point for messages located by an OR-address. This may be a P7 accessed message store, a P3 accessed user agent or a proprietary protocol accessed user process. The distinction between these variants is irrelevant to MHS routing.
- 3.6.9 OR-address plan:** A plan of OR-address attribute types used to identify an organization's departments, divisions and users of MHS. An OR-address plan is specified by the Organisational and MHS administrators to select the particular OR-address forms and attributes for use within the organisation from all those possibilities that are specified in the MHS base standards. Organizational Administrator are persons or roles which manage non-MHS aspects of an organization, but who place requirements on an MHS system.
- 3.6.10 reference OR-address subtree:** an OR-address subtree that contains a routing collective's internal routing and message delivery information.
- 3.6.11 registration Authority:** An administrative role which ensures that OR-addresses are unambiguous, i.e. that each OR-address is allocated to one and only one user.
- 3.6.12 routing Information:** Information held in OR-address subtrees which instructs an MTA on how to process a message for a particular OR-address.
- NOTE – This is more general than the ASN.1 construct 'routing advice', since it is oriented to the discrete actions that an administrator will take, and includes Recipient MD Assigned Alternate Recipients etc.
- 3.6.13 top level routing collective:** A routing collective which does not belong to a superior routing collective.

4 Abbreviations

For the purposes of this Recommendation | Technical Report, the following abbreviations apply.

ISDN	Integrated Services Digital Network
LAN	Local Area Network
PSDN	Packet Switched Digital Network
PSTN	Public Switched Telephone Network

5 Overview

5.1 The role of MHS Routing

Message Handling Systems exchange messages between users on a store-and-forward basis. A message submitted by one user (the originator) will be transferred through one or more Message Transfer Agents (MTAs) in the Message-Transfer-System (MTS) and delivered to one or more other users (the recipients). The sequence of MTAs through which a message is transferred on its way from originator to recipient is the message's route. The originator does not specify

the route but identifies the recipient in the MTS by means of an unambiguous Originator/Recipient Address (OR-address) or directory name which is translated by the directory into the recipient's OR-Address. A recipient's OR-Address indirectly specifies which MTA the recipient is attached to. Each user is supported by a single MTA that is responsible for delivering messages addressed to the user. Each MTA that handles a message on its way towards its destination uses the recipient's OR-address on the message's envelope to select the most appropriate subsequent route. This process eventually leads to the delivering MTA.

To achieve store and forward messaging, each MTA is configured with routing information indicating the OR-addresses for which it has delivery responsibility and also the routes through adjacent MTAs which should be taken towards all other OR-addresses. In general, each MTA requires different routing information to reflect its location and connectivity within the MTS with respect to MTAs supporting other OR-addresses.

5.2 Administrative Roles

MTAs are configured with routing information by MHS administrators. The information that (possibly different) administrators supply to their MTAs should be co-ordinated to ensure that each message is correctly and efficiently routed towards its recipients, and that routing conflicts and loops do not occur.

MHS routing is influenced by a number of different organizational roles:

- an organizational administrator who is primarily concerned with the operational aspects of the organisation, and regards MHS as a resource. Organisational administrators specify requirements of MHS, but do not get involved in the realisation of MHS;
- an MHS Administrator who is directly responsible for all aspects of the installation and operation of the MHS, including the connectivity between MTAs etc.;
- an administrative role to ensure the allocation of unambiguous OR-addresses to MHS users. This role is formally identified as a 'Registration Authority'. However its realisation, and the way it is administered will vary from organization to organization, and it will also vary according to the type of OR-address attribute being administered and the type of users the attributes are being registered for.

The MHS requirements are determined primarily by organizational administrators, and the MHS routing strategy supporting those requirements is designed and configured by MHS administrators. In some cases, the MHS administrator's role may be further subdivided e.g. into those who deal with creating and managing mailboxes and those who are primarily involved in managing MTAs and their interconnectivity.

In addition, a registration authority is also assumed to be responsible for registering the OR-addresses assigned to users, ensuring that they are unambiguous and that the OR-addresses conform to the OR-address plan. Each registration authority will administer values for some OR-address attributes within the scope of one or more well defined sets of OR-addresses. Some registration authorities will act at higher levels in the global OR-address space (e.g. to register Country Names, PRMD Names, ADMD Names and organization Names). Others will act at lower levels, (e.g. to register organization Names, organizational Unit Names, Personal Names etc. within ADMDs, PRMDs, organizations etc.). Registration of Country names, MD names and Organisation names is performed by a hierarchy of formal registration authorities. However, the way that the remaining OR-address attributes are administered within an organization may be less formal, and it will differ dependent on the type of organization.

The sole collective technical requirement of registration authorities as far as MHS routing is concerned is that all OR-addresses should be unambiguous – i.e. no two MHS users should be allocated (or granted) the same OR-address.

In some organizations, one or more of the above roles may be carried out by the same department or person, however they will often be dispersed and, for the purposes of this Recommendation | Technical Report the functions they carry out are regarded as separate and independent of each other.

Organizational administrators are concerned with the day to day business of the organization and regard the MHS as a facility to support the organization. They are not directly involved in MHS administration. They determine the internal structure of the organization into departments, and the distribution of the organization and the departments over different geographically dispersed sites. In doing this, they determine the geographical site at which each MHS user is located. Their primary concerns with MHS routing are:

- that the MHS supports an OR-address plan which reflects each user's departmental or site location¹⁾;
- that the OR-address plan makes it easy to guess a user's OR-address; and
- that OR-addresses are stable and do not need to be changed when users are relocated or connected to a different MTA.

It is essential that the OR-address plan is determined by these factors alone, and that these requirements are not compromised by any MHS configuration choices or limitations imposed by MHS products or services. The organizational administrator's input into the MHS routing design is therefore a specification of an OR-address plan suited to the organization, and, for each user, a specification of the geographical and departmental locations and the user's OR-address.

MHS administrators are concerned with supporting the messaging requirements specified by organizational administrators. There are two independent aspects of the MHS administrator's role:

- To develop and maintain the interconnectivity of MTAs together with a message routing strategy which supports the organization's OR-address plan. These must take into consideration the geographical distribution of the organization and the available MHS systems and data communications links connecting the organization's different sites;
- Configuring each user's MHS OR-addresses and mailbox at a particular location and with an OR-address specified by the organizational administrator.

It should be noted that an organization's operational internal structure, which should be the sole factor determining its OR-address plan, is often quite different to the organizational structure of the MHS administration, and it may be different again to the topology of MTA interconnections. It is therefore important that the three following aspects of MHS routing design should be maintained independently of each other, and that:

- the OR-address plan can be constructed to reflect the requirements of the organizational administrator, and to provide as short an OR-address as is possible whilst remaining intuitively 'guessable' by users. This aim should not be compromised by any aspects of the design or configuration of the MHS;
- an MHS administrative structure can be constructed to fulfil the requirements of the organizational administration, recognizing that it may be centralised or devolved, and may need to be quite different to the organization's structure as specified by the organizational administrator (e.g. it must allow for cases where individual organizational departments are distributed over different geographical locations and served by different MTAs);
- the OR-address plan remains independent of the topology of the MTS.

5.3 The Role of the MHS Routing Standard

In the absence of ITU-T Rec. X.412 | ISO/IEC 10021-10, different MTA products tended to adopt different approaches to routing and often, no tools were provided to support the specification of an overall routing strategy among groups of MTAs. This was particularly so in multi-vendor environments in which MTS designs were often compromised because the specification of the OR-address space and MTS topology and message routing could not be done independently of each other. In these cases, to achieve a workable routing strategy, each MTA had to be assigned the delivery responsibility for a complete OR-address space (e.g. by assigning delivery responsibility for all OR-addresses containing a particular OU Name to a particular MTA). This strategy creates an unfortunate binding between the OR-address space and the topology of the MTS, and had a number of disadvantages:

- 1) Distribution of an organizational department's users over several MTAs forced MHS administrators to introduce unnecessary extra OR-address attributes into the OR-address plan. These attributes were of no real significance within the organization but were necessary to identify the delivering MTA (e.g. by adding extra Organizational Unit attributes to distinguish between the different MTAs which supported a single organizational department, so that each MTA had delivery responsibility for all OR-addresses containing a particular value of the extra attribute);

¹⁾ Organizational administrators can choose whether to develop an OR-address plan reflecting the departmental structure or geographic distribution of the Organization, or they can choose a mixture of both.

- 2) Routing strategies often had to be constructed that concentrated routing knowledge in central MTAs, and required that all message traffic should pass through that MTA. This inhibited direct routing between the originator and recipient's MTAs, it introduced a single point of failure, and was the source of potential traffic overload situations and unnecessarily long delivery delay times;
- 3) When users had to be relocated to another MTA (e.g. because they were moved to another building or site or department), their OR-address had to be changed to reflect the fact that they were served by a different MTA;
- 4) It was sometimes not possible to determine whether an OR-address was valid, other than by passing the message through all MTAs in an area of the MTS until all possibilities had been tried.

MHS administrators were also confronted with a variety of different mechanisms to configure different MTAs, and many MTA products placed restrictions on the OR-address plans that they could specify.

5.4 The advantages of using MHS Routing

ITU-T Rec. X.412 | ISO/IEC 10021-10, and this Recommendation | Technical Report, specify MHS routing and routing techniques using the directory. The directory enables the information required for MTA routing decisions to be held in a standardized, distributed store which can be accessed by each MTA using a standardized access protocol. Using the directory in this way to support MHS routing has a number of advantages for the administration of ADMDs and PRMDs:

- a) It allows MHS and organizational administrators to allocate shorter OR-addresses which remain stable over longer time periods. This is particularly important where a department within an organization (e.g. identified by a particular organizational Unit OR-address attribute value) is geographically distributed, or where it is necessary to spread the processing load of a large department over two or more MTAs;
- b) It allows enhancement of MTS performance, since although it might be considered that use of the directory is an overhead, the method can improve the performance of MTS routing because of its ability to select more direct routes and reduce the number of MTAs which have to be traversed by messages;
- c) The MHS administrator can avoid single points of failure (i.e. by avoiding the need to configure star connectivity to centralised MTAs, or having to route in a 'hierarchic' fashion as dictated by the OR-addresses attributes);
- d) It is possible to construct alternate routes from any MTA to any destination in order to avoid congested areas or system failures;
- e) It allows MHS administrators to avoid routing strategies which lead to traffic congestion situations;
- f) It provides a uniform approach and a single method to ease the MHS administrative tasks of routing by providing a standard (abstract) interface to the routing controls of MTAs through a single point of access to the directory which may be remote from the MTAs being configured;
- g) It can be used for ADMDs and PRMDs to exchange routing information identifying which MTAs should be used as entry points into the MD and it will be of particular use where the single space country name or 'XX' country code are in use. Such information may assist other MDs to make optimum routing choices.

Because of these characteristics, MHS administrators using systems that conform to the MHS-Routing standards and this methodology can provide a better user service to the organization and MHS users.

6 MHS Routing Concepts

6.1 Introduction

This clause introduces MHS Routing administrators to the concepts of the MHS Routing Standard (ITU-T Rec. X.412 | ISO/IEC 10021-10). Clause 6 of ITU-T Rec. X.412 | ISO/IEC 10021-10 contains a more formal and detailed description. The following concepts are explained here:

- 1) MHS;
- 2) MHS Routing and the MHS Routing problem;
- 3) The Directory;

- 4) The rationale for a Directory based solution;
- 5) The Model of Directory Based MHS Routing;
- 6) Directory schema configuration for MHS Routing.

6.2 MHS

Message Handling Systems (MHS) 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 other users, the recipients. The MTS comprises a collection of Message Transfer Agents (MTAs). The MTAs are highly distributed, and connected directly or indirectly to each other in a networked fashion. A message will traverse one or more MTAs on its journey from its originator to its recipient.

MHS routing takes place in the OSI Application Layer (i.e. in MTAs), and is distinct from network layer routing (MHS routing deals with complete messages, the Network layer routes data streams that carry messages between different MTAs).

A more complete description of the MHS is provided in ITU-T Rec. F.400/X.400 | ISO/IEC 10021-1.

6.3 The Message Routing Problem

In MHS, the message originator does not specify a path through different MTAs to reach a recipient, but specifies the recipient's OR-address (or a directory name that is used to determine the OR-address from the directory). An OR-address consists of a set of OR-address attributes, each of which identifies a particular characteristic of the recipient, e.g. Country, Management Domain, Organization, Personal Name etc.

It is the responsibility of each MTA to determine the next MTA to which the message should be transferred to progress the message's journey to its recipient. Any given MTA is connected to a number of other MTAs, and an MTA routing a particular message must choose another MTA to which it will forward the message towards its destination. Some choices will be more efficient than others. The selection of the next MTA is based on the recipient's OR-address. Routing is therefore the process of selecting, given an OR-address, the MTA to which the message should next be transferred. The path taken between an originator and recipient may vary on different occasions, since there will in general be a number of possible paths through different MTAs between them, and factors such as congestion and availability may influence route selection dynamically.

MHS Routing presents a number of significant problems:

- MHS is envisaged as a global service and therefore has a very large OR-address space. It is not feasible to configure MTAs with direct routing knowledge for all possible OR-addresses because of the size of the information and the logistics of distributing OR-address update information;
- MTAs are often sparsely connected (i.e. they do not all directly connect to each other) because it is not feasible to provide a single underlying data communications network to connect all of them. In the absence of complete MTA interconnectivity, this has the implication that different MTAs need to have different routing information dependent on their location with respect to each recipient's delivering MTA and OR-address;
- Different MTAs require different levels of detail about different OR-addresses. An MTA that is responsible for a set of OR-addresses must have all the delivery information in detail. However, an MTA that is remote from the delivering MTA will only need to know the identity of the delivering MTA or some other MTA on a path towards the delivering MTA.

