Network Working Group
Request for Comments: 1465

U. Eppenberger
SWITCH

May 1993

Routing Coordination for X.400 MHS Services Within a Multi Protocol / Multi Network Environment Table Format V3 for Static Routing

Status of this Memo

This memo defines an Experimental Protocol for the Internet community. Discussion and suggestions for improvement are requested. Please refer to the current edition of the "IAB Official Protocol Standards" for the standardization state and status of this protocol. Distribution of this memo is unlimited.

1. Introduction

The usage of the X.400 Message Handling System (MHS) is growing rapidly, especially in the commercial world but much interest can also be found in the academic and research community. New networks and new addresses come into use each and every day. The underlying technology for different X.400 networks can vary depending on the transport network and the X.400 MHS implementations used. As a large number of X.400 implementations now support multiple stacks, this offers the chance of implementing a world wide message handling service using the same electronic mail standard and, therefore, without the need of gateways with service reduction and without the restriction to a single common transport network. This, however, leads to several problems for the MHS manager, two of which are:

- Where do I route new X.400 addresses and

- How do I connect to a MHS domain that uses an underlying technology that I do not support.

This document proposes short term solutions to these problems. It proposes a strategy for maintaining and distributing routing information and shows how messages can travel over different networks by using multi stack MTAs as relays. Document formats and coordination procedures bridge the gap until an X.500 directory service is ready to store the needed connectivity and routing information. The format has been designed to allow the information to be stored in an X.500 directory service while managers without directory service access may still use a table based approach.

The routing structure proposed can be applied to a global MHS service

but may also be used at a national level or even within an organisation.

Many experts from IETF X.400-Operations Group and RARE Working Group 1 on Message Handling Systems have read drafts of this document and contributed ideas and solutions. I would especially like to thank Harald Alvestrand, Erik Huizer, Marko Kaittola, Allan Cargille and Paul-Andre Pays.

This is the third version of a table format. The first one was in use within COSINE-MHS for about two years. A second version with major enhancements was then proposed which has been in use for the past year. The third version will probably be the last one before it will be possible to switch to dynamic, directory service based routing.

2. Terminology

MHS community

One or more MHS domains form an MHS community. Mail exchange between these MHS domains is defined by the coordination procedures within this document. Examples of such communities are the Global Open MHS service GO-MHS and the COSINE-MHS service.

MHS domain

One or more MHS subtrees form an MHS domain. This is a purely administrative grouping of MHS subtrees. It is helpful, if someone is responsible for several MHS subtrees, to refer to an MHS domain instead of listing all the subtrees.

MHS subtree

An MHS subtree consists of the total of the mailboxes addressable within a subtree of the X.400 OR address space.

Example: O=SWITCH; P=SWITCH; A=ARCOM; C=CH;

MHS domain of SWITCH in Switzerland, consisting of all mailboxes with O=SWITCH; P=SWITCH; A=ARCOM; C=CH; in the OR address.

RELAY-MTA

An X.400 MTA serving one or several MHS domains. Note that the term WEP -Well Known Entry Point- has been used since the early X.400ies (1987/88) until now, giving the wrong impression of a

single entry point (and therefore a single point of failure). This document proposes to use the term RELAY-MTA, reflecting more clearly the functionality of the MTA.

COSINE-MHS

The COSINE-MHS community is mainly formed by European X.400 service providers from the academic and research area, each of which is a member of RARE. The COSINE-MHS community is used in the annex as an example for the usage of this document in a multinational environment.

3. Requirements

X.400 MTAs can communicate using different transport and network protocol stacks. For this document the stacks used in a WAN environment need to be considered:

	Stack 1	Stack 2	Stack 3	Stack 4
Transport Layer 4	TP0	TP4	RFC1006	TP0
Networkservice 1-3	X.25	CLNS	TCP/IP	CONS

A common protocol stack is not the only requirement to enable communication between two MTAs. The networks to which the MTAs belong need to be interconnected. Some well known networks are listed together with the stacks they use.

Network	Stack	Abbreviation
Public Switched Packet Data Networks	1	Public-X.25
International X.25 Infrastructure EMPB	1,4	EMPB-X.25
US and European connectionless pilot	2	Int-CLNS
Internet	2,3	Internet

Note that several stacks may be supported over a single network. However communication between MTAs is only possible if the MTAs share at least a common stack AND a common network.

Unlike SMTP/TCP/IP systems, there is no directory service available which would allow an MTA to look up the next MTA to which it should submit a message. Routing within X.400 will continue to be table based until a solution using X.500 directory services is available.

Furthermore it is not generally allowed to connect to any MTA even on the same network without being registered on the destination MTA. These restrictions require a large coordination effort and carefully configured and updated systems.

4. Outline of the proposal

This proposal offers a solution for describing information about X.400 message routing within an MHS community in RELAY-MTA and DOMAIN documents. Basic information on the MHS community is documented in the corresponding COMMUNITY document. All contact persons and RELAY-MTA administrators can be found in PERSON documents. A future X.500 based solution may need extended information to overcome still unsolved problems like optimal routing or traffic optimization for messages with multiple recipients. The information collected for the intermediate solution however is very basic. All established coordination procedures will help and even speed up the future introduction of an X.500 based solution.

4.1 The COMMUNITY document

For each MHS community there exists one single COMMUNITY document containing basic information. First the contact information for the central coordination point can be found together with the addresses for the file server where all the documents are stored. It also lists network names and stacks to be used in the RELAY-MTA and DOMAIN documents. An MHS community must agree on its own set of mandatory and optional networks and stacks.

4.2 The RELAY-MTA document

Every MHS domain in the community may designate one or more MTAs as RELAY-MTAs. These RELAY-MTAs accept incoming connections from the RELAY-MTAs of the other MHS domains and in return are allowed to send messages to these RELAY-MTAs. A RELAY-MTA is documented with all the necessary connection information in the corresponding RELAY-MTA document.

4.3 The DOMAIN document

An MHS domain has a responsible person who sets up the routing entries for the domain in the DOMAIN document. The primary RELAY-MTAs listed in the DOMAIN document as serving this MHS domain must, TOGETHER, offer at least connectivity to all networks and stacks listed as mandatory in the COMMUNITY document. Optional RELAY-MTAs may be added, generally with higher priority, to allow more precise routing.

An MHS domain may also decide not to operate a RELAY-MTA. It will then only need agreements with one or more RELAY-MTAs from other MHS services which will relay for this domain. The domain itself, however, must either create its own DOMAIN document or document its MHS subtrees jointly with another MHS domain.