Because of these problems, MTA implementations have tended to simplify their routing strategy by binding the delivery responsibilities for complete OR-address spaces summarized by a particular OR-address attribute value to a particular MTA (e.g. where an MTA has delivery responsibility for a complete Organizational Unit). Whilst this approach may be efficient in terms of administration, it causes a number of problems e.g.:

- 1) Where the personnel of a particular organization's department (e.g. the sales force) are geographically widely distributed, MHS administrators may need to impose further unnecessary layers of Organizational Unit attributes (or some other attribute) in the OR-address plan to represent that geographic distribution, and to enable MTAs to identify the correct delivering MTA;
- 2) When user's move location (geographically), their OR-address may need to change to one served by the MTA at the new location;

- 3) A user's OR-address may be constrained by the MTA to which the user connects;
- 4) the distortion of the relationship between OR-addresses and delivering MTAs means that message traffic must often be routed through intermediate MTAs unnecessarily (e.g. where a star network has been implemented to overcome routing difficulties). This leads to traffic congestion, single points of failure and unnecessary throughput delays.

Because of this, many organizations have had difficulties in establishing an acceptable OR-address scheme and supporting it with an efficient routing strategy. There has always been a trade off between address scheme tailored to the organization and routing strategy imposed by MTA products. These difficulties have primarily been caused by the binding of particular OR-address spaces to specific MTAs, leading to inappropriate OR-address schemes being imposed on organizations, restrictions on mobility of OR-addresses, traffic bottlenecks and conflicts between management of the organization and management of messaging systems.

6.4 A Directory Solution to Message Routing

Use of ITU-T Rec. X.412 | ISO/IEC 10021-10, MHS Routing, avoids these problems by utilising standard directory systems in a special way. It specifies how directories are used by MHS administrators to configure routing information for groups of MTAs, and how individual MTAs are provided with access to routing information which is appropriate to their role and location in the MTS.

The standard directory technology has been used to support a directory schema specific to the MHS routing task. This technique can be used in a number of different ways to support different MHS routing objectives:

- to support a global routing strategy, in which case a global directory is required;
- to support a local routing strategy (e.g. within an MD), in which case a local directory is required, but a global directory is not.

Both of these are the subject of this Recommendation | Technical Report. MHS routing encompasses all global and local routing requirements.

6.5 General Directory Services and Functions

The Directory is specified in ITU-T Rec. X.500 | ISO/IEC 9594. Directories are traditionally used to provide a service for storing and retrieving a wide range of types of information to support global communications services such as Telephony, Fax, MHS and EDI in a similar fashion to paper based White Pages and Yellow Pages.

A Directory typically contains communications related information held in one or more Directory System Agents (DSAs) which users access using a Directory-User Agent (DUA). The DUA aids users to form directory updates and queries and returns the results to the user.

DSAs are generic and can be configured to support many different types of schemas reflecting the different requirements of different communications services, whereas DUAs are generally specific to a particular task.

The directory stores information concerning objects (e.g. people, machines in the real world that communicate with each other by telephone/fax/MHS etc.). The tree structure is referred to as the Directory Information Tree (DIT). Each node in the DIT locates a Directory Entry. The topmost entry in the DIT is referred to as "The Root"²⁾.

Each object's directory entry is identified in the directory by a Directory Name that is determined by its location in the tree structure. Each entry is allocated a name in the context of its superior entry. This name is known as a Relative Distinguished Name (RDN). Each of an entry's subordinate entries is allocated a different RDN in order to distinguish between the subordinates. The sequence of RDNs, starting from the Root, down to a particular entry is the entry's Directory Name. In this way, entries, and the corresponding real world objects they represent, are identified by a Directory Name which corresponds to the entry's location in the DIT. The DIB can be constructed to hold many different types of directory names forms to identify different types of objects (e.g. persons, groups, houses and buildings) and different entry content types to reflect those different objects.

Each Directory entry contains information about a single real world object in the form of a set of Directory Attributes. Each attribute holds a particular type of information about the object, e.g. a telephone number.

Directory service administrators may create and delete entries, and perform read, write and modify operations on entries identified by a directory name. Search operations can be applied to a whole DIT subtree that has been identified by its directory name to retrieve information from all entries with specified attribute values.

²⁾ In practice, the root will not actually exist, however, its subordinates will.

The Directory also provides a comprehensive set of access controls which can be applied to allow read, write and modify access to identified directory users to directory attributes, entries and whole directory subtrees. These can be used to prevent un-authorized read or write access.

Several Directory applications have already been developed, each specifies a directory schema, e.g. for:

- Telephony information;
- Fax Information;
- Security Information;
- OSI data communications applications, including EDI, Transaction Processing, FTAM and the Virtual Terminal.

However, the schema components provided in ITU-T Rec. X.500 | ISO/IEC 9594 can be extended and adapted or replaced with new ones for other tasks. MHS Routing specifies such a directory schema to support the particular task of MHS message routing. It consists of attributes, collections of attributes for routing entries, and directory name forms specific to the routing task.

6.6 MHS-Routing use of the Directory

MHS routing uses directory systems in a slightly different way to the global communications services use it. However, as far the MHS administrator is concerned, each Directory Name locates a particular directory entry containing MHS routing related information, and that entry contains a set of attributes of the real world object identified by the directory name.

A DUA customised to the particular task of MHS routing will considerably simplify the administrative task of constructing and maintaining directory information for MHS routing.

6.7 Scenario

In MHS Routing, the MHS administrator can use a Directory User Agent (DUA) or some other mechanism (such as exporting/importing or shadowing) to configure MHS routing information into the directory for the set of MTAs administered by the routing collective. Each MTA has a local embedded DUA to read the directory, and possibly builds a cache of directory information which is most relevant to its own operation. Clause 10 of this Recommendation | Technical Report provides details of how an MHS administrator should identify and configure the directory information.




Figure 1 – MHS Routing Information flows

Figure 1 illustrates MHS routing in terms of the flow of information controlling message routing. OR-addresses, MTS topology and configuration, and the routing information for specified MTAs are prepared by the organizational administrator, the MHS administrator and the Registration Authority. The directory entries are read by MTAs during the message routing process to determine whether the MTA can deliver the message or to determine the next MTA to which the message should be transferred if the message is to be relayed.

The MHS administrator's task is to configure the MHS Routing Directory Information prior to initializing MTAs and putting them into service. It should be noted that different types of information may be prepared and input to the directory by different administrators; (e.g. MHS system administrative structures and MTA interconnectivity information may be prepared by one administrator, and user mailbox creation and deletion may be carried out by a different administrator). The routing standard and methodology do not specify how user mailboxes are administered.

Each MTA is configured to access this information to configure its internal routing knowledge and to be able to access routing information for each OR-address. Each MTA uses this directory information to route messages to adjacent MTAs which are 'nearer'³⁾ the message's destination.

The MHS Routing Model consists of the following elements:

- Routing Collectives, which provide a hierarchic model of the MTS administrative structures. Each routing collective is a portion of the MTS consisting of a group of MTAs managed by an MHS administrator;
- Connection Groups, which model the topology of interconnection of MTAs in the MTS. They represent connectivity between two or more MTAs, and hence the topology of the MTS;
- The OR-address plan;
- OR-address subtrees;
- Routing Information, held in OR-address subtree entries, and which contain instructions to MTAs. Each instruction may be to either deliver, Non-Deliver, DL-expand, Redirect or to Route the message through another MTA;
- MTA definitions, which hold basic configuration information for each MTA from which it MTA can 'bootstrap' its routing information from the directory. Each MTA is considered to be a routing collective in its own right.

The following subclauses elaborate on these concepts and outline how the necessary information is configured in the directory. Clause 7 onwards provide details on each of the tasks which administrators must perform with respect to each of these concepts.

6.8 Routing Collectives and Connection Groups

Figure 2 illustrates a routing collective 'X' that is internally divided into a number of subordinate routing collectives (A, B and C). These, in turn, are divided into further subordinate routing collectives (A.1, A.2, A.3, B.1, B.2, B.3, and C.1, C.2, C.3 etc.). Each of the lowest level routing collectives is supported by a single MTA, which assumes the name of the lowest level routing collective it supports. So, each routing collective (X, A, B, C, A.1, A.2, A.3, B.1, B.2, B.3, C.1, C.2, C.3, C3.1 and C3.2) might each be administered by a different MHS administrator

Y and Z are other parts of the MTS outside the routing collective which, for the purposes of MHS routing are modelled in the directory as routing collectives (irrespective of whether they practice MHS routing or not). In this Recommendation | Technical Report, these are called External Routing Collectives.

Also, some of the internal routing collectives might not consist of Routing MTAs that conform to ITU-T Rec. X.412 | ISO/IEC 10021-10 and they cannot be configured with routing information as specified in that Recommendation | International Standard. However, these routing collectives will be represented in the directory as routing collectives for the benefit of all other MTAs needing to route to these non-routing MTAs.

A number of connection groups connecting MTAs within and outside the routing collective are illustrated (CG1, CG2, etc.).

³⁾ In terms of the number of hops through subsequent MTAs.




Figure 2 – Routing Collective structure example

Figure 3 illustrates the resultant routing collective hierarchy as a tree structure. It does not model the connection groups. The routing collective hierarchy is directly represented as a directory subtree, the entries of which contain the connection groups available to each routing collective.




Figure 3 – Routing Collective Hierarchy example

Thus, from an MHS administrator's point of view the MTS is partitioned along administrative boundaries, and the MTAs are formed into a hierarchy of routing collectives. The MTAs directly supporting a routing collective are all managed as a single unit by the same MHS administrator. A routing collective may be partitioned into a number of subordinate routing collectives. The MHS administrator of a superior routing collective delegates administrative authority for the subordinate routing collective to a different MHS administrator. Routing collectives thus form a hierarchy reflecting MHS administrative boundaries.

Each routing collective is supported by one or more MTAs. Each MTA is also modelled as an individual routing collective. All MTAs within a routing collective are configured to know of one or more direct or indirect routes to each other. The Routing Collective concept also provides a convenient way of summarizing and identifying the routing capability of groups of MTAs by hiding the internal structure and connectivity of those MTAs (this allows other administrators to construct routes to it without getting involved in its internal structure). The primary characteristic of a routing collective is its message delivery and routing capability. Examples of routing collectives include:

- A Management Domain;
- The group of MTAs within an organization operated by a particular department or division;
- A collection of co-operating MDs;
- A single MTA (each MTA is considered to be a routing collective in its own right).

Parts of the MTS outside a routing collective (e.g. in another MD) are also modelled as routing collectives with distinct routing and delivery responsibilities for specified sets of OR-addresses.

Where a number of MTAs are subject to a common routing strategy (e.g. to support a Management Domain, or the MTAs supporting an organization), they may be defined to belong to a single superior routing collective. This will result in a two level routing collective hierarchy, where the top level is the collection of all the MTAs, and the subordinate routing collectives represent the individual MTAs.

However, in complex cases where a more devolved management of the MHS is necessary, a subordinate routing collective may be subdivided into further subordinate routing collectives. Subordinates may need to be defined for the following cases:

- 1) Departments within the organization which have autonomous control over one or more of their own MTAs or which fund their own private data communications links to external parts of the MTS;
- 2) Departments who wish to maintain anonymity of their user population (i.e. operating an ex-directory facility).

Routing collectives are designed and specified by MHS administrators independently of connection groups and OR-address plans.

6.9 Routing Collective Directory Representation

A routing collective directory subtree is constructed to model the hierarchy of routing collectives within each top level routing collective. For instance, the tree of Figure 3 can be implemented in the directory to model the administrative structure of routing collective X in Figure 2. The routing collective administrator allocates a directory name to each of the routing collectives, and sets up a corresponding directory entry to contain information about it. Routing collectives are hierarchically related to each other, and this is reflected by their directory entries and their respective directory names that form a routing collective subtree in the directory. The name of each subordinate routing collective is derived from its superior's directory entry. Each routing collective's entry contains a list of one or more *connection groups* through which it can access MTAs in other routing collectives, and through which other MTAs can pass messages into the routing collective. A routing collective's connection groups therefore express the topological connectivity of a routing collective within the MTS as is detailed in 6.8.

6.10 Connection Groups

A connection group is defined as a number of MTAs which can exchange messages directly with each other and which therefore represent an element of connectivity between MTAs, i.e. the MTS topology.

A connection group:

- requires a set of data communications links (or a common data network) over which each pair of MTAs in the group can exchange messages directly;
- implies the existence of an agreement⁴⁾ between the MHS administrators of each of the MTAs to allow their MTAs to exchange messages, and it implies the existence of agreements on data communications security, passwords and security contexts for the exchange of messages between the MTAs;
- requires a common message interchange protocol and underlying data communications protocol;
- is assigned a directory name and a corresponding directory entry to hold information about the connection group.

There are two types of connection group:

- 1) Enumerated connection groups, for which all the MTAs in the connection group can be listed and are known to each other;
- 2) Un-enumerated connection groups, where all of the connection group's member MTAs can exchange messages with all other MTAs attached to some common data communications network (e.g. an IP network or a public packet switched network) and it is not generally possible to list all of the MTAs which are members of a connection group attached to that network. Also, the MTA membership of the connection group may change in time.

Connection groups are specified by MHS administrators independently of routing collectives and OR-address plans. The concepts of connection groups are further explained in diagrams in 8.4.

6.11 Connection Group Directory Representation

Each connection group is represented by a directory entry. The entry may provide a list of the MTAs that it connects (only if it is an Enumerated Connection group), and it will contain a set of communications related parameters to enable the MTAs in the connection group to exchange messages with each other.

6.12 OR-address Plan

Organizational administrators develop OR-address plan in conjunction with MHS administrators to fulfil the requirements of the organization. The OR-address plan specifies the selection of OR-address attribute types used to identify MHS users within a routing collective. It might also impose constraints on some OR-address attribute values, e.g. by naming Organizational Units etc. An OR-address plan usually represents all of the OR-addresses for which a routing collective has responsibility for delivery or DL-expansion. This is the routing collective's internal OR-address space. There is an exception to this with Secret OR-address subtrees (see 6.17.3.5.2).

This is done within the constraints of the MHS standards and is unconstrained by the particular MHS products (MTAs) which support the MHS.

OR-address plans are implemented by the registration authorities that allocate attribute values within the constraints of the OR-address plan to create user's OR-addresses.

6.13 OR-address plan directory representation in OR-address subtrees

OR-addresses are represented in the directory as OR-address subtrees that hold MTA routing information. A number of OR-address subtrees are constructed in the directory to hold routing information used by one or more MTAs. Each directly reflects the OR-address plan and may also reflect a summary of OR-address spaces which are defined for other parts of the MTS outside the routing collective. OR-address subtrees are discussed more fully in 8.7.

MHS routing defines a method to represent each OR-address as a Directory Name. This is achieved by assigning a Relative Distinguished Name to represent each OR-address attribute value. MHS administrators construct one or more OR-address subtrees using these directory names. Each entry contains routing information for MTAs concerning the OR-address represented by the entry. This provides a useful way to organise and store routing information for access by MTAs. MTAs are configured to access one or more OR-address subtrees.

⁴⁾ Such an agreement may be implicit (e.g. where an administrator manages a number of routing collectives), informal, or supported by a formal contract between the administrators of different routing collectives. However, in all cases there is a need to agree on the technical details of message interchange.

6.14 MTA definitions

Each Routing-MTA is defined by two different directory entries:

- 1) A directory entry and name to hold information on the protocols it uses, the network locations for each different data communications network it is attached to, the OSI Application Entity Title of the MTA, the MTA's name and Global domain Name;
- 2) A directory entry which represents it as a routing collective, and which contains its routing collective identity and the names of the OR-address subtrees that it must use during routing.

The reasons for each MTA having two entries are purely historical. It was necessary to define the first entry [(1) above] at the very beginning of MHS standardization process (i.e. in 1984), and the information contained in 2) above only became necessary with the more recent standardisation of MHS routing.

6.15 The MTA's Message Routing Process

An MHS administrator configures the directory with routing collective subtrees, connection groups, MTA definitions and OR-address subtrees and then configures each MTA to reference its own directory definition entry and a sequence of OR-address subtrees which it should refer to. Each MTA then runs an initialization process to obtain its working information and gain access to the OR-address subtrees it should use. During this initialization process, it obtains knowledge of its own connectivity to other adjacent MTAs and obtains the minimum knowledge necessary to be able to route messages to all other parts of the routing collective and through Exit MTAs to destinations outside the routing collective.

As a result of this process, the MTA can relate any non-adjacent target MTA in the routing collective to one of its adjacent MTAs. Once this process has completed, the MTA can enter active routing service.




Figure 4 – MTA Roles in a routing collective

Figure 4 illustrates the various roles that an MTA may assume with respect to a particular message being routed by an MTA in the routing collective.

- MTA A is called the Current MTA because it is actually processing a message. Its OR-address subtrees will suggest transferring it to a particular Target MTA on the basis of the recipient's OR-address (e.g. in this case either B, C, or D).
- MTA B is an Adjacent MTA to A because it is directly connected to A. In this case MTA B is the next hop on the journey to either C or D.
- D is considered to be an exit MTA because it has connection groups to MTAs outside the routing collective which support other OR-addresses (unlike A, B, and C which only have connections to other MTAs within the Routing Collective). Because of this external connectivity, any messages leaving the routing collective must pass through D. Exit MTAs must have access to either at least one exit connection group or a transit exit connection group.

As MTA A routes a message, it converts each recipient's OR-address into its corresponding directory name. It uses this to read its configured OR-address subtrees in sequence until it finds appropriate routing information (the name of a target routing collective) for that OR-address. The MTA then uses its internal map of its local routing collective connectivity, which it constructed during initialization. This enables the MTA to determine which of its adjacent MTAs it should transfer the message to through and which local connection group should be used for the transfer. So, in summary, there are three steps to determine the adjacent MTA for any given OR-address:

- For any given OR-address, one of the MTA's configured OR-address subtrees provides the identity of a target routing collective to which the message should be transferred;
- If the target routing collective is within the same parent routing collective then the MTA's internal tables enable it to select an appropriate adjacent MTA to which the message should be transferred towards the target routing collective;
- If the target routing collective is not known (e.g. it is not a key routing collective) the target routing collective's entry is read from its directory to determine a list of the connection groups to which the target routing collective is attached. This list of connection groups is then used to select either an entry MTA to that routing collective or an exit connection group from the current routing collective according to the available connectivity.

6.16 MHS Routing Administrative Roles

MHS Routing implies that the tasks of designing and implementing a routing strategy for a part of MHS are divided into the following three roles:

- 1) The organizational administrator who makes decisions about how to identify staff, where they are located, the structure and geographical distribution of departments within the organization and the designation of organizational roles, and whose primary concerns are the specification of an OR-address plan reflecting the internal structure of the organization being supported by an MHS. They also arrange for individual users to be provided with access to the MHS in co-operation with the MHS Administrator;
- 2) The OR-address Registration Authority, which is responsible for allocating OR-addresses (according to the requirements of the organizational administrator) and ensuring that each user's OR-address is unambiguous (i.e. ensuring that no two users are allocated the same OR-address);
- 3) The MHS administrator, who is responsible for the connectivity and routing strategy of the MTAs which directly support a routing collective, and for co-operating with MHS administrators of other routing collectives and other parts of the MTS to achieve the desired connectivity and routing capabilities of the MTS as a whole. However, an MHS administrator is not considered to be responsible for the administration of subordinate routing collectives to which management responsibility has been devolved.

These distinct administrative roles are co-ordinated within the organization to establish an MHS service that is appropriate to the user community. The roles may all be performed by a single person, or may be performed by different people.

6.17 MHS Routing Administrative Tasks

In this description, it is assumed that an MHS administrator is going to establish and maintain a top level routing collective and that the complete task must be carried out. In the case of a subordinate routing collective, some of the tasks may already have been completed by the MHS administrator of the superior routing collective and the subordinate's administrator is expected to work within the context of the superior routing collective administrator's definitions. After having established a routing collective, it will need to be maintained and changed to reflect the registration of new OR-addresses, and addition of further MTAs, connection groups and subordinate routing collectives.

The approach specified in this Recommendation | Technical Report may be applied to a routing collective managed by a single or devolved administration. An administrator may need to apply the tasks selectively to:

- implement one or more routing MTAs;
- upgrade one or more existing MTAs to routing MTAs;
- implement a complete new installation;
- extend an existing routing collective by adding new OR-addresses, adding routing MTAs or upgrading existing MTAs or adding subordinate routing collectives.

The end effect of the administrator's work is to ensure that:

- all MTAs in the routing collective are able to route messages to each other by provision of an adequate data communications infrastructure;
- each MTA knows how to reach Exit MTAs which have access to external connections to other parts of the MTS;
- each MTA can identify the subordinate routing collective responsible for each OR-address supported by the routing collective.

Establishing MHS Routing for a part of the MTS should be carried out by the following sequence of tasks:

- a) Analysis of the organization's MHS requirements;
- b) Design the MTS topology, the OR-address plan and the routing collective structure as three independent tasks based on the information obtained from the analysis;
- c) Configure the directory to hold entries for a routing collective subtree, one or more OR-address subtrees, MTA definitions and connection group definitions;
- d) Configure MTAs to access their own directory definitions and routing information;
- e) Initialize each MTA to read its directory information and to create sufficient internal information for it to carry out its routing task;
- f) Place initialized MTAs into service.

An administrator performs these tasks as outlined in the following subclauses.

6.17.1 Organizational Requirements Analysis

The organizational and MHS administrators should co-operate in performing an analysis of the organization's MHS requirements and identify the scope and structure of each of the routing collectives. Clause 7 gives details of this task which is oriented to establishing the following information:

- the user community to be provided with messaging services;
- the specific messaging requirements of the organization and users served by the routing collective, including OR-addressing requirements;
- all of the existing component MDs, MTAs which will form a part of the routing collective;
- the data communications links and networks accessible by each MTA;
- the external MTS environment within which the routing collective is to operate (i.e. other Management Domains).

The organizational requirements analysis task should be completed before construction of an MHS, or when it is decided to convert all or some MTAs to using the MHS Routing standard. It may also need to be carried out if the MTS or the organization undertakes a major revision of its MHS requirements. The information resulting from this task will be used in the subsequent MHS routing design tasks.

6.17.2 Design Tasks

The design of a routing collective must deal with the topology of MTA interconnections, the administrative structure of the routing collective and its subordinates, and the OR-address plan.

6.17.2.1 Topology Design

This task identifies the MTAs that are to support the routing collective and their interconnections. The administrator should take into account the information obtained from the organizational requirements analysis on existing messaging systems and the MTS environment (e.g. other Management Domains) to:

- ensure that MTAs are available to serve all users;
- ensure that the routing collective's MTAs are adequately connected by data communications links;
- ensure that the routing collective's MTAs are adequately connected to other parts of the global MTS and to ensure that the data communications links are adequately sized for the expected traffic loads;
- identify any messaging systems components (MTAs and data communications links) which should be procured to ensure that sufficient MTA and data link capacity is available;
- identify all DUAs and DSAs, including the necessity to maintain mirror DSAs and multiple MTA-DSA connectivity to meet resilience and availability requirements.

The result of this step should be a topology of MTA connectivity both within the routing collective and to other parts of the MTS. In particular, it will identify the connection groups that are accessible by each MTA and each subordinate routing collective. This information will be used to construct the routing collective's directory entry. It will also identify all of the connection group entries that need to be defined in the directory.

Topology design is dealt with in more detail in clause 8. The topology will subsequently need to be maintained to reflect addition or deletion of any MTAs or connection groups.

6.17.2.2 OR-address Plan

The organizational and MHS administrators should develop an OR-address plan to reflect the addressing requirements of the organization by:

- choosing appropriate OR-address forms and attribute types for each group of users served by the routing collective;
- establishing appropriate address registration procedures and registration authorities;
- obtaining the OR-address attribute value registrations used to identify all of the users of the routing collective from external OR-address registration authorities (e.g. the Country Names, MD Names, Organization and Organizational Unit Names).

The result of this step will be a specification of an OR-address plan for the routing collective and the establishment of OR-address registration authorities. This task is detailed further in clause 9.

6.17.2.3 Routing Collective Design

The administrator should develop a hierarchic administrative structure for the routing collective. Routing collectives and the conditions under which they are formed are explained in more detail in clause 10. The administrator should decide how or whether the routing collective is devolved to subordinate routing collectives. Design of the routing collective results in the following:

- Identification of each subordinate routing collective;
- Identification of each MTA supporting the routing collective;
- Identification of the connection groups attached to the routing collective and those accessible through its subordinates.

Once this substructure has been decided, the top level, and any subordinate routing collectives can be configured into the directory. A suggestion as to how this may be achieved is detailed in 8.2.

6.17.3 Configuring the Directory

Once the three design tasks have been completed, the routing collective administrator can proceed with configuring the directory to reflect these designs. The following subsections give an overview of how this is done. The details of these tasks are provided in clause 10. The configuration is based on information obtained from the routing collective, the OR-address plan and the topology design tasks.

6.17.3.1 Directory Information Base Configuration

The administrator should prepare the directory with base entries to locate the various types of directory information specific to the routing collective. The base entries should include the following:

- the base entry of a routing collective subtree to represent the routing collective and its subordinates;
- the base entries of OR-address subtrees specified for use by the routing collective, its subordinates and their supporting MTAs;
- locations at which the MTA application entity definition entries are to be held;
- locations at which the definitions of connection groups defined for the routing collective are to be held.

Subclause 8.2 suggests one way of organizing this that allows the administrators of different routing collectives to easily locate the entries they require.

Subordinate routing collective administrators will need to co-ordinate the way they do this with the superior routing collective's administrator.

6.17.3.2 Configuring Routing Collective Directory Entries

A top level routing collective and each of its subordinates should be defined by an entry in a directory subtree to reflect the hierarchical relationships between the routing collectives. The primary information content of each routing collective entry is a list of the connection groups that the routing collective has access to, and the way each connection group can be used. Connection groups are classified for use by each routing collective as follows:

- Entry connection group, implying that messages can enter the routing collective through the connection group;
- Exit connection group, implying that messages originating from users served by the routing collective can exit the routing collective through the connection group;
- Transit exit connection group, implying that messages originating from users supported outside the routing collective may be relayed internally through the routing collective to the exit connection group.

Routing Collective configuration is further detailed in clause 10.

6.17.3.3 Configuring Connection Group Directory Entries

There are two types of connection groups which the administrator will deal with: those defined and used solely to connect MTAs within the routing collective, and those connecting the routing collective's MTAs to MTAs in other parts of the MTS. The difference between them is that the routing collective administrator has control over the design and identification of those defined for internal use, but will need to agree the definition and identification of each external connection group with other MTS administrators (or adopt existing definitions where they exist).

If a connection group definition entry does not exist then the administrator should create one.

If a connection group already has a directory definition agreed among those who use it (i.e. administrators of other parts of the MTS) and it already has an accessible directory definition, then it will be sufficient to simply reference that definition by its directory name. The administrator does not need to configure an entry for it. Access to a common directory is required for this to be possible. For example, a group of PRMDs may establish and define a connection group based on a common public packet switched network that they can all access.

However, if an externally defined connection group does not have an accessible definition in the directory (i.e. it has been defined by the management of some other parts of the MTS and it has no accessible definition in the directory, or a common directory service does not exist), then the MHS administrator should configure a proxy entry for the connection group and use this to configure routing collective directory entries which have access to the connection group. A proxy connection group definition will require synchronisation with the actual state and membership of the real connection group on an ongoing basis, i.e. the proxy definition of enumerated connection groups will have to be updated as and when MTAs join or leave it.