The structure of the DOMAIN document is very straightforward. It starts off with one or more MHS subtrees, each on its own line. After the domains follows a line indicating the responsible person for the MHS subtrees mentioned. Finally the responsible RELAY-MTA(s) are listed with appropriate priorities.

4.4 The PERSON document

All administrators and responsible persons are documented in PERSON documents. The RELAY-MTA and DOMAIN documents contain just keys pointing to a PERSON document. If such a person can already be found in an X.500 directory service, then the key consists of a Distinguished Name, else the key is just its OR address.

4.5 Coordination

This approach requires an identified coordination point. It is up to the MHS community to decide on the level of coordination and support to be provided and on the funding mechanisms for such activities. Basic information can be found in the COMMUNITY document. The following list of support activities is considered mandatory for an operational service:

- New RELAY-MTAs joining the service are tested and support is given to create the RELAY-MTA document.

- New MHS domains joining the MHS community get assistance to set up RELAY-MTA(s) and/or find appropriate RELAY-MTA(s) and to create DOMAIN documents.

- Updated documents are announced to the RELAY-MTA managers and responsible persons for the DOMAIN documents unless automatic distribution is used.

- All the RELAY-MTA, DOMAIN and PERSON documents are made available on a file server together with the COMMUNITY document. The file server must at least be reachable via email. MHS communities with a big number of documents may consider additional access methods like ftp and FTAM.

- Tools should be made available to manage routing tables for the X.400 software used on the RELAY-MTAs or to fill in and check the documents. The format of the documents has specifically been chosen to enable the use of automated tools.

The RELAY-MTA managers must be aware that a large number of RELAY-MTAs in an MHS community may require significant operational resources to keep the local routing tables up-to-date and to

constantly monitor the correct functioning of the connections. On the other hand more than one RELAY-MTA with a good connectivity to an MHS domain improves the overall robustness of the domain and thus the QOS.

MHS communities may decide on additional mandatory requirements for the operation of a RELAY-MTA. These may include a hot line, echo services, exchange of statistics, response time to problem reports, uptime of the RELAY-MTA, etc. This will ensure a certain quality of service for the end users.

4.6 Routing

The proposal addresses MHS communities spanning several organisations. But it may also be used to manage routing within a single organisation or even a global MHS community.

Two kinds of mail relays are defined, the primary RELAY-MTAs and the secondary RELAY-MTAs. A primary or secondary RELAY-MTA must allow incoming connections from all other primary and secondary RELAY-MTAs with a common stack. Primary RELAY-MTAs must be able to connect to all other primary RELAY-MTAs which share a common stack. A secondary RELAY-MTA must connect to at least one primary RELAY-MTA.

Each MHS community must define update procedures for the routing based on the documentation. Automated update has to be studied carefully.

An MHS community should also define procedures for new RELAY-MTAs and MHS domains joining the service. Since the usage of X.400 is growing rapidly a flexible but well coordinated way of integrating new members into an MHS community is needed. The proposed documentation format supports this by allowing primary and secondary RELAY-MTAs. All RELAY-MTAs accept incoming connections from each other. Sending messages can be done by using the primary RELAY-MTAs only. This allows new RELAY-MTAs to join the community as secondary and to get primary status when traffic flow increases. Secondary RELAY-MTAs may also require a longer testing period.

5. The documents

The definition is given in BNF-like syntax. The following conventions are used:

| | means choice

\ is used for continuation of a definition over several lines

[] means optional

{} means repeated one or more times

() is used to group choices

\" is used for double quotes in a text string

<CR> is a Carriage Return and means that the next section starts on its own line.

The definition is complete only to a certain level of detail. Below this level, all expressions are to be replaced with text strings. Expressions without more detailed definition are marked with single quotes '. The format and semantics should be clear from the names of the expressions and the comments given.

Wherever the BNF definition requires a single blank, multiple blanks may be used to increase the readability. Please note that for some field values the number of spaces is significant.

Lines exceeding 80 characters should be wrapped at any convenient blank except at blanks which are significant. The line is continued with at least one leading blank.

Comments may be placed anywhere in the document but only on separate lines and without splitting wrapped lines. Such a comment line must either start with a '#' sign followed by white space and the comment or consist of a single '#' on a single line.

The documents must follow the case of the strings defined in BNF. Note that some values, especially connection parameters like TSEL or MTA password are case dependant too.

The BNF definitions are ordered top-down. See Appendix B for an alphabetically sorted list.

A set of one COMMUNITY document and several RELAY-MTA, DOMAIN and PERSON documents belong together. The detailed definitions can be found in the following chapters.

```
<X.400 routing coordination document set> ::= \
                        <COMMUNITY-document> \
                        { <RELAY-MTA-document> } \
                        { <DOMAIN-document> } \
                        { <PERSON-document> }
```

5.1 Common Definitions

<DirectoryName> ::= 'Distinguished Name'

The string representation of a Distinguished Name is defined in the RFCxxxx. If a Distinguished Name is used as a key in the documents, then the information can be fetched from the directory instead of checking the appropriate document. But as long as not all managers in the same community have directory access, the same information must also be present in a document. Note that Distinguished Names in the context of the routing documents are just used as key strings to point to other documents.

<Community-Identifier> ::= "Community: " \
('community name' | <DirectoryName>) <CR>

The 'community name' is a string identifying the MHS community to be used in the first line of all documents.

<UniqueRELAY-MTAkey> ::= ((["P=" 'PRMDname' "; "] \
["A=" 'ADMDname' "; "] \
"C=" <Country-Code> "; " \
"MTAname=" 'MTAname')
| <DirectoryName>)

A unique key is needed to identify the RELAY-MTA. In addition to the MTA name itself, it is proposed to use OR address attributes of the management domain where the RELAY-MTA resides. ADMD and PRMD fields are both optional and may be used to guarantee uniqueness of the key. The values used are irrelevant. Even non-printable characters like @ or ! are acceptable. The result is not an address but a key string. A Distinguished Name may be used instead.

<UniquePersonKey> ::= (<X.400 address> | <DirectoryName>)

A unique key is necessary to make the links from the documents where a responsible person or an administrator is needed, to the PERSON documents. It is either the OR address of the person or a Distinguished Name (if the person is already registered in the directory).