The administrator may choose to define connection groups for use by:

- the particular MTAs to which the connection group is directly attached;
- subordinate routing collectives.

The definition of connection groups in the directory is detailed in 8.4. Although the directory names of connection groups may be arbitrarily allocated, subclause 8.2 suggests a way of allocating them in a consistent manner.

6.17.3.4 Configuring OR-address subtrees

Each MTA obtains its routing information from one or more OR-address subtrees which it has been configured to read from the directory. Although all of the OR-address subtrees are structured in essentially the same way and have similar contents, for the purposes of this Recommendation | Technical Report, it is useful to distinguish between different types, each constructed for a different purpose:

- A Reference OR-address subtree is configured for each top level routing collective to hold internal routing and message handling information for every OR-address supported within the routing collective. This subtree is accessible from all of the MTAs supporting the top level routing collective and all of its subordinates. Each entry in the OR-address subtree represents an OR-address, and identifies the target MTA which can either deliver, DL expand, non-deliver or redirect the messages for that OR-address. In some circumstances, a Reference OR-address subtree may be partitioned, access controlled and spread across different DSAs to reflect different operational strategies;

- Secret Reference OR-address subtrees are identical in function to Reference OR-address subtrees, but they are administered to allow the OR-addresses they contain to remain confidential (e.g. Ex-directory MHS users);
- A sequence of one or more External OR-address subtrees to hold routing information regarding connections to other parts of the MTS outside the routing collective. The first Exit OR-address subtree contains the most preferred routing information to all destinations outside the routing collective. It may be combined with the Reference OR-address subtree to reduce the number of directory reads required. Subsequent External OR-address subtrees hold alternate, less preferred routing information to destinations outside the routing collective.

The contents and use of these are all identical, however they fulfil different roles.

6.17.3.5 Multiple OR-address subtrees

In many simple cases, it will only be necessary for the administrator to construct and use a single OR-address subtree which consists of the combination of the Reference OR-address subtree and a single External OR-address subtree. However, an MHS administrator may need to configure MTAs to read a number of different OR-address subtrees for various reasons. The following paragraphs give general guidelines on when multiple OR-address subtrees may need to be created and configured into MTAs.

6.17.3.5.1 Splitting Reference OR-address subtrees for administrative or operational reasons

A Reference OR-address subtree may need to be split for the following reasons:

- Different parts of the Reference OR-address subtree must be supported on different DSAs, and the DSAs are not capable of using the DSP;
- Because different parts of the OR-address subtree are managed by different administrators, and the existing directory write access control mechanisms are not considered adequate (e.g. in the absence of Strong Authentication DSA facilities);

6.17.3.5.2 Creating Secret OR-address subtrees

MTAs which support ex-directory OR-addresses that are implemented on DSAs only accessible to those MTAs implies that the secret OR-address subtree should be specified as a separate OR-address subtree from the Reference OR-address subtree.

6.17.3.5.3 Implementation of Firewalls and multilevel Security Systems

A further reason why a split in a Reference OR-address subtree may be necessary is to ensure that message traffic between different parts of the MTS supporting a routing collective must pass through a firewall (i.e. to separate different parts of the MTS which operate at different security contexts in a multi-level security system).

6.17.3.5.4 Setting default routes

Each MTAs OR-address subtree list must terminate with a default OR-address subtree, which specifies a route to be taken if a message's OR-address has not been found in any of the previous OR-address subtrees. This default route might be a target MTA that leads to either an ADMD, or to another MTA within the routing collective that has more comprehensive routing information.

6.17.3.5.5 Controlling access to locally funded connection groups

For instance, although the administrator of C in Figure 3 has specified that C.1, C.2, and C.3 should route all messages to Z through Y (assuming Y and Z to be connected in some way), the administrator of C3.1 may have established an alternative direct and locally funded route to Z using CG6. It is not intended to allow traffic originating from any MTA other than C3.1 to use it. The administrator of C3.1 constructs a separate OR-address subtree representing the OR-address spaces reachable through CG6 and Z and uses it as a preferred route.

6.17.3.5.6 Defining multiple External OR-address subtrees for alternative routes to the same external OR-address space

In simple cases, a top level routing collective's administrator will specify a single target MTA (e.g. B1 in Figure 3 to reach the OR-address spaces of routing collective Y) within the routing collective for each external OR-address space. This is used by all MTAs supporting the routing collective.

However, to avoid a potential single point of failure at the exit MTA B1, he may define a second preference Exit MTA (e.g. B.3 or B.2) to the external OR-address spaces in Y. In this case, the routing collective administrator must define a separate OR-address subtree to carry the routing information for each alternative.

6.17.3.5.7 Defining multiple External OR-address subtrees to achieve efficient routing or ensure balanced traffic loading

The administrator of a large and complex routing collective may need to instruct different sets of MTAs to route messages for the same OR-address space to different Exit MTAs (e.g. in Figure 3, telling all MTAs in routing collective A to route messages for Y via MTA B.1, and telling all MTAs in routing collective C to route messages to Y through B.3).

6.17.3.6 Configuring MTAs to read multiple OR-address subtrees

A consequence of defining multiple OR-address subtrees covering identical OR-address spaces, is that it will often be necessary to configure MTAs to read two or more OR-address subtrees in a particular sequence. The sequence should reflect the following arrangement of OR-address subtrees:

- The first set should either be the reference subtree, or the component OR-address subtrees where the reference subtree is split. Arranging the largest Reference OR-address subtree portion to be first may give a slight advantage in efficiency;
- The second set should include Secret OR-address subtrees and separate OR-address subtrees implemented to support multi-level security and firewalls;
- The third set should be External OR-address subtrees that contain locally defined routes.
- The fourth set should be OR-address subtrees defined by superior routing collective administrators. The most locally defined ones should be placed earliest in the list;
- The fifth set should consist of the single default OR-address subtree.

Reference and External OR-address subtrees may be combined under certain circumstances.

6.17.3.7 Publishing Routing capabilities

Administrators need to inform the administrators of other parts of the MTS of the routing and message handling capabilities of their routing collective. The OR-address subtree structure provides a convenient way of publishing this information. This technique is detailed further in 8.9.

6.17.3.8 Connecting to non-MHS systems

There are two types of non-MHS systems that the administrator needs to consider:

- non-routing MTAs and messaging systems which are embedded within the routing collective, and form a part of it;
- other parts of the MTS which are outside of the routing collective.

Both must be represented to the MHS system through MTAs, and the administrator should treat them as if they were routing collectives and define them with routing collective entries in the directory. It will also be necessary to provide an MTA directory entry to be able to hold the authentication information for the MTA representing the routing collective.

6.17.4 MTA Configuration

Each MTA has two entries in the directory:

- 1) its MHS Message Transfer Agent entry, which holds information concerning its communications, authentication and network addresses;
- 2) its definition as a routing collective. Each MTA is represented by an entry to define it as a routing collective. From this single directory entry, the MTA can initialise itself with all of its routing information. The information includes the directory name of the MTA's definition as an MHS Message Transfer Agent as described in 1) above, and the sequence of OR-address subtrees that the MTA should use in routing.

MTA configuration and organising its OR-address subtree sequence is detailed in 8.8.

6.17.5 MTA Initialization

A Routing-MTA automatically initializes itself when invoked by the MTA operator after it has been configured to access its own routing collective directory entry. From this entry, the MTA reads:

- its own routing collective definition configuration, including the base entry of the routing collective subtree of which it is a member;
- the name of its MHS Message Transfer Agent entry;

- the connection groups to which it is directly attached;
- the sequence of OR-address subtrees which it should read for routing information during message routing.

It then searches through the routing collective subtree of which it is a member (the top of which it can determine from its own routing collective directory name) to examine the connection groups to which all other MTAs in the routing collective are attached.

From this information, it constructs internal tables which tell the MTA which of its connection groups and adjacent MTAs it should pass messages to for each recipient's OR-address, i.e. it builds a local map of its connectivity and possible routes through its connection groups to other MTAs in the routing collective.

NOTE – This is not a complete map of entries of the top level routing collective, but it is sufficient to ensure that all parts of the routing collective can be reached.

Re-initialization of the MTA will only be necessary if one or more of its configuration parameters change.

After initialization, the MTA knows of one or more routes to every other part of the top level routing collective (including all Exit MTAs provided by superior routing collectives) and all of its subordinates, and it can enter service and proceed to route messages as outlined in 8.11.

6.18 Prerequisites for MHS Routing

The following systems and resources are required to support a routing collective:

- At least one of the MTAs being managed should conform to the requirements of ITU-T Rec. X.412 | ISO/IEC 10021-10;
- Directory systems shall conform to the requirements of ITU-T Rec. X.500 | ISO/IEC 9594;
- The MHS administrator should ensure that all MTAs can exchange messages with all other MTAs within the routing collective using direct or indirect routes by providing an adequate data communications infrastructure;
- The MHS and organizational administrators should ensure that administrative roles are established for the routing collective including OR-address attribute registration authorities;
- The MHS administrator should ensure that OR-address registrations (e.g. Country, ADMD, PRMD and Organization Names) which are used as parts of any OR-addresses supported by the routing collective are obtained from the appropriate external OR-address registration authorities. These may be national authorities for the case of country names, MD names and Organization names, and Organization internal authorities for Organizational Unit and other locally administered names;
- The MHS administrator should obtain configuration data for all MTAs and connection groups directly supporting the routing collective, and the communications related data for all external MTAs to which the routing collective has direct connection groups. This will include network addresses of MTAs, passwords etc. The internal information may already exist, or it must form part of the configuration process. The external information may be obtained from administrators of other parts of the MTS or it may already be available in the directory;
- The MHS administrator should obtain knowledge of all existing external messaging system connectivity requirements and agreements.

6.19 Organizational Requirements Analysis

Organizational and MHS administrators should complete a requirements analysis to gather and document all of the information which will influence the design of the MHS. The analysis will need to be done for all parts of the organization for which an MHS is to be installed. In the following subclauses, it is assumed that the scope of the task is that of a complete top-level routing collective. The required information is summarized in the following list. It will be used in further steps of the methodology.

6.20 The user community

The identity of the user communities affected and their respective roles must be established. The information required, and its impact on the design of MHS routing is given as follows:

Factor	Impact
The geographic locations of users to be served by messaging systems.	This may impact the design of the OR-address plan, and the geographical location of messaging systems needed to support the user community.
The operational organizational structure of each user community into departments and divisions.	This may also impact the design of the OR-address plan, and might additionally impact the administrative structure of the MHS systems, particularly where MHS management will be devolved.
Distribution of Organizational units over geographic locations.	If Organizational units are highly distributed over distinct geographical locations, where the location itself is significant in Organizational terms, then this suggests that the OR-address plan might benefit from attributes which reflect the different locations. However, this would not be the case if the organization's administrative structures ignored locations.
User mobility (both geographically and within the organization).	This will indicate whether users are more likely to change between departments or between geographic locations, and it will help the administrators to design an OR-address scheme with maximum stability, (e.g. if users are more likely to change geographic location, then geographic attributes should be avoided; if users are more likely to change departments, then departmental attributes might best be avoided).
Anonymity requirements.	Some user groups may need to remain anonymous. They should be identified and either be represented through secret OR-address subtrees, or be allocated non-informative OR-address attribute values which do not identify the person or the role.

6.21 Cultural requirements and constraints

This primarily has impact on the OR-address schemes implemented. The organization will already have an internal postal addressing structure, and this may impact the OR-address plan. The organization may have requirements for role addressing, personal addressing, anonymous addressing, operational organizational unit addressing and for geographical location addressing (e.g. site identity, building name, floor number).

6.22 Organizational messaging administrative structures

If other messaging systems already exist, then their current administrative structures should be investigated. This should cover any remaining telex systems, proprietary systems and office automation systems that must be integrated or connected to the MHS system under design. In many organizations, messaging and communications infrastructure are centrally managed; in others the organization's departments act autonomously and procure and manage their own messaging systems, particularly if the organization is highly distributed and diverse, e.g. these might be organizational departments (and in some cases their subordinate divisions) which, by tradition, operate, fund and manage their own messaging systems.

The primary impact of this information will be on the administrative structures applied to the MHS system and the resultant design of the routing collective hierarchy. If departments or divisions operate their own messaging systems, then they should be represented as distinct routing collectives to provide clear divisions of responsibilities for operation of the different systems.

6.23 Existing, non-standard and other Messaging Systems

The organization may have a number of existing or planned messaging systems which should either become a part of the total messaging system under design, or to which connections must be established. Each should be represented by a routing collective which that must be integrated into the MHS design. These will include:

- Access units to other telematic services (Telex, Fax, Physical Delivery, Teletex);
- Proprietary messaging systems which must interconnect into the MHS through MTA gateways;
- MHS systems which, through proprietary limitations in their OR-address support or deficiencies in their protocols must be defined as separate MDs, Organizations or Organizational Units;
- Legacy MHS systems (e.g. X.400 1984 systems);

- Existing MHS MDs;
- Connections to other MHS systems in other organizations;
- ADMD services used by the organization and their individual departments;
- Other planned messaging systems that must be integrated or connected to.

6.24 Available Communications Infrastructure

The requirement analysis should document the availability of data communication networks or point to point links which might be of use to connect MTAs to each other or connect users to MTAs. The connectivity provided may be a strong indication of the most appropriate topology of connection between an organization's sites, and it will also highlight any need to procure further data network switches and communications links.

6.25 Messaging Traffic Patterns and Volumes

Knowledge of the expected traffic patterns and volumes can be potentially derived from the applications that are to be supported by MHS. If the traffic volumes are available, or can be estimated, the parameters might be of use in judging where MTAs should be located (to best handle local message flows) and the requirements for data communications throughput capacity between them.

6.26 Security requirements

Security requirements for different classes of application and different users will affect the topology and routing design in the following ways:

- Provision of access control information, passwords etc. with respect to each pair of interconnected MTAs;
- Control of the flow of messages over routes on the basis of origin (i.e. for the creation of firewalls) will impact on the definition of connection groups;
- The definition of routing collectives to manage secret or 'ex-directory' user groups.

6.27 Specific message routing requirements or policies

The messaging systems may need to support specific internal or external message transfer requirements (e.g. where there are agreements for providing services to external users). This will primarily have an impact on the MTS topology design and the specification of data communication links to support the requirement.

6.28 Alternative routing capability

Alternative routes may be necessary to improve the MHS's resilience and avoid single points of failure. The cost/benefit case for alternative routes should be analysed to ensure that the topology designed can meet the resilience requirements of the organization. The requirement will vary considerably between organizations and applications. Some organizations and applications may be able to tolerate occasional brief interruptions in MTS connectivity (resulting in delay to messages) and might not benefit from a high number of alternative routes. Other organizations and applications will not be able to tolerate interruptions, and will need to plan for the provision of alternative routes.

This primarily has an impact on the design of the connection groups supporting the topology, the provision of data communications connectivity, and the provision of extra Exit MTAs and External OR-address subtrees.

6.29 Directory systems availability

MHS routing is based on directory technology, and if a corporate directory already exists it might be possible to extend the schema to support MHS routing. If not, a new directory system will have to be procured. Redundancy may need to be introduced by provision of mirror DSAs with Directory Information Base (DIB) replication arrangements to meet the organization's resilience and availability requirements. Also, in some cases, the routing collective might be supported by a number of distributed DSAs. There are several ways of arranging this distribution:

- by holding a master copy of the DIB on a master DSA and distributing copies to other DSAs which are accessed by MTAs;
- by partitioning and distributing the organization's DIB over several DSAs.

Whilst designing directory access, it is important to realise that failure of DSAs can compromise an MTA's ability to route messages. Designers should therefore take the following into consideration:

- Mirror images of DIB contents may need to be provided on different (remote) DSAs. If so, the MTA should be configured to access either or both, and also be configured to have alternative network connections to provide alternative data communications paths between the MTAs and the DSAs;
- Some degree of resilience can be obtained where the MTA is allowed to cache directory information, however, this cache must be constantly updated to reflect changes in the real DSA's contents. This might be achieved by using the Directory Information Shadowing Protocol (DISP) and specifying an appropriate update policy.

6.30 Distribution lists

Many organizations have established manual distribution lists that might usefully be configured into the MHS routing system. These should be documented for inclusion into the MHS routing configuration.

6.31 MTS Topology Design

In this step of the methodology, the MHS administrator designs the topology of interconnection of MTAs directly supporting the routing collective, and documents all connections to MTAs outside the routing collective. The outcome of this step is an MTS topology. It should be noted that the design of the MTS Topology is performed independently of designing the OR-address plan and the routing collective hierarchy, since there is effectively no necessity for any dependency between them.

6.32 General Guidance

MHS administrators are advised to establish as high a degree of direct interconnectivity as is possible between MTAs supporting the routing collective. This avoids the need for indirect routing paths and it achieves the maximum transmission efficiency and minimum message transit delays. However, total connectivity may not be achievable where:

- there are cost constraints on providing networks or there is a lack of communications links (e.g. with very widely dispersed organizations);
- there are operational restrictions, e.g. security policy requirements which limit the use of communications links to certain types of traffic;
- the use of gateways between subordinate routing collectives or to other routing collectives is mandated to monitor or audit traffic flows.

6.33 Input information

The MTS topology should be determined using the following basic information:

- The geographical distribution of users and user clusters. This will lead to the identification of potential MTA and Message Store locations;
- The geographical location of existing messaging systems and gateways. This also indicates potential locations of MTAs;
- Messaging system resilience requirements. Many messaging applications are not of the mission critical type, and occasional delays in message delivery which might result from MTA or communications link failures can be tolerated. However, if any application is mission critical, in that short message delivery delays cannot be tolerated then it will be necessary for the topology and routing strategy to provide alternative routes between any two MTAs, duplication of communications links between MTAs and, for particularly critical applications, dual connectivity of MTAs to two or more supporting data networks;
- The predicted traffic volumes between the various user communities, and between the user communities and external messaging systems can be used to generate a traffic volume matrix between each source and destination MTA. This can be used to indicate the connectivity and throughput requirements required within the MTS, and will also give another indication of suitable MTA locations;
- Maximum allowable traffic delay times for message delivery. This will place requirements on MTA connectivity and communications bandwidth;
- Security requirements, e.g. some security policies require messages to pass through particular gateways to external destinations, or require certain classes of information to avoid use of specified (un-trusted) messaging systems and communications links, or require double enveloping techniques to be used. These all impact on the MTS topology design;

- The communications links, networks and services available to support MTA interconnectivity, and the availability of funds to procure more resources may place constraints on the topology;
- The availability of systems support personnel at locations where messaging systems can possibly be located. MTAs and Message Stores require maintenance, and unless the organization has a centralised management department with remote access to management interfaces of each MTA and MS, this factor may place constraints on where messaging systems can be located.

6.34 Results of the Topology design process

The following decisions and actions must result from the topology design process:

1) *Connect users*

Each user should be provided with access to a 'mailbox' (e.g. co-located with an MTA or a Message store – the distinction is not of importance to message routing). A user is ideally connected directly to a mailbox supported by an MTA at the user's location. Highly dispersed individual users do not easily justify installation of MTAs, and it may be preferable to use dial up connections (ISDN or PSTN) to centrally operated MTAs, possibly using the Asynchronous Protocol Specification (ITU-T Rec. X.445). Another alternative is to provide them with access through an ADMD. However, these users can then be considered to be outside the scope of the MHS being designed, and they would also not fit within the organization's OR-address plan.

2) *Choose locations for MTAs*

MTAs may usefully be co-located with clusters of users, or at large organizational sites where there is a high level of local messaging traffic. This will typically be the case where there is devolved, autonomous management of messaging systems.

However, where central control of messaging systems is desirable (e.g. in ADMDs) it may be preferable to install central MTAs at a single site and connect users through data networks. Since this introduces a single point of failure, in mission critical applications, and in cases where minimum availability levels are mandated, the MTAs may need to be duplicated (mirrored) at two or more sites.

It will sometimes be useful to install several MTAs at the same geographical location, each of which serves a small group of users who exchange a lot of traffic, or where there are particularly high traffic loads.

3) *Interconnecting MTAs*

When the siting of each MTA has been chosen, their interconnections can be planned. Development of a traffic matrix, which expresses the predicted traffic between source and sink MTAs will give a guide to the required connections and their throughput dimensions.

This methodology, and ITU-T Rec. X.412 | ISO/IEC 10021-10, allow full and direct connections to be established between all MTAs concerned. This reduces the number of store and forward hops that messages must make and consequently reduces congestion and transit times. Such an approach will require an underlying data communications infrastructure that provides either a fixed or switched connection (e.g. LAN or PSDN/ISDN) between each MTA.

Where total connectivity between MTAs is not possible MTAs should be interconnected as fully as is possible.

Single points of failure (e.g. a single MTA or communications link joining two parts of the MTS) should be avoided by ensuring that multiple paths (direct or indirect) exist between each pair of MTAs. In mission critical installations each MTA should be duplicated at a different geographical site, and each MTA should have multiple connections to different networks or network access points.

Congestion situations may be identified from the traffic matrix. Congestion may be avoided by installing alternative paths, and dimensioning the communications links and MTAs to deal with expected traffic loads.

Excessive transit time delays may be avoided by reducing the number of store and forward hops between source and destination MTAs. In the totally connected case, the transit delays are minimized.

4) *Allocate Directory Names to MHS Message Transfer Agents*

Each MTA is assigned an MTA Directory Name, and an mHSMMessageTransferAgent directory entry is configured for it. (Note that each MTA will additionally be allocated a separate routing collective Directory Name and entry at a later stage in the methodology);

For each MTA-to-MTA connection, it will be necessary to document the Presentation Service Access Points, Security Contexts and Passwords to be used in establishing and using the connection.

5) *Identify and define Connection Groups*

A connection group consists of two or more MTAs that can exchange messages directly with each other. This requires that:

- they have direct communications links between them;
- there is administrative agreement between the MTA’s managements that they can exchange messages directly;
- all of the relevant security, network addressing and protocol information is available to the administrator of each MTA in the connection group.

Each connection group should be allocated a directory name and an entry containing its definition.

The processes may highlight the need for extra data link or switched network connections dependent on the predicted traffic volumes, availability and message transit time requirements.

6.35 Identifying Connection Groups for MTAs

The results of the topology design will be the documentation of the topology for use in the routing strategy design later. The topology could usefully be documented using Table 1 for each MTA in the routing collective.

Table 1 – Connection Groups

(1) Member MTA name	(2) Network or connection type	(3) Group MTA Passwords	(4) Connection Group Name – common name	(5) enumerated or un-enumerated	(6) Description	(7) Security Context
(1) Indicates the connection group’s member MTAs identified by their MTA’s mHSMMessageTransferAgentNames; (2) Indicates the type of communications link, e.g. leased line, switched network; (3) Holds MTA passwords and access control information; (4) Indicates the Connection Group Name allocated to the communications link; (5) Enumerated if the connection group links an identifiable set of MTAs, un-enumerated if it links an un-quantifiable set of MTAs (e.g. on a public network); (6) A textual description of the connection group; (7) The security contexts which govern the use of the connection group.						

The connection groups should include all those supporting the routing collective’s internal connections and all those connecting the routing collective’s MTAs to external parts of the MTS.

The tables of 8.4.1 distil this information.

6.36 Results

Two sets of results flow from this step in the methodology:

- the MTA connectivity for all MTAs comprising the routing collective;
- the definition of all connection groups supporting the routing collective.

Together, these provide all of the necessary topological information necessary for all MTAs in a routing collective to be able to pass messages to each other, and to pass messages outside the routing collective through Exit MTAs.

7 OR-address Plan Design

7.1 General

Another step in the method is to specify the addressing forms and attribute types required for the routing collective. This should be done independently of the design of the topology. It will require the participation of the organizational administrator and the MHS administrator.

MHS Routing allows the freedom to design a routing scheme that suits the organization, and is:

- independent of the MTS topology;
- independent of the MTA that users are attached to;
- may be independent of the physical address scheme chosen by the organization.

It allows the MHS administrator to design an OR-address plan which minimizes the depth of Organization/Organizational units in the OR-address hierarchy consistent with reducing the probability of user naming clashes. It also allows the design of an address scheme which is 'intuitive' to the user (i.e. the OR-addresses of users can be guessed at with a reasonable likelihood of success and based on the organizations operational culture).

This allows user OR-addresses to be allocated according to the requirements of those who organize departmental structures and roles (i.e. not normally those responsible for managing messaging systems). There is a need for communication between the organizational administrators and the MHS administrators to ensure co-ordination of the allocation of OR-addresses to users. This is achieved by specification of an OR-address plan for the routing collective that is implemented as the routing collective's Reference OR-address subtree.

7.2 Information Input

The organizational administrator or departmental manager should evaluate the following information in designing an OR-address plan:

- User addressing culture. Organizations may traditionally be sub-structured on a departmental/divisional basis or on a geographical basis or by a mixture of the two. This will usually be evident in the organization's current internal physical mail addressing structure. Another aspect is the way in which an organization names its departments. Government departments are usually allocated cryptic internal codes according to the departmental organigram. Commercial organizations often allocate 'user friendly' departmental names which are descriptive of each department's role and geographic location or building name. MHS allows the organizational administrator to emulate any of these address schemes.
- Address stability considerations. Users quite often change role, department and geographic location within an organization. The OR-address plan should be organised to minimise the impact of these changes on the OR-addresses allocated to users where possible. For instance, if users often change departments within a particular geographical location, but rarely change geographic location, then it is wise to exclude 'department' name from the OR-address if at all possible.

Address confidentiality requirements. Many organizations need to keep the real identity or role of their members or employees confidential. There are a number of ways of achieving this. The first is by allocating cryptic or numeric attributes in a flat OR-address space (i.e. with minimal OR-address attributes) and suppressing the real identity of departments. This prevents roles being deduced from OR-addresses; the second is by maintaining secret OR-addresses, which can take delivery of messages if the address is right, but the OR-addresses are not accessible through normal directories or routing information.

7.3 OR-address plan specification

The resulting OR-address plan must include the following statements:

- The high level OR-address Attribute Registrations which are to be used in the routing collective e.g. Country Names, ADMD Names, PRMD Names, Organization Names;

- The choice of OR-address forms and their use. There are four forms of OR-address which may be used, and each specifies a set of OR-address attribute types which constitute a valid OR-address. Each of the following forms may be used to locate users within a routing collective:
 - 1) the mnemonic form, which is most widely used and has an element of 'user friendliness' associated with it. The address structure includes names of Organizations, Organizational Units, Personal Names and Roles;
 - 2) the numeric form locates users by means of an integer. It is rarely used except for legacy systems and in cases where user identities should remain confidential.
 - 3) the terminal form, which is used to identify terminals attached to data networks. Its main use is in addressing messages to telematics services (e.g. Telex, Fax), so, it is rarely used within user organizations;
 - 4) the postal form, is provided to address messages to postal patrons via postal or courier services. It may find use in organizations which have a large community of people who do not have access to MHS or any other messaging service, and to whom the organization must send physical mail via an internal courier.