<Country-Code> ::= 'Two Character Country Code ISO-3166'

<X.400 address> ::= 'OR address, ISO 10021-2 Annex F'

It has been used consequently all over the document. This explains also the syntax of the

```
            <UniqueRELAY-MTAkey> and the <MHS-subtree>. Examples:
            S=user; O=org ltd.; OU1=sect1; P=org; A=rel400; C=aq;
            DDA:RFC-822=we(a)sell.it; P=internet; A= ; C=xx;
            G=john; I=w; S=doe; P=org; A=rel400; C=aq;

      <EMail> ::= "Address: " <X.400 address> <CR>

      <tel-no-list> ::= <tel-number> [{"; " <tel-number>}]

      <tel-number> ::=  {"+" <int-pref> " " <national number> \
                           [" x" <extension>]}
            This syntax follows the attribute syntax of the X.500
            directory based on CCITT E.123.

      <int-pref> ::= 'international prefix'

      <national number> ::= 'national telephone number'
            A national number may be written with spaces and
            hyphens to group the figures.

      <extension> ::= 'local extension'

      <Phone> ::= "Phone: " <tel-no-list> <CR>
            One or more phone numbers

      <Fax> ::= "Fax: " <tel-no-list> <CR>
            One or more FAX numbers

      <Mail> ::= "Mail: " 'postal address information' <CR>
            The items of the postal address are separated by ' /'
            wherever the next item goes onto the next line for a
            printed address label.  If the document is for an
            international community, the address should include the
            person's country.
            Example:
            Mail: SWITCH Head Office / Urs Eppenberger /
                  Limmatquai 138 / CH-8001 Zurich / Switzerland
            results in the following mailing label:
            SWITCH Head Office
            Urs Eppenberger
            Limmatquai 138
            CH-8001 Zurich
            Switzerland

      <Update-info> ::= "Update: FORMAT=V3; DATE=" 'yymmdd' \
                           "; START=" 'yymmdd' \
                           ["; END=" 'yymmdd'] <CR>
            The <Update-info> contains also the format identifier.
```

The date of the last update of a document is given in the form 'yymmdd'.
A start date must be set. A document can be published this way before the information in it is valid. (This is especially useful in absence of automated tools. RELAY-MTA managers get more time to prepare their systems.)
An end date is used to set an expiration date for the document.

<P-address> ::= 'String encoded Presentation Address'
The format of this string follows RFC1278, A string encoding of Presentation Address and RFC1277, Encoding Network Addresses to support operation over non-OSI layers. See chapter 5.2 about the usage of macros in a Presentation Address.

<Service-type> ::= <Network-name> "/" \
<Network-service> "/" \
<Transport-Protocol>
The service type consists of a string with three parts concatenated with a "/": Network-name/Network-service/Transport-Protocol.

<Network-name> ::= 'Name of a network'
The network-name string identifies a network. A well known key word should be chosen. (No '/' character is allowed.)
Examples: Public-X.25, Internet, EMPB-X.25, Int-CLNS, WIN, Janet,

<Network-service> ::= 'Name of a network service'
Examples: X.25, CONS, CLNS, TCP

<Transport-Protocol> ::= 'Name of a transport protocol'
Examples: TP0, TP2, TP4, RFC1006

Since network and stack information forms one string, it identifies in an easy way a connection possibility between two RELAY-MTAs. The COMMUNITY document defines the strings to be used in the RELAY-MTA and DOMAIN documents. Some examples:
Internet/TCP/RFC1006
Public-X.25/X.25/TP0
RARE-IXI/CONS/TP0
RARE-CLNS/CLNS/TP4
(It is probably important to mention here that this string has nothing to do with the format of a

presentation address as defined by Steve Hardcastle-Kille in RFC1278. The problem of networks using the same address structure (X.121 DTEs, 4 Byte Internet addresses) but not being connected is not addressed in RFC1278 but solved by using the proposed service identifier above in addition to the presentation address. As long as there are network islands, there is no other way than the addition of an 'island'-identifier.

```
<MHS-subtree> ::= ["O=" 'Organization-name' "; "] \
                      ["OU1="'OrganizationalUnit'"; "\
                      ["OU2=" 'OrganizationalUnit' "; " \
                      ["OU3=" 'OrganizationalUnit' "; " \
                      ["OU4=" 'OrganizationalUnit' "; "]]]] \
                      ["P=" 'PRMDname' "; "] \
                      "A=" 'ADMDname' "; " \
                      "C=" <Country-Code> ";"

<Operation> ::= "Reachable: "  {<time> "-" <time> "; "} \
                      <Time-zone> <CR>

<time> ::= 'hh:mm'

<Time-zone> ::= ("UTC+" | "UTC-") 'hhmm'
```

The operation information is needed to know the time someone is reachable. This information is important for communities spanning several time zones.
'hhmm' is a four digit value, the first two digits indicate hours, the second two digits indicate minutes.
Use "UTC+" for time zones east of Greenwich. A simple formula helps to calculate the current time at the remote place:
local-time - local-displacement + remote-displacement = remote-time
18:00 - (UTC + 0100) + (UTC - 0800) = 09:00
The <Time-zone> entry may be followed by a comment line indicating when Daylight Saving Time is in effect. This is especially reasonable for MHS communities spanning continents on the northern and southern hemisphere.

5.2 The COMMUNITY document

```
<COMMUNITY-document> ::= <Community-Identifier> \
                      <Update-info> \
                      <COMMUNITY-document-body>
```

The first line of the COMMUNITY document specifies the

<Community-Identifier> to be used in the header of all other documents belonging to the same community. It is recommended to add a few comment lines to describe the members of this MHS community.

<COMMUNITY-document-body> ::= <Coordination> \
[{<Macro-definition>}] \
{<Connections>}

<Coordination> ::= <EMail> <Phone> <FAX> \
<Mail> <Operation> <File-server>
Set of contact information for the coordination point

<File-server> ::= <email-server> [{<FTP-server>}] \
[{<FTAM-server>}]
All documents must be available at least to the managers of the MHS domains and the RELAY-MTAs through an email server. If FTAM and FTP are also available, the generation of automated update tools is much easier.
It is suggested to have a single server. If there is only one, knowing a single X.400 OR address will allow you to reach the server. However for FTP and FTAM more system addresses may be possible depending on the number of available network connections (or service types as they are called in this document).

<email-server> ::= "Mail-server: "<X.400 address> <CR>
The email address of the file server.

<FTP-server> ::= "FTP-server: " 'domain name' "; " \
'account-name' ["; " 'password'] <CR>
In addition to the domain name of the server, an account name and a password is given. In most cases this will probably be something like "anonymous" and "guest".
Some servers request the RFC822 address of the user. This is documented by using the string 'user@domain' as password entry. The meaning is not to use 'user@domain' literally as password while accessing the server (even if this would generally work too since the software often just checks the presence of an @ sign,) but to use ones own RFC822 email address.

<FTAM-server> ::= "FTAM-server: " <P-address> "; "\
'account-name' ["; " 'password'] \
["; X.500 " <DirectoryName>] <CR>
The account name is often case sensitive. Some FTAM

servers offer anonymous access with the account-name ANON. Documenting an FTAM server with a Distinguished Name is only allowed if the server is registered in the directory.

<Macro-definition> ::= "Macro: " 'Macro name' " " \
'Macro value' <CR>

Presentation addresses without the usage of macros are generally unreadable. RFC1278 suggests a few macros. All macros which are allowed in a community must be defined in the COMMUNITY document. It is recommended to use the proposed macros in RFC1278 and add new ones if necessary:

```
Macro: Int-X25(80)        TELEX+00728722+X.25(80)+01+
Macro: Janet-X25(80)      TELEX+00728722+X.25(80)+02+
Macro: Internet-RFC-1006  TELEX+00728722+RFC-1006+03+
```

<Connections> ::= {<mandatory-service>} \
{[<optional-service>]}

Note that at least one mandatory service type is needed.

<mandatory-service> ::= "Mandatory-Service: " \
<Service-type> <CR>

<optional-service> ::= "Optional-Service: " \
<Service-type> <CR>

5.3 The RELAY-MTA document

<RELAY-MTA-document> ::= <Community-Identifier> \
<Update-info> \
<RELAY-MTA-document-Identifier> \
<RELAY-MTA-document-body>

A RELAY-MTA document contains the full description of a single RELAY-MTA. Only one community is allowed. Since some of the information is community dependent, it would not be easily possible to have a single RELAY-MTA document used in different communities.

<RELAY-MTA-document-Identifier> ::= \
"RELAY-MTA: " <UniqueRELAY-MTAkey> <CR>

<RELAY-MTA-document-body> ::= <Status> <connection-info> \
<contact-info>

<Status> ::= "Status: " ("primary" | "secondary") <CR>

This defines if the RELAY-MTA has 'primary' or

'secondary' status. See section 4.3 and 6 for more information.

```
<connection-info> ::= <password> <RTS> \
                      {<called-connection><calling-connection>}\
                      [<system>] \
                      [<local-domain>] \
                      [<echo-server>]
```

More than one set of connection information may be present for RELAY-MTAs supporting several networks and protocol stacks.

```
<password> ::= "Password: " \
                      ("secret" | "none" | \
                      "value=\"" 'password' "\"") <CR>
```

If the keyword none is present, then no password is sent with the MTAname when this MTA initiates an RTS connection or responds to an incoming connection.
Password: none

If the keyword secret is present, then the connection needs a password which is not made publicly available. (For example, a community might keep a list of the passwords at the central coordination point. The list would then be faxed to the RELAY-MTA managers.)
Password: secret

A password must be documented using the value="password" notation. The double quotes around the password are needed, consider the case of a single blank as a password.
Password: value=" "
Password: value="nume-n-ine"

```
<RTS> ::= <dialog-mode> \
                      [<checkpoint-size> <window-size>]

<dialog-mode> ::= "RTS-dialog-mode: " \
                      ("TWA" | "MONOLOGUE") <CR>

<checkpoint-size> ::= "RTS-checkpoint-size: " \
                      'checkpoint size' <CR>

<window-size> ::= "RTS-window-size: " \
                      'window size' <CR>

<called-connection> ::= "Called-address: " \
                      <Service-type> "; " \
```

```
                        <P-address> "; " <MTS> \
                        ["; " <Service-priority>] <CR>

      <MTS> ::= "MTS-T" | "MTS-TP" | "MTS-TP-84"
                MTS-T:     mts-transfer
                MTS-TP:    mts-transfer-protocol
                MTS-TP-84: mts-transfer-protocol-1984
                See ISO 10021-6, Section 3, chapter 11.1 for more
                details on this matter.  X.400(84) systems only support
                mts-transfer-protocol-1984.

      <Service-priority> ::= 'Integer 0..99'
                The lowest Integer corresponds to the highest priority
                as in DNS.  It is possible to set different priorities
                for each service type.  This may be chosen, for
                example, to distribute the load amongst different
                networks according to their available bandwidth.

      <calling-connection> ::= "Calling-address: " \
                        <Service-type> "; " \
                        <P-address> <CR>
                Since called and calling network addresses may differ
                in certain configurations and some X.400 systems do
                validation on calling network addresses, it is
                important to have this information in the RELAY-MTA
                document.  (Note: a calling X.121 address might change
                if the X.25 switch is reconfigured.  This will stop a
                RELAY-MTA from connecting to other RELAY-MTAs using
                address validation without having changed anything at
                the higher layers!)

      <system> ::= "System: HW=" 'computer type' "; " \
                        "OS=" 'operating system' "; " \
                        "SW=" 'MHS  software' <CR>
                It is optional to provide HW/SW information.
                Experience, however, has shown that a number of
                communication problems were more easily identified and
                solved with this information present and up-to-date.

      <local-domain> ::= "LocalDomain: " <MHS-subtree> <CR>
                This is a useful but optional extension to the
                documentation.
                The <MHS-subtree> is local to the RELAY-MTA.  The <MHS-
                subtree> attributes might be used together with
                S=nosuchuser; to do connectivity and availability
                tests.
```

<echo-server> ::= "EchoServer: " <X.400 address> <CR>
Some of the RELAY-MTAs might offer an echo server functionality. It does make sense to document this in the RELAY-MTA document for test purpose. This field is optional.

<contact-info> ::= {"Administrator: " <UniquePersonKey> <CR>}
The contact details for the RELAY-MTA administrator can be found in the appropriate PERSON document. It is possible to document a whole team using a distribution list if this is desired. It is generally better to document one or more 'real' persons.

5.4 The DOMAIN document

<DOMAIN-document> ::= <Community-Identifier> \
<Update-info> \
<DOMAIN-document-body>

<DOMAIN-document-body>::= {<Domain-entry>} <responsible> \
{<Relay>}

<Domain-entry> ::= "Domain: " <OR-matching> <MHS-subtree> <CR>
Note that it is not allowed to have equal <Domain-entry> lines in different DOMAIN documents belonging to the same MHS community. A Domain-entry line can only appear in one DOMAIN document.