7.4 Special Cases

There are a number of special cases that the organizational administrator should consider in the address plan:

- A single PRMD may inherit multiple names if it connects to more than one ADMD, and if it connects to ADMDs in more than one country, or if the '-' (single space) ADMD name, '0' ADMD name convention, and the 'XX' Country Code conventions are in use. These all result in different OR-addresses for the same user. These need to be dealt with using aliases in the routing collective's Reference OR-address subtree, or by using the alias redirection routing information to ensure that only a single entry is required for each user's preferred OR-address is configured in the Reference OR-address subtree;
- The OR-address plan must absorb any limitations imposed by existing or other planned messaging systems, and recognize that legacy systems addressing constraints may distort parts of the overall OR-address plan;

7.5 OR-address Registration Authority Roles

The OR-address Attribute Registration authority has the following roles:

- To obtain externally registered address attribute values, e.g. for the country name(s), ADMD Name(s), PRMD Name(s) associated with the users served by the routing collective;
- To establish and maintain a Reference OR-address subtree (or secret OR-address subtree) for the OR-address spaces that it is responsible for;
- Specify an appropriate role for each entry in the OR-address subtree, e.g. MTS user, Alias Redirection, DD-Expansion, non-delivery, and double enveloping by means of the appropriate routing information.

The Reference OR-address subtree should always be maintained to reflect the current state of registered OR-addresses. It establishes and maintains the relationship between each MTS user's OR-address and the MTA that delivers to the MTS user.

7.6 Results of the OR-address plan design

The result of this methodology step should be:

- 1) Identification of the OR-address form for each MTS user;
- 2) The OR-address of each MTS user;
- 3) A Reference OR-address subtree;
- 4) Routing advice for each OR-address subtree entry.

8 Routing Collective Design and Configuration

8.1 Introduction

A Routing Collective is the basic unit of design of a routing strategy for a part of an MTS for which an administrator is responsible. Each routing collective is configured into a directory information base and the entries are read by each MTA that supports the routing collective. The directory information base must be configured to contain the following information:

- Information concerning the connectivity of MTAs via data communications links and data networks. This information is expressed as definitions of 'connection groups' as defined in 8.4. Each connection group is defined in a directory entry, and identified by a directory name;
- Information concerning the administration of different parts of the Message Transfer Service. This information is expressed as 'routing collectives'. Each routing collective is defined by a directory entry. The entry contains the directory names of all connection groups accessible to the routing collective. Routing collectives are hierarchically related to each other to represent the potentially devolved management of the MTS to autonomous routing collectives within a 'superior' routing collective. This hierarchy is modelled in the directory as a subtree of routing collectives. Subordinate routing collectives are defined to cover areas of the MTS which are under autonomous (devolved) management, or legacy systems, or to reduce the amount of routing information which needs to be held by MTAs about complex or large parts of the MTS;
- Information instructing each MTA on how it should process messages addressed to each different OR-address. This information is contained in OR-address subtree directory entries, and, for each OR-address space, it either indicates that a message should be delivered, non-delivered, redirected, DL-expanded or whether it should transfer the messages to another specified routing collective. Each routing collective may establish one Reference OR-address subtrees and one or more Exit OR-address subtrees, and each MTA may be configured to read one or more OR-address subtrees;
- Information defining each MTA. Each MTA has two representations in the directory: one defines the MTA's communications identity and capabilities (i.e. its Open Systems Interconnection addresses, passwords etc. defined in its MHS Message Transfer Agent entry). The other defines the MTA as a routing collective (as above), which includes pointers to the various directory entries that contain its definition and routing information.

Although the DIT location of the directory subtree containing connectivity group entries is 'arbitrary', the entries may usefully be kept in a single subtree. Subclause 8.2 suggests an appropriate location for this information. For convenience, the subtrees should be located as follows:

- a) Where a routing collective has sole control of an enumerated connection group, and the routing collective's subordinates all have access to the same DIT, the subtree should be named and located as a subordinate of the top level routing collective's entry. They may be declared as entry or exit connection groups;
- b) Where an enumerated connection group is controlled jointly by several routing collectives which are not related by a superior enclosing routing collective, and the routing collective's subordinates all have access to the same DIT, then its location should be established by mutual agreement between those routing collectives;
- c) Where an enumerated connection group is controlled jointly by several routing collectives which are not related within a superior routing collective, and the routing collectives do not agree on a common location for the connection group definition, (e.g. the MTAs involved cannot access a common directory subtree), then each routing collective shall define the connection group as a proxy connection group at the same level as its top level routing collective directory entry;
- d) Where an un-enumerated connection group is administered within a routing collective and access to a common directory tree is possible, this should be handled as in a) above.
- e) Where an un-enumerated connection group spans unrelated routing collectives, the definition of the connection group is (usually) managed by an organization which freely publishes the connection groups details. If this information is not already accessible from the directory then it should be imported into the routing collective's connection group subtree as a proxy connection group and be maintained in synchrony with the published version.

After the above directory entries have been configured, each MTA is initialised to access its own definition (i.e. the information configured in subclauses 6.17.4 and 8.9).

After the initialisation process has been completed, the MTA is ready to process messaging traffic. The MTA converts each OR-address on a message's envelope into its directory name form, and uses it to read the MTA's OR-address subtrees in sequence until it obtains a suitable routing information for the OR-address. This routing information tells the MTA what action it should take i.e.:

- 1) Route the message to another MTA where that MTA is nearer to the destination;
- 2) Deliver the message;
- 3) Distribution List Expand the message, i.e. where the OR-address identifies a distribution list for which the MTA has expansion responsibility;
- 4) Non-deliver the message, i.e. where the OR-address does not identify an MS user;
- 5) Redirect the message to a designated Recipient MD assigned alternate recipient, i.e. to an OR-address representing a 'dead letter point', where the OR-address is incomplete;
- 6) Redirect the message to a designated Alias Redirection OR-address, specified to manage changes to user OR-addresses.

The MHS administrator creates and updates each routing collective's definition. This is done in the following steps:

- a) Prepare an MHS Routing Directory Information Base entry to hold the definition of the routing collective;
- b) Add subordinate routing collective definitions;
- c) Populate and maintain the directory information base with the routing collective's routes and delivery information. This will involve:
 - Adding/deleting subordinate routing collectives;
 - Adding/deleting connection groups;
 - Adding/deleting MTA definitions;
 - Adding and updating OR-address subtree entries.
- d) Configure each MTA to be able to access its own definitions and those parts of the OR-address subtrees which it should use in the message routing process;
- e) Initialize each MTA so that it reads its own routing information from the directory. During this process, the MTA will automatically discover information about the OR-addresses for which it has delivery or DL-expansion responsibility, the topology of its local MTS, routes to all other MTAs within the routing collective and routes to all other OR-address spaces in the MTS through other MTAs to which it can pass messages.

In addition to this, OR-address attributes must be registered by an appropriate registration authority. The registered (or de-registered) attributes, together with their definitions must be added/deleted from the Reference OR-address subtree or Secret Reference OR-address subtree if applicable.

In the following subclauses, the data required to be input by the MHS administrator is expressed as a set of tables. The contents of these tables are related to the directory entries and attributes at the heading of each column and may be used directly to create the required directory entries.

8.2 Directory Information Base Preparation

To establish a new top level routing collective, the administrator should allocate some directory entries to act as a 'skeletal DIT framework' in which all MHS routing entries will be placed. It is suggested that this entry should be established as a subordinate of an Organizational Role entry in the directory which has been specifically created for the

task of configuring the routing collective, e.g. MHS-routing-collective-xxx where xxx is the routing collective's name. Definitions of the routing collective's MTAs and those connection groups defined for the sole use of the routing collective and its subordinate routing collectives should be held as subordinate entries to the MHS-routing-collective-xxx entry.




Figure 5 – Directory entries representing a top level routing collective

The top level routing collective's entry is a subordinate located somewhere in the global or organizational directory created by the MHS administrator specially for the purpose of MHS routing administration. Its subordinate entries will be:

- The Reference OR-address subtree for the whole routing collective;
- Possibly one or more Secret Reference OR-address subtrees (not shown);
- Possibly one or more external OR-address subtrees;
- Connection group definition entries for any connection groups which are available throughout the routing collective;
- Subordinate routing collective definitions (see Figure 5);
- A single MHS MTA definition if the routing collective consists of only a single MTA;
- Note that Proxy routing collective definitions which represent external parts of the MTS must be defined at the same level as the top level routing collective.

Figure 6 illustrates the structure of the directory subtrees to contain the definition of subordinate routing collective. ST-1 to ST-n, are the bases of its OR-address subtrees.

The connection group definitions are those defined by the administrator of the routing collective. Connection group definitions should be located under the top-most routing collective throughout which they are available for use. It is suggested that MHS MTAs definitions should be located under the routing collective that they support, i.e. only at the lowest level of the routing collective hierarchy.

8.3 Directory Information Base Configuration

Once the MHS Routing Collectives DIB Base Entry (MHS-Routing-collective-xxx) exists, the following procedures can be used to add and delete:

- Connection Groups;
- Subordinate Routing Collectives;

- OR-address subtrees;
- Routes;
- MTAs.

under the routing collective definition entry.




Figure 6 – Generic Routing Collective DIB structure

The information required to do this is:

- 1) The topology of the MTS directly supporting the routing collective and connections to other parts of the MTS outside the routing collective;
- 2) The Reference OR-address subtree;
- 3) Routing information for each entry in the Reference OR-address subtree.

8.4 Connection Groups

8.4.1 Defining Connection Groups

Connection groups are defined by an administrator for each group of MTAs that can exchange messages directly with each other. Connection groups will include at least one MTA within the routing collective, and may contain one or more external MTAs. The administrator will have control over the definitions and names given to all 'internal' connection groups.

However, connection groups including external MTAs will either have no definition, or have already been defined by administrators of other parts of the MTS. The MHS administrator should identify all of these external routing collectives as follows:

- If the external connection group has a formal definition and directory name, and its directory entry is accessible by all of the routing collective's MTAs through some directory service, then that definition and directory name should be used in configuring the routing collective;
- If no connection group definition exists, or one exists but its definition cannot be accessed by all MTAs supporting the routing collective (e.g. because the directory service is partitioned), then a proxy connection group definition and directory entry should be constructed to represent it to the routing collective's MTAs.

To define a connection group, it is allocated a directory name and a directory entry containing the information in Table 2. This information can be deduced from the topology design:

Table 2 – Constructing a Connection Group Definition (to build an entry of the connectionGroup Object Class)

(1) common Name	(2) enumerated Flag	(3) description	(4) groupMTA Password	(5) member MTA	(6) Connection Type	(7) Security Context	(8) Access Controls
<p>(1) Identifies the connection group by directory name. This name should be subordinate to the connection group base entry;</p> <p>(2) Indicates whether the connection group is enumerated (i.e. has a known set of MTA members) or is un-enumerated;</p> <p>(3) Describes the connection group in text;</p> <p>(4) Provides a password which is shared among all member MTAs;</p> <p>(5) Identifies a set of member MTAs;</p> <p>(6) Provides a set of details about the protocols and connections used to connect the MTAs;</p> <p>(7) Is the object identifier of a security context defined among the MTAs;</p> <p>(8) Provides information on access controls applied to the passage of messages over the connection group.</p>							

8.4.2 Adding a Connection Group to a Routing Collective

Each routing collective must be specified with the entry, local exit and transit connection groups through which it can transfer messages to and from other routing collectives. Each proxy routing collective must be associated with one or more exit connection group names. A routing collective’s available connection groups may be specified as in Table 3:

Table 3 – Specifying a Routing Collective’s Connection Groups (related to the routingCollective Object Class)

(1) routingCollectiveName	(2) connectionGroupName	(3) type		
		entry	local exit	transit exit
<p>(1) Identifies a routing collective by its directory name;</p> <p>(2) Identifies a connection group attached to an MTA supporting the routing collective identified in column (1);</p> <p>(3) Specifies the type of the connection group as follows:</p> <ul style="list-style-type: none"> – entry: signifies that the connection group can be used to transfer messages into the routing collective from external routing collectives and that the connection group can be included in the entryConnectionGroupName of the routing collective entry; – local exit: signifies that the connection group may be used to transfer messages to one or more other routing collectives, and that the connection group can only be used to transfer messages which originated within the routing collective. It also indicates that the connection group can be included in the localExitConnectionGroupName of the routing collective entry; – transit exit: signifies that the connection group may be used to transfer messages to one or more other routing collectives, and that the connection group can be used to transfer messages which originated both within the routing collective and outside the routing collective. It also indicates that the connection group can be included in the transitExitConnectionGroupName of the routing collective entry. <p>NOTE – A tool could be used to manage the allocation of connection groups automatically to an MTA’s superior routing collectives. Manual intervention can then be used to remove the connection group from superior routing collective entries that are to be excluded from access to the connection group.</p>				

8.4.3 Deleting a Connection Group Definition

To delete a connection group:

- Remove the connection group name from any of the routing collective definitions that reference it. (This may be achieved using a directory search through the routing collective subtree);
- Delete the connection group’s defining directory entry.

8.4.4 Routing Collectives

This subclause outlines procedures for identifying, defining, adding and deleting routing collective directory entries. A routing collective subtree base entry should have previously been created as recommended as a subordinate entry to MHS-routing-collective-xxx.

8.4.5 Identifying Routing Collectives

8.4.5.1 The Top Level Routing Collective

The main characteristic of the top level routing collective is that it has no superior routing collective. It is managed by a routing collective administrator and includes the portion of the MTS to which this methodology is being applied. By definition, this may be:

- The MHS systems supporting one or more organizations;
- An organizations department or division or a group of departments;
- An MD;
- A group of MDs;
- Part of an MD;
- An MTA.

8.4.5.2 Subordinate Routing Collectives

The MHS routing standard models an MTS devolved administrative structures as a hierarchy of routing collectives. Subordinate routing collectives may be:

- The MHS systems supporting one or more organizations;
- An organizations department or division or a group of departments;
- An MD;
- A group of MDs;
- Part of an MD;
- An MTA;
- Non-routing or proprietary messaging systems;
- Complex groups of MTAs which can be represented within the superior routing collective by means of a few entry and exit MTAs; this ploy can be used to hide unnecessary information and reduce the amount of information that other MTAs have to hold regarding the subordinate.