<OR-matching> ::= ("* " | "= ")
This qualifier defines how the following OR address attributes should be handled for the routing algorithm.
If a '*' is present, a destination address of a message is matched by the "Domain:" entry if at least the OR address attributes in the "Domain:" entry are equal to the destination address.
If a "=" is present, a destination address of a message is matched by the "Domain:" entry if there are exactly the same OR attributes in the destination address as in the "Domain:" entry. (This restriction works for OU4, OU3, OU2, OU1, O, P, A and C only.)
Example:
a) Domain: * P=switch; A=arcom; C=ch;
b) Domain: = P=switch; A=arcom; C=ch;
The address S=eppenberger; P=switch; A=arcom; C=ch;
matches both cases, a) and b).
The address S=eppenberger; O=unibe; P=switch; A=arcom;
C=ch; matches only case a).

```
<responsible> ::= {"Administrator: " <UniquePersonKey> <CR>}
```
This is the person responsible for the listed domains. His task is to get the agreement of the relaying RELAY-MTAs and keep the DOMAIN document up-to-date. This person is the only one authorized to make changes to this document. Note that multiple administrators may be listed.

```
<Relay> ::=         "Relay: " \
                      ( 'UniqueRELAY-MTAkey' | \
                      "Internet-SMTP" ) "; " \
                      <RELAY-MTA-Priority> <CR>
```
The priority is used to define the sequence in which different RELAY-MTAs may be tried in case of failure. A lower integer corresponds to a higher priority as in DNS. Priorities 0..49 are used to indicate backup RELAY-MTAs. Priorities 50..99 are used for RELAY-MTAs not acting as backup but as relay service provider for a network service type not supported by the main RELAY-MTA.
The keyword "Internet-SMTP" is a placeholder for an RFC1327 gateway connected to Internet. The RELAY-MTA manager selects a gateway of his choice.

```
<RELAY-MTA-Priority> ::= <Integer 0..99>
```

5.5 The PERSON document

```
<PERSON-document> ::= <Community-Identifier> \
                      <Update-info> \
                      <PERSON-document-identifier> \
                      <PERSON-document-body>

<PERSON-document-identifier> ::= "Key: " <UniquePersonKey> <CR>

<PERSON-document-body>::= <Name> {<EMail>} {<RFC822>} \
                      <Phone> <Fax> <Mail> <Operation>

<Name>  ::= "Name: " 'name of person' <CR>
```
The name of the person is given. The issue of the character set problem is not addressed in this document. Especially international communities should restrict themselves to IA5 or ASCII.

```
<RFC822> ::= "RFC822: " <RFC-822-address> <CR>
```
This is the RFC-822 address of the person. It is often a big help to know the RFC822 address of someone, for example if the X.400 system is not reachable. This is

also the reason why it is possible to provide multiple OR and RFC822 addresses. The first one is considered the primary one.

6. Routing rules

All the users within the MHS community have the right to send messages to each other. The general agreement is that the RELAY-MTA infrastructure is used according to the following routing rules. More direct connections based on bilateral agreements are fully accepted.

A primary or secondary RELAY-MTA must allow incoming connections from all other primary and secondary RELAY-MTAs with a common stack. Primary RELAY-MTAs must be able to connect to all other primary RELAY-MTAs which share a common stack. A secondary RELAY-MTA must connect to at least one primary RELAY-MTA.

A message arriving at a RELAY-MTA must either be sent to the next RELAY-MTA based on the DOMAIN documents of the MHS community or it is sent to an MTA closer to the destination based on local routing decisions. The following algorithm must be used when forwarding a message to the next RELAY-MTA:

1) Select the relevant DOMAIN document by searching for a match of the Recipient address in the message with the entries in the document.

If your own RELAY-MTA appears in this list, this indicates one of the following:

- You offered relay services for another RELAY-MTA with higher priority. Continue with step 2 to decide on the next RELAY-MTA.

- Your RELAY-MTA is the final destination according the DOMAIN document of your community. You need to forward the message to the final destination according local routing information.

2) From the list of RELAY-MTAs select those that have at least one common network service type with your own RELAY-MTA.

3) Now delete all secondary RELAY-MTAs from the list where no direct connection is desired. For remaining RELAY-MTAs in the list no difference is made anymore between primary and secondary status.

4) Select from this reduced set of RELAY-MTAs the one with the highest RELAY-MTA-priority. If there is more than one RELAY-MTA

having the same priority, then select a RELAY-MTA of your choice. If your own RELAY-MTA appears in that list, then you are not allowed to send to a RELAY-MTA with lower or equal priority.

5) If there are no service-priorities set in the corresponding RELAY-MTA document indicating which service type to use, you are free to choose the service type for connecting to the RELAY-MTA. However, if there are specific priorities set then select the service type with the highest priority.

6) If the connection fails then try other service types in the sequence of descending priority.

7) If no connection works for the selected RELAY-MTA, then check in the list for the one with the same priority, if possible, or else select one with the next lower priority. If there is another RELAY-MTA with a RELAY-MTA-priority between 0..49, then select it and proceed at step 5. Without another RELAY-MTA to try the currently selected RELAY-MTA will be retried.

6.1 How to use RELAY-MTA-priorities

An example helps to explain the usage of RELAY-MTA-priorities. Internet/TCP/RFC1006 and Public-X.25/X.25/TP0 are mandatory service types in the community REMOTEmail. The MHS domain P=REMOTE; A=ARCOM; C=CH; operates MTA-B, only connected to public X.25. A second RELAY-MTA which is connected to both, Internet and public X.25 is needed to offer relay services. A connection using Internet is considered cheap in this example. Both service types are available at MTA-A. Since MTA-B is the only RELAY-MTA responsible for relaying messages to P=REMOTE; A=ARCOM; C=CH; to the final destination it must have the highest priority.

```
Community: REMOTEmail
Domain: * P=REMOTE; A=ARCOM; C=CH;
RELAY-MTA: P=REMOTE; A=ARCOM; C=CH;MTAname=MTA-B; 20
RELAY-MTA:  P=MTA-C; A=ARCOM; C=CH;MTAname=MTA-C; 80
```

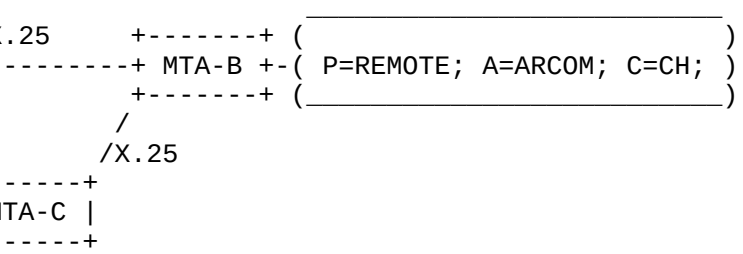

If MTA-A needs to relay a message for P=REMOTE; A=ARCOM; C=CH; then the rules of chapter 6 are evaluated:

1. MTA-B and MTA-C are possible destinations

2. MTA-B and MTA-C are reachable from MTA-A

3. MTA-B is a primary RELAY-MTA (not relevant in this example)

4. MTA-B has the highest priority.

5. MTA-B doesn't have specific service type lines documented. MTA-A chooses public X.25, since it is the only possibility to connect to MTA-B.