8.4.5.3 'Non-routing MTAs'

An MTA, or groups of MTAs, within the routing collective that do not participate in MHS routing (i.e. which do not conform to ITU-T Rec. X.412 | ISO/IEC 10021-10), should be modelled as one or more routing collectives (as determined by their respective management structures and routing capabilities). However, the internal structure or strategy of those routing collectives is not modelled further, since it is assumed that their routing strategy is internally consistent.

The administrator should ensure that the routing strategy of each non-routing MTA or proprietary messaging system embedded within the routing collective is compatible with that of the superior routing collective. This may involve modification of each of the embedded system's routing strategies using the proprietary configuration mechanisms specified for those MTAs.

Such routing collectives should conform to the internal connectivity requirement of routing collectives. If not, then either they should be provided with sufficient internal connectivity, or they should be considered as multiple routing collectives.

8.4.5.4 External MTAs

It is advantageous if the definitive directory entries of these MTAs can be accessed directly (instead of making copies) since this will avoid synchronization problems when the MTA definitions are updated. This can only be achieved if the other administrations make their directory information visible. However, MTAs and other routing collectives which do not form a part of the routing collective being designed, (i.e. which support other parts of the MTS) must be modelled as proxy routing collectives to ensure that routes to them can be configured.

8.4.6 Routing Collective Definition

Each routing collective identified should be defined in a directory entry subordinate to its superior routing collective's entry. This automatically models a hierarchic relationship between routing collectives, and implements the routing collective's subtree in the directory.

A routing collective's entry should be configured to contain the information in Table 4:

Table 4 – Routing Collective Definition

(1) routingCollectiveName	(2) Superior's directory name	(3) description
(1) RoutingCollectiveName is the relative distinguished name of the routing collective allocated by the MHS administrator in the context of the superior's routing collective in (2);	(2) The superior's routing collective name, is the directory name of the superior routing collective within the context of which the RDN in (1) is allocated. If it is a top level routing collective, then the MHS administrator should specify a convenient directory entry name under which the routing collective subtree can be established;	(3) A textual description of the routing collective (e.g. London, or Sales).

The intended content of the routing collective entry is the set of connection group identities to which the routing collective has access. These are added subsequently using procedures to add and delete connection groups.

A proxy routing collective should be created for each external MTA to which the routing collective has a connection. These definitions should be located at the same level as the top level routing collective's definition.

8.4.7 Creating the top level Routing Collective

If there is no superior routing collective, then the base entry of the routing collective's subtree is MHS-routing-collective-xxx that was created in the preparation phase.

8.4.8 Adding a Subordinate Routing Collective

If the routing collective is subordinate to an existing routing collective, then the superior's entry in the routing collective subtree should be used as the base, and the subordinate routing collective entry should be added as a leaf to this tree with a relative distinguished name selected by the administrator of the superior routing collective. The administrator should:

- Register the routing collective's name and create a routing collective entry and add it to the routing collective subtree;
- If the routing collective is an MTA, register the MHS Message Transfer Agent Name, create a directory entry for it and configure its OR-address subtrees.

8.4.9 Deleting a Routing Collective

The MHS administrator deletes a subordinate routing collective by:

- Moving the routing collective users to other routing collectives and updating their entries in the Reference OR-address subtree;
- Deleting the routing collective entry from the routing collective tree;
- Deleting any OR-address subtrees which were used exclusively by the routing collective's MTAs;
- If it was the most superior routing collective, then remove the MHS-Routing-collective-xxx entry.

8.4.10 Adding Proxy routing collectives

MTAs in other parts of the MTS outside the routing collective should be modelled by proxy routing collective directory entries so that routing MTAs can recognize and route to them. Each proxy routing collective should be represented in the routing collective subtree as a sibling of the base entry of the top level routing collective.

Proxy routing collective entries should be created for ADMDs, PRMDs, Organizations and Organizational Units which are not part of the routing collective under construction to enable direct message transfer to MTAs lying outside the routing collective.

An ADMD may be able to supply some of this the routing collective information in the form of a publication (see 8.9).

8.5 Configuring Routes for MTAs in OR-address subtrees

8.5.1 OR-address subtree types

Explain the different types, and what their functions are, and when they are created.

8.5.2 OR-address subtrees Model

MTAs require instructions on how to deal with messages for each OR-address or OR-address space. Instructions are configured for each OR-address or OR-address space in one or more OR-address subtrees as outlined in 6.13. MTAs are configured to read these OR-address subtrees sequentially to obtain the optimum routing information for any OR-address carried in a message. A routing collective's administrator organises an MTA's routing information into a number of different types of OR-address subtrees:

- Reference OR-address subtree(s), which contains instructions for routing messages to destinations within the routing collective and for the completion of message processing (i.e. delivery, DL expansion);
- Secret OR-address subtrees, which are confidential components of a Reference OR-address subtree used to implement secret (Ex-directory) OR-addresses;
- External OR-address subtrees, which contain instructions to route messages out of the routing collective for Exit MTAs which have connections outside the routing collective.

Figure 7 illustrates the flow of routing knowledge, consisting of OR-address spaces and their corresponding target MTAs to different parts of the MTS:




Figure 7 – OR-address subtree organization

MTAs A, B, C and D support routing collective RC1. They all have access to RC1's single Reference OR-address subtree, which contains routing information enabling A, B, C and D to exchange messages with each other and with any other MTA defined to support RC 1.

X, Y and Z are MTAs outside RC1 which can route messages to external OR-address spaces (x, y, and z respectively). X, Y, and Z can exchange messages directly with an Exit MTA of RC1 (MTA A) through one or more communications links. From the point of view of RC 1, and MTA A, these data communications links may be modelled as one or more Proxy connection groups. MTAs A, B,C and D are configured to read the External OR-address subtree which tells A to route messages to X, Y and Z for OR-address spaced x, y, and z respectively. Either the actual routing collective definitions for X, Y and Z should be available to RC1, or proxy entries should be created.

MTAs B, C and D are not able to pass messages directly to X, Y and Z because they are not directly connected through connection groups. However, the MTA initialization process will enable MTAs B, C and D to associate the external connection groups with MTA A, so that they can correctly route messages to external OR-address spaces.

8.5.3 Routing Information

Each OR-address subtree entry is configured to contain one or more routing information for the OR-address represented by the entry. These instructions are read and performed by MTAs. The routing information is listed in Table 5.

Table 5 – Use of Routing Information in OR-address subtrees

Function	Reference OR-address subtree	External OR-address subtree
Target MTA	y	y
Non-delivery	y	exceptionally
Alias Redirection	y	exceptionally
Distribution List	y	exceptionally
Double Enveloping	y	y
Recipient MD assigned Alternate Recipient	y	no
Expression Matches	y	y
Next level complete	y (almost always true)	y (frequently false)

The purpose and details of each of these instructions are given in 7.8.

8.5.4 Specifying OR-address subtree bases

The first step in building OR-address subtrees is to provide each with a base entry that is considered to be its root. Typically, a routing collective will define a number of different OR-address subtrees to support its MTAs and an OR-address subtree Base Entry should be established for each identified OR-address subtree and allocated a convenient directory name. It is recommended that the entry is subordinate to the routing collective's defining entry in the routing collective subtree. Subclause 8.2 illustrates this. The base entry does not normally carry a routing information unless it contains a default route (e.g. to an ADMD). In general, such default routes should be avoided. It is preferable that the ADMD provides information about the destinations that it can route to avoid having to non-deliver messages because it has no route to the OR-address.

Typically, the following OR-address subtree bases should be configured for a top level routing collective and its subordinates:

- A single base entry for the top level routing collective's Reference OR-address subtree (subordinates do not have distinct Reference OR-address subtrees);
- Additional bases for Secret OR-address subtrees and where the Reference OR-address subtree is partitioned;
- Potentially a single External OR-address subtree base entry for each MTA which has connections to other parts of the MTS outside the routing collective.

8.5.5 Building OR-address subtrees

Each OR-address subtree entry represents an OR-address attribute. These are listed in Table 6:

Table 6 – OR-address subtree entry types

Base entry	MHS Network Address	MHS Postal Code
MHS Country	MHS Terminal Identifier	MHS Surname
ADMD	MHS Terminal Type	MHS Given Name
PRMD	MHS Numeric ID	MHS Initials
Organization	MHS PDS Name	MHS Generation Qualifier
Organizational Units	MHS PD Country	MHS Common Name

Figure 8 illustrates how the entries representing different OR-addresses are organized and related to each other in an OR-address subtree.




Figure 8 – OR-address subtree structure

8.5.5.1 The Reference OR-address subtree

A top level routing collective's internal routes are defined in its Reference OR-address subtree and zero or more Secret OR-address subtrees. In some circumstances outlined in clause 6, a top level routing collective may have several Reference OR-address subtrees representing partitions in the subtree. Their contents enable all of the MTAs supporting a top level routing collective to pass messages to each other.

The Reference OR-address subtree(s) exactly reflects the OR-address plan provided by the organizational administrator and should contain a directory entry for each OR-address registered and supported by the top level routing collective. Each entry contains a routing information for the OR-address that the entry represents. This tells MTAs to route messages to the particular MTA (i.e. the Target MTA) which deals with and can process messages for that OR-address. All MTAs supporting the same top level routing collective are configured to access the same Reference OR-address subtree.

Reference OR-address subtrees will probably be constructed in several stages:

- 1) Development of the base entry and structure of each subtree to hold an entry for each of the organization's departments and divisions present in the initial OR-address plan. Each entry should be configured with a routing information appropriate for that entry. If no other routing information is applicable, then a non-delivery should be configured;
- 2) Addition of entries for each of the routing collective's initial user population and distribution lists;
- 3) Maintenance, to reflect changes in the OR-address plan such as addition or deletion of organizational departments and division entries, and addition/deletion of individual MHS user's entries.

Reference OR-address subtree entries may be configured to contain any one or more of the routing information identified in 8.7.

The Reference OR-address subtree may also be used as a 'register' of OR-addresses allocated within the routing collective to ensure that no two MHS users acquire the same OR-address. When users are allocated an MHS mailbox, or departmental/divisional OR-address attributes are allocated, they are allocated an OR-address in an appropriate part of the OR-address plan if an entry representing that OR-address is not already present. If an entry is already present then there is an OR-address clash. The organizational administrators must resolve this. When a user no longer requires a mailbox, the entry is removed from the Reference OR-address subtree and the corresponding OR-address becomes 'available' for re-allocation. The role of registering MHS users may be carried out either by the organizational administrator or the MHS administrator or some other nominated person. The role is referred to in this Recommendation | Technical Report as the OR-address Registration Authority.

Adding an OR-address subtree entry happens when a registration authority adds an attribute entry to the OR-address plan. The administrator should decide which type of routing information is appropriate for the new OR-address and configure it into a new entry of the OR-address subtree.

In general, copies of OR-address subtrees should not be made, since this introduces an unnecessary directory synchronisation problem. The Reference OR-address subtree should only be truncated when a secret subtree is implemented by a 'secret' subordinate routing collective, or where the routing collective has an embedded non-routing or proprietary messaging system (for truncation, see 8.6.13), or where a firewall within a multilevel security MHS is implemented.

8.5.5.2 The Reference OR-address subtree

This Recommendation | Technical Report defines the concept of a Reference OR-address subtree. Conceptually, a Reference OR-address subtree models the complete OR-address space of the top level routing collective (including all of its subordinates) from country level down to the leaf entries representing MTS users. It is modelled in one or more directory information tree structures maintained by the OR-address attribute registration authorities.

Ideally, a Reference OR-address subtree should be a single subtree valid for the whole of a top level routing collective and all its subordinates. This will mean that each MTA in the routing collective need only reference a single subtree for all internal addresses. However, there may be exceptions to this as outlined in clause 6. In these cases, MTAs must be configured with two or more OR-address subtrees, each containing a component of the complete Reference OR-address subtree.

The reference OR-address subtree always identifies the delivering MTA within the routing collective except where the user population requires that a secret OR-address subtree is implemented or where embedded non-Routing MTAs or for MTAs which use non-standard addressing schemes. In these cases, the OR-address subtree is truncated, and the identity of a target MTA is provided which can route the message onwards.

Where secret OR-address subtrees or embedded messaging systems are involved, the user registration and the registration authority must be administered by those who manage the subordinate OR-address spaces.

OR-address registration authorities may use the Reference OR-address subtree to aid the registration process and to control the routing collective's address space. Management and implementation of parts of a Reference OR-address subtree may be devolved to subordinate OR-address Registration Authorities for organizational departments that implement Secret Reference OR-address subtrees. This makes use of normal Directory Information Tree conventions. Where OR-addresses are required to remain confidential, the Reference OR-address subtree entries should either be truncated or be subject to read access controls.

8.5.5.3 Secret Reference OR-address subtrees

These are similar in structure, purpose and content to Reference OR-address subtrees. They are implemented if a part of a routing collective's OR-address space needs to remain secret (i.e. ex-directory). Different Secret Reference OR-address subtrees may be administered by different administrators.

Secret OR-address spaces may be arranged by defining separate subordinate routing collectives which have responsibility for those portions of the superior routing collective's OR-address spaces which must remain secret. The subordinate routing collective is allocated one or more OR-address spaces (within the context of the Reference OR-address plan) and the subordinate then implements its own local Reference OR-address subtrees within this context. The subordinate must also establish its own OR-address registration authority to allocate OR-addresses within the context of the OR-address spaces allocated to it by the superior.

The administrators of the subordinate routing collective construct separate 'secret' OR-address subtrees according to their own OR-address plan and allocate OR-addresses within it in exactly the same way as for the superior's Reference OR-address subtree. These subtrees are then configured into the routing collective's MTAs.

The superior's Reference OR-address subtree is truncated at the point where a secret Reference OR-address subtree starts. The name of the MTA which can deal with the message (and which should have access to the Secret Reference OR-address subtree) is configured into the target MTA of the routing information. The entry should have the 'next level complete' attribute removed, to signal that the entry is not a proper 'leaf' entry, and that subordinate OR-addresses do exist. This forces the MTA to adopt the routing information that the entry contains.

8.5.5.4 Personal Name Aliases

Directory Aliases may need to be configured to represent each possible representation of a personal name – i.e. to represent different combinations and sequences of personal names, initials and given names (see 8.6.4 for a discussion on the various forms of aliasing).

8.5.6 Establishing External Routes to destinations outside the routing collective

Establishing external routes to OR-addresses supported outside the routing collective (i.e. in other Management Domains) requires the creation of one or more External OR-address subtrees to contain routing information for Exit MTAs. This allows Exit-MTAs to pass messages to other MTAs outside the routing collective (which are represented as proxy routing collectives).

8.5.6.1 Building External OR-address subtrees

External OR-address subtrees contain entries for each external OR-address space. Each MTA in the routing collective should be configured to read External OR-address subtrees to obtain external routing information. An External OR-address subtree holds the identity of one or more target MTAs outside the routing collective which are known to be able to route messages to those OR-address spaces.

In general, the top level routing collective should seek to provide an External OR-address subtree indicating routes to all parts of the MTS outside the routing collective. However, as outlined in clause 6, multiple OR-address subtrees may need to be defined for overlapping OR-address spaces.

The administrator obtains the information to build External OR-address subtrees from the administrators of distant MTAs. This should preferably be done by configuring the OR-address subtree entry representing the external OR-address space with an alias entry to point to the reference OR-address subtree of the distant MTA or to a subtree created specially for this purpose. This ensures that the information is always up to date. The following should be noted about this approach:

- the external parts of the MTS (MTAs, Directories etc.) must support ITU-T Rec. X.412 | ISO/IEC 10021-10;

- The administrators of the external parts of the MTS must either provide read access to the whole of their Reference OR-address subtree, or provide a special 'censored' version which can be made externally visible. The depth of the externally visible tree should be sufficient to allow MTAs to determine the most efficient entry MTA for each distinct OR-address space where more than one entry MTA exists. The administrators may also bar search access to the externally visible subtrees to avoid un-authorized access to the complete internal structure of the subtree;

However, the distant MTA's OR-address subtrees might not be visible. If this is so, then the information must be transferred in another way. Subclause 8.9 outlines ways of publishing and obtaining this information to deal with this situation.

Each entry in an External OR-address subtree contains the target routing collective name for the external OR-address space that the entry represents. Each External OR-address subtree entry will usually hold:

- The target MTA: the name of the MTA outside the routing collective to which the message should be sent;
- Optionally, Double Envelope information used to secure messages in transit to that MTA.

In exceptional circumstances, the entry may also carry the following information:

- 1) Non-Delivery: if the administrator wishes to bar access to particular OR-address spaces);
- 2) Alias Redirection: if the administrator wishes to 'trap' messages for particular external OR-address spaces;
- 3) DL-Expansion information: however, this would be quite unusual;
- 4) Expression matches information: if the external system is a non-MHS system with multiple entry points - e.g. a Network address. The Expression Matches information would contain a different OR-address attribute match requirement for each external target MTA.

External OR-address subtrees are normally truncated to summarize the OR-address spaces that they hold routes to. However, they should contain sufficient depth to be able to correctly select the optimum Target MTA for entry for each distinct OR-address space in the external messaging system.

If an Exit MTA has access to two or more alternative external MTAs which lead to the same OR-address space, then the administrator should create a different External OR-address subtree for each alternative and place the two trees on their order of preference.

8.5.7 Distributing Access to External routes through a routing collective

An Exit MTA's routes to external OR-address spaces may be shared with superior and sibling routing collectives. To do this, an internal MTA (i.e. one that must route messages to external OR-address spaces via some Exit MTA) is granted access to read appropriate Exit OR-address subtrees. The information that these contain will be the identity of the external MTA represented as a proxy routing collective. This provides the internal MTA with sufficient information to be able to select an appropriate Exit MTA within the routing collective, and this is sufficient information for the internal MTA to be able to route messages to the appropriate Exit MTA.

Access to external routes may be controlled to prevent unauthorized traffic from using the External Routes specified in External OR-address subtrees. This is achieved by indicating whether the connection groups available to each routing collective can be used as an entry, local exit or transit exit connection group. For instance, if an organization's department sets up a private link to another MD, or to an ADMD financed by the department's funds, the administrators may need to restrict access either to all MTAs in the department's routing collective or to a selection of them.

To make external routes available to all of the MTA's siblings within the context of one of its superior routing collectives (i.e. either its immediate superior or one higher up the routing collective hierarchy) then:

- the base directory entry of the Exit MTA's External OR-address subtree is located as a subordinate to the superior's routing collective entry and
- the access controls are set to grant read access to all subordinates.
- read access may also be restricted to allow only specified subordinates to access the External routes held in an External OR-address subtree.

A superior routing collective may inherit External OR-address subtrees from one or more of its subordinate MTAs. These External OR-address subtrees may contain alternate routes to the same external OR-address spaces. The sequencing of these alternates must be evaluated and sequenced as appropriate by the administrator before they are configured for use by each of the routing collective's subordinate MTAs.

8.5.8 Establishing Routes to non-Routing MTAs or proprietary messaging systems

A routing collective may include messaging systems that are not MHS or MHS routing compliant. There are three main types, each of which must be dealt with in the same way:

- One or more Non-routing MTAs;
- Proprietary electronic messaging systems;
- Other 'standard' telematic systems.

Some of these systems may use different internal addressing structures, i.e. they do not use standard OR-addresses. To deal with these cases, the administrator should:

- 1) represent each one in the directory as a subordinate routing collective;
- 2) allocate them one or more OR-address spaces in the context of the organization's OR-address plan (their internal addressing may be based on Domain Defined Attributes or portions of standard MHS attributes);
- 3) truncate the Reference OR-address subtree at each entry where an allocated OR-address space begins. The target MTA of each entry in the Reference OR-address subtree should contain the identity of the MTA representing the proprietary or non-routing MTA to the rest of the routing collective. Where multiple entry points are available, they should each be modelled as a subordinate routing collective, and the OR-address subtree must have sufficient depth to be able to relate each distinct OR-address space with the appropriate entry routing collective.

8.5.9 Default Routes

A default route should be configured into the base entry of the last subtree read by each MTA. This specifies the default routing information for any message with an OR-address for which no routing information has been found in the rest of the MTA's OR-address subtrees. The target MTA should either have access to information concerning that OR-address, or it should generate an error report or non-delivery diagnostic for those OR-addresses of which it also has no knowledge. There are several approaches to configuring defaults:

- Do nothing; This results in a non-delivery and the originating user will be left to rectify the situation;
- Configure a route to an ADMD, and trust that the ADMD can determine a route to the OR-address space. This can eventually result in a non-delivery, but it represents a low maintenance approach for the administrator;
- Configure a Recipient MD assigned Alternate Recipient that will be able to investigate the cause of the routing failure. This is a high maintenance approach, but also offers a higher quality service.

The default configured for the target MTA itself, i.e. which has been reached as a consequence of other MTA's default routes, should be configured as a Recipient MD Alternative Recipient routing information to ensure that the message's OR-address can be investigated by an MHS administrator.

Default routing information should never be configured in the Reference OR-address subtree. This is because the Reference OR-address subtree and its subordinate secret subtrees should already hold the complete internal routing knowledge for the top-most routing collective and all of its subordinates.

Any OR-address carried in a message that falls within the OR-address space supported by the top level routing collective or any of its subordinates, but for which there is no entry in the Reference OR-address subtree should therefore result in a non-delivery. This can be arranged by configuring each non-leaf entry with a non-delivery instruction that is used if no other instruction is valid for the entry.

8.6 OR-address subtree Entry Routing Information Configuration

Each OR-address subtree contains the routing information for particular OR-addresses. The following subsections outline the procedures for creating and maintaining OR-address subtree entry instructions.

8.6.1 OR-address attribute registration

The organizational and MHS administrators agree on the details of the MHS user, including the OR-address and the MTA on which the user's mailbox will be established. They should check with the registration authority that the newly allocated OR-address is not present in the Reference OR-address subtree to ensure that it has not already been allocated to another user.

Each valid OR-address attribute value should result in creation of an entry in the Reference OR-address subtree. This action is performed by the OR-address Registration Authority to create each leaf and non-leaf entry in the Reference OR-address subtree.

The entry remains empty until it is configured with routing information.

8.6.2 The Target Routing Collective Instruction

A target routing collective name indicates the directory name of the routing collective towards which a message addressed to the OR-address represented by the entry should be transferred. Target routing collective instructions may be configured into any OR-address subtree type. Each entry, with the exception of those which indicate Recipient MD assigned Alternate Recipient or and Alias Redirection, should be configured with a target routing collective name. Within a routing collective’s Reference OR-address subtree, this might be:

- The routing collective which can complete the processing of the message e.g. by delivering it, or DL expanding it (if it represents a distribution list);
- A routing collective which has access to a secret Reference OR-address subtree;
- A routing collective which represents a non-routing MTA or proprietary messaging system gateway modelled as a proxy routing collective;
- An external MTA in another part of the MTS represented as a proxy routing collective.

An instance of the following Table 7 may be used to express each MTA’s routing choices:

Table 7 – MTA routing specification

(1) OR-address subtree base	(2) To OR-address-space	(3) target-routing-collective
(1) The directory name of the OR-address subtree base which will contain the entry, i.e. the OR-address subtree to which the Target Routing Collective Instruction is to be added; (2) The OR-address space which can be reached through the target MTA in (3); (3) The directory name of the routing collective towards which the messages shall be routed. This information forms part of the routingAdvice attribute of the respective OR-address subtree entry (see the oRAddressElement Object Class).		

An instruction may be deleted by either removing the entry, or by replacing it with a non-delivery or the Recipient MD Alternate Recipient instruction.

8.6.3 MHS User Instruction

As far as MHS Routing is concerned, specification of an MHS user is a special case of configuring a target routing collective for a particular OR-address. However, it is only configured into the Reference OR-address subtree or secret Reference OR-address subtrees, and it is associated with other steps to establish the MHS user’s delivery information. To add an MHS user to the Reference OR-address subtree, the administrators should complete the following steps:

- The users entry is configured to contain the identity of the target routing collective (this will normally identify the MHS user’s delivering MTA);
- The entry might also contain some information of local (proprietary) interest to the delivering MTA used in message delivery. This information will be in a proprietary format and will be derived from the documentation of the MTAs or systems supporting the target routing collective;
- The Reference OR-address subtree may be configured to hold one or more alias entries for the user (this is explained in 8.6.4). Each of these represents a common variant of the users personal name, e.g. to represent different variants and orderings of Given names and initials OR-address attributes. These aliases all refer to the user’s preferred entry which, in turn, holds the user’s delivery information;
- As a separate task, the administrator of the delivering MTA (or embedded proprietary system) should create the user’s mailbox either on the MTA or a message store;
- A Double Enveloping instruction may also be configured if required by the organizational security policy;
- It may be necessary to create an entry in the directory to hold the MHS User information, providing the user’s OR-addresses. (This refers to ISO/IEC 12073).

The following procedures should be carried out for the deletion of each MTS user:

- 1) delete the MTA/MS mailbox entry;
- 2) delete the user's subtree entry from the Reference OR-address subtree if it is a leaf entry;
- 3) delete the MHS User entry.

NOTE – MHS users may be configured for any leaf or non-leaf Reference OR-address subtree entry below a country entry. This enables specification of MHS users which represent Management Domains, Organizations, and Organizational Units.

An MHS user entry in an OR-address subtree may be specified as in Table 8.

Table 8 – Specification of an MHS User in the OR-address subtree

(1) Delivering Routing Collective Name	(2) MHS user's OR-address	(3) target-routing-collective name	(4) local-user-identifier
<p>(1) The delivering routing collective directory name;</p> <p>(2) The user's OR-address;</p> <p>(3) The routing collective name of the delivering MTA;</p> <p>(4) The local-user-identifier may carry information to instruct the MTA to deliver to the addressed MTS user. The encoding of this value is of a proprietary nature, and will be specified by the MTA supplier's documentation.</p>			

8.6.4 Aliases

8.6.5 Aliasing Techniques

There are two distinct types of aliasing which have different effects:

- Directory Alias OR-address Instructions for Personal Names; and
- The Alias Redirection Instruction.

These are described in the following two subclauses.

8.6.6 Alias OR-address Instructions for Personal Names

Message originators often use different combinations of Personal Name OR-address elements at different times and specify initials in different sequences to identify personal recipients. To deal with this the administrator should configure a number of alias directory entries for the MHS Common Name, MHS Surname, MHS Given Name, MHS Initials, MHS Generation Qualifier attributes in the Reference OR-address subtree so that an entry exists for any likely combination of Personal Name elements.

This may be accomplished as follows: each Personal Name branch of the tree is examined, and the highest level entry that unambiguously denotes the user is configured to contain the user's actual entry. All of the subordinate entries become alias entries that refer to this object entry. Some examples of modelling Personal Name OR-address elements are given in Figure 9.

The alias technique may be used to resolve cases where a Management Domain inherits multiple country codes, or multiple ADMD codes (where the single space ADMD name is not used) so that OR-addresses only have a single OR-address subtree entry under the preferred country code or ADMD name. All of the other country codes and ADMD names are represented in the directory as aliases.

8.6.7 The Alias Redirection Instruction

The Alias Redirection instruction is used to manage changes of OR-address. Alias redirection instructions may be created by the administrator where one or more MTS users have had all or part of their OR-address changed (e.g. where Organizational Units have merged or changed name).

Alias Redirection should normally only be configured in the Reference OR-address subtree. However, in some circumstances, it may be configured in External OR-address subtrees in order to intercept messages addressed to specific OR-address spaces.