6. No other service types are available if the connection fails.

7. MTA-C has a priority of 80, it is not a backup RELAY-MTA. MTA-A must spool the message and try to connect directly to MTA-B.

The organisation running MTA-A could save money by sending messages for the MHS domain P=REMOTE; A=ARCOM; C=CH; via MTA-C. Since the connection between MTA-C and the destination MTA-B is also expensive, the organisation running MTA-C would have to pay for external relay traffic. Setting a lower priority for MTA-C forces the other RELAY-MTAs with public X.25 connectivity to take their share of the cost.

Note that forcing other RELAY-MTAs to use a specific stack should be avoided wherever possible by offering RELAY-MTAs with equal priority for all mandatory network services. This can be an important financial issue for MHS communities spanning several organisations, it is not relevant in general for an MHS community within a single organisation.

6.2 How to use RELAY-MTA-priorities for backup RELAY-MTAS

Two RELAY-MTAs offer real backup connectivity for the MHS domain P=REMOTE; A=ARCOM; C=CH;. It is therefore possible to set RELAY-MTA priorities in the range of 0..49 for both RELAY-MTAs. MTA-B will be the preferred one since it has the higher priority. If the connection to MTA-B fails, a sending RELAY-MTA may immediately try to connect to MTA-C.

```
Community: REMOTEmail
Domain: * P=REMOTE; A=ARCOM; C=CH;
RELAY-MTA: P=REMOTE; A=ARCOM; C=CH;MTAname=MTA-B; 10
```

```
RELAY-MTA: P=REMOTE; A=ARCOM; C=CH;MTAname=MTA-C; 30
```

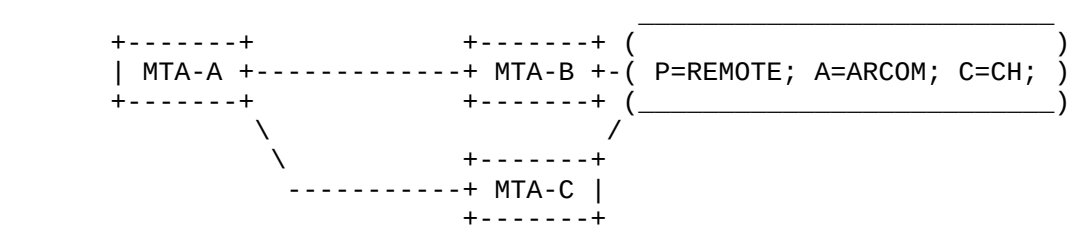

6.3 Load Sharing

It is possible to set equal priorities to do some sort of load sharing. However, most implementations are not able to do random selection of RELAY-MTAs with equal priorities but have a fixed configuration. If load sharing is really needed then it is suggested to split up the MHS domain into several MHS subtrees and document them separately with a set of RELAY-MTAs with different priorities.

An example is provided for illustration of the first possibility with equal RELAY-MTA-priorities:

```
Community: REMOTEmail
Domain: * P=REMOTE; A=ARCOM; C=CH;
RELAY-MTA: P=REMOTE; A=ARCOM; C=CH;MTAname=MTA-B; 10
RELAY-MTA: P=REMOTE; A=ARCOM; C=CH;MTAname=MTA-C; 10
```

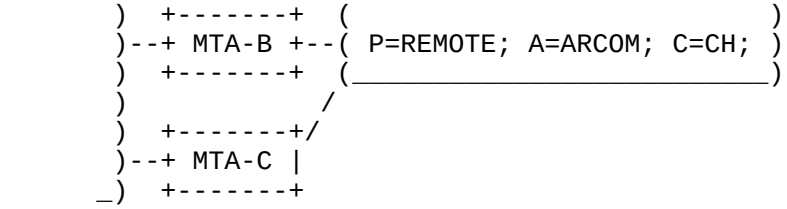

And here is an example where the MHS domain C=CH;ADMD=ARCOM;PRMD=REMOTE;O=Big-Org is documented with its own DOMAIN document: Note that in this example both RELAY-MTAs serve as backup RELAY-MTAs.

```
Community: REMOTEmail
Domain: * P=REMOTE; A=ARCOM; C=CH;
RELAY-MTA: P=REMOTE; A=ARCOM; C=CH;MTAname=MTA-B; 10
RELAY-MTA: P=REMOTE; A=ARCOM; C=CH;MTAname=MTA-C; 30

Community: REMOTEmail
Domain: * O=Big-Org; P=REMOTE; A=ARCOM; C=CH;
```

```
RELAY-MTA: P=REMOTE; A=ARCOM; C=CH;MTAname=MTA-C; 10
RELAY-MTA: P=REMOTE; A=ARCOM; C=CH;MTAname=MTA-B; 30
```

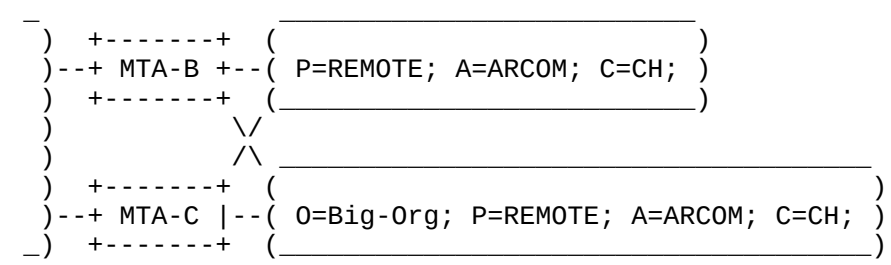

7. Open issues

Currently there are about 35 RELAY-MTAs within the COSINE-MHS service. A rough guess is that a network with about 60 RELAY-MTAs is still manageable provided there are automated tools for MTA configuration. If there are more MTAs applying for RELAY-MTA status in an MHS community, then either an X.500 based solution should be ready or a core set of about 30 well operated super-RELAY-MTAs should form a backbone, documented within a specific MHS community.

Since the RELAY-MTA document contains information about the supported X.400 version (84 or 88), it is possible for an X.400(88) system to select with higher priority an (88)RELAY-MTA. The rules in chapter 6 could be modified to select X.400(88) systems first if the sending RELAY-MTA is an (88) system itself. The issue of how to establish an X.400(88) RELAY-MTA infrastructure within an MHS community is for further study.

Appendix A: Document examples for the COSINE-MHS community

Instead of creating artificial documents to show an example document set, this appendix contains extracts from a real operational X.400 service. The research and development community in Europe has used X.400 for several years. This proposal was initially written to address this community only and solve the urgent routing and coordination problems. Contributions from different experts have made it more flexible and therefore hopefully useful for other MHS communities.

Appendix A1: The COMMUNITY document

```
Community: COSINE-MHS
# The COSINE-MHS service is a MHS community formed by the European
# academic and research networks with additional contacts in all
# other continents.
#
# The coordination is done by the COSINE-MHS project team.
#
Update: FORMAT=V3; DATE=921218; START=930201
#
Address: S=Project-Team; O=SWITCH; P=SWITCH; A=ARCOM; C=CH;
#
Phone: +41 1-262-31-43
Fax:   +41 1-261-81-88
#
Mail:  SWITCH Head Office /
       MHS Coordination Service /
       Limmatquai 138 /
       CH-8001 Zurich /
       Switzerland
#
Reachable: 09:00-12:00; 14:00-17:30; UTC+1
#
Mail-server: S=mhs-server; O=switch; OU1=nic;
             P=SWITCH; A=ARCOM; C=CH;
FTP-server:  nic.switch.ch; cosine; user@domain
#
Macro: Int-X25(80)        TELEX+00728722+X.25(80)+01+
Macro: Internet-RFC-1006  TELEX+00728722+RFC-1006+03+
Macro: IXI                TELEX+00728722+X.25(80)+06+
#
Mandatory-Service: Public-X.25/X.25/TP0
# The public X.25 network.  X.25 is supported in most X.400
# applications and mandatory in X.400 anyhow.
#
Mandatory-Service: Internet/TCP/RFC1006
```

```
  # The Internet, standing for the global TCP/IP network of the
  # research and development community
  # RFC1006 is considered a solution to ease migration to OSI. It will
  # be replaced by TP4/CLNS as soon as a reliable service is
  # available.
  #
  Optional-Service: Int-CLNS/CLNS/TP4
  # The RARE Connectionless pilot service.  Current participants are
  # NORDUnet, SURFnet, CERN, NSFnet and SWITCH.
  #
  Optional-Service: EMPB-X.25/X.25/TP0
  # The International X.25 Infrastructure covering most countries in
  # Europe.  The absence of volume tariffs make it a preferred choice.
```

Appendix A2: Example of a RELAY-MTA document

```
  Community: COSINE-MHS
  #
  Update: FORMAT=V3; DATE=921218; START=930201
  #
  RELAY-MTA: P=SWITCH; A=ARCOM; C=CH; MTAname=chx400.switch.ch
  #
  Status: primary
  #
  Password: none
  RTS-dialog-mode: MONOLOGUE
  #
  Called-address:  Public-X.25/X.25/TP0;
                   "591"/Int-X25(80)=22847971014520+CUDF+03010100;
                   MTS-TP-84
  Calling-address: Public-X.25/X.25/TP0;
                   Int-X25(80)=22847971014520;
  #
  Called-address:  Internet/TCP/RFC1006;
                   "591"/Internet-RFC-1006=chx400.switch.ch;
                   MTS-TP-84
  Calling-address: Internet/TCP/RFC1006;
                   Internet-RFC-1006=chx400.switch.ch
  #
  Called-address:  EMPB-X.25/X.25/TP0;
                   "591"/IXI=20432840100520+CUDF+03010100;
                   MTS-TP-84
  Calling-address: EMPB-X.25/X.25/TP0;
                   IXI=20432840100520;
  #
  Calling-address: Int-CLNS/CLNS/TP4;
                   "591"/NS+39756F1111111101000014560AA00040005E100;
                   MTS-TP-84
```

```
  Called-address:  DCC+756+x11111111010000014560AA00040005E100
  #
  # For X.400(88) over CLNS
  #
  Calling-address: Int-CLNS/CLNS/TP4;
                   "592"/NS+39756F11111111010000014560AA00040005E100;
                   MTS-T
  Called-address:  DCC+756+x11111111010000014560AA00040005E100
  #
  System: HW=SUN 4/690MP; OS=SunOS 4.1.1; SW=PP-6.0
  #
  LocalDomain: O=switch; OU1=chx400; P=switch; A=arcom; C=ch;
  #
  EchoServer:  S=echo; O=switch; OU1=chx400; P=switch; A=arcom; C=ch;
  #
  Administrator: CN=Felix Kugler, O=SWITCH, C=CH
  Administrator: CN=Christoph Graf, O=SWITCH, C=CH
```

Appendix A3: Example of a DOMAIN document

```
  Community: COSINE-MHS
  #
  Update: FORMAT=V3; DATE=921218; START=930201
  ##
  Domain: *     P=SWITCH; A=ARCOM; C=CH;
  Domain: *     P=SANDOZ; A=ARCOM; C=CH;
  Domain: *        P=ABB; A=ARCOM; C=CH;
  Domain: *        P=UBS; A=ARCOM; C=CH;
  Domain: *      P=ISREC; A=ARCOM; C=CH;
  Domain: *    P=ALCATEL; A=ARCOM; C=CH;
  Domain: *        P=ITU; A=ARCOM; C=CH;
  Domain: * P=OSILABMAIL; A=ARCOM; C=CH;
  Domain: *        P=WHO; A=ARCOM; C=CH;
  Domain: *       P=CERN; A=ARCOM; C=CH;
  Domain: *   P=CERBERUS; A=ARCOM; C=CH;
  #
  Administrator: CN=Christoph Graf, O=SWITCH, C=CH
  Administrator: S=postmaster; O=SWITCH; P=SWITCH; A=ARCOM; C=CH;
  #
  RELAY-MTA: P=SWITCH; A=ARCOM; C=CH; MTAname=chx400.switch.ch; 0
  #
  RELAY-MTA: P=SWITCH; A=ARCOM; C=CH; MTAname=vms.switch; 10
```

Appendix A4: Example of a PERSON document

```
  Community: COSINE-MHS
  #
  Update: FORMAT=V3; DATE=921218; START=930201
```

```
#
Key: CN=Christoph Graf, O=SWITCH, C=CH
#
Name:    Christoph Graf
#
Address: S=Graf; O=SWITCH; P=SWITCH; A=ARCOM; C=CH;
RFC822:  Graf@switch.ch
#
Phone:   +41 1 2565454
Fax:     +41 1 2618133
#
Mail:    SWITCH Head Office /
         Christoph Graf /
         Limmatquai 138 /
         CH-8001 Zurich /
         Switzerland
#
Reachable: 09:00-12:00; 14:00-17:30; UTC+0100
```

Appendix B: BNF Definitions

```
<called-connection> ::= "Called-address: " \
                        <Service-type> "; " \
                        <P-address> "; " <MTS> \
                        ["; " <Service-priority>] <CR>

<calling-connection> ::= "Calling-address: " \
                         <Service-type> "; " \
                         <P-address> <CR>

<checkpoint-size> ::= "RTS-checkpoint-size: " \
                      'checkpoint size' <CR>

<COMMUNITY-document> ::= <Community-Identifier> \
                         <Update-info> \
                         <COMMUNITY-document-body>

<COMMUNITY-document-body> ::= <Coordination> \
                              [{<Macro-definition>}] \
                              {<Connections>}

<Community-Identifier> ::= "Community: " \
                           ('community name' | <DirectoryName>) <CR>

<connection-info> ::= <password> <RTS> \
                      {<called-connection><calling-connection>}\
                      [<system>] \
                      [<local-domain>] \
                      [<echo-server>]

<Connections> ::= {<mandatory-service>} \
                      {[<optional-service>]}

<contact-info> ::= {"Administrator: " <UniquePersonKey> <CR>}

<Coordination> ::= <EMail> <Phone> <FAX> \
                      <Mail> <Operation> <File-server>

<Country-Code> ::= 'Two Character Country Code ISO-3166'

<dialog-mode> ::= "RTS-dialog-mode: " \
                      ("TWA" | "MONOLOGUE") <CR>

<DirectoryName> ::= 'Distinguished Name'

<DOMAIN-document> ::= <Community-Identifier> \
                      <Update-info> \
```

```
                        <DOMAIN-document-body>

    <DOMAIN-document-body>::= {<Domain-entry>} <responsible> \
                        {<Relay>}

    <Domain-entry> ::= "Domain: " <OR-matching> <MHS-subtree> <CR>

    <echo-server> ::= "EchoServer: " <X.400 address> <CR>

    <EMail> ::= "Address: " <X.400 address> <CR>

    <email-server> ::= "Mail-server: "<X.400 address> <CR>

    <extension> ::= 'local extension'

    <Fax> ::= "Fax: " <tel-no-list> <CR>

    <File-server> ::= <email-server> [{<FTP-server>}] \
                        [{<FTAM-server>]}

    <FTAM-server> ::= "FTAM-server: " <P-address> "; "\
                        'account-name' ["; " 'password'] \
                        ["; X.500 " <DirectoryName>] <CR>

    <FTP-server> ::= "FTP-server: " 'domain name' "; " \
                        'account-name' ["; " 'password'] <CR>

    <int-pref> ::= 'international prefix'

    <local-domain> ::= "LocalDomain: " <MHS-subtree> <CR>

    <Macro-definition> ::= "Macro: " 'Macro name' " " \
                        'Macro value' <CR>

    <Mail> ::= "Mail: " 'postal address information' <CR>

    <mandatory-service> ::= "Mandatory-Service: " \
                        <Service-type> <CR>

    <MHS-subtree> ::= ["O=" 'Organization-name' "; "] \
                        ["OU1="'OrganizationalUnit'"; "\
                        ["OU2=" 'OrganizationalUnit' "; " \
                        ["OU3=" 'OrganizationalUnit' "; " \
                        ["OU4=" 'OrganizationalUnit' "; "]]]] \
                        ["P=" 'PRMDname' "; "] \
                        "A=" 'ADMDname' "; " \
                        "C=" <Country-Code> ";"
```

```
<MTS> ::= "MTS-T" | "MTS-TP" | "MTS-TP-84"

<Name>  ::= "Name: " 'name of person' <CR>

<national number> ::= 'national telephone number'

<Network-name> ::= 'Name of a network'

<Network-service> ::= 'Name of a network service'

<Operation> ::= "Reachable: "  {<time> "-" <time> "; "} \
                      <Time-zone> <CR>

<optional-service> ::= "Optional-Service: " \
                      <Service-type> <CR>

<OR-matching> ::=  ( "* " | "= " )

<P-address> ::= 'String encoded Presentation Address'

<password> ::= "Password: " \
                      ("secret" | "none" | \
                      "value=\"" 'password' "\"") <CR>

<PERSON-document> ::= <Community-Identifier> \
                      <Update-info> \
                      <PERSON-document-identifier> \
                      <PERSON-document-body>

<PERSON-document-identifier> ::= "Key: " <UniquePersonKey> <CR>

<PERSON-document-body>::= <Name> {<EMail>} {<RFC822>} \

<Phone> ::= "Phone: " <tel-no-list> <CR>

<Relay> ::=         "Relay: " \
                      'UniqueRELAY-MTAkey' "; " \
                      <RELAY-MTA-Priority> <CR>

<RELAY-MTA-document> ::= <Community-Identifier> \
                      <Update-info> \
                      <RELAY-MTA-document-Identifier> \
                      <RELAY-MTA-document-body>

<RELAY-MTA-document-body> ::= <Status> <connection-info> \
                      <contact-info>

<RELAY-MTA-document-Identifier> ::= \
```

```
                        "RELAY-MTA: " <UniqueRELAY-MTAkey> <CR>

    <RELAY-MTA-Priority> ::= <Integer 0..99>

    <responsible> ::= {"Administrator: " <UniquePersonKey> <CR>}

    <RFC822> ::= "RFC822: " <RFC-822-address> <CR>

    <RTS> ::= <dialog-mode> \
                        [<checkpoint-size> <window-size>]

    <Service-priority> ::= 'Integer 0..99'

    <Service-type> ::= <Network-name> "/" \
                        <Network-service> "/" \
                        <Transport-Protocol>

    <Status> ::= "Status: " ("primary" | "secondary") <CR>

    <system> ::= "System: HW=" 'computer type' "; " \
                        "OS=" 'operating system' "; " \
                        "SW=" 'MHS  software' <CR>

    <tel-no-list> ::= <tel-number> [{"; " <tel-number>}]

    <tel-number> ::=  {"+" <int-pref> " " <national number> \
                        [" x" <extension>]}

    <time> ::= 'hh:mm'

    <Time-zone> ::= ("UTC+" | "UTC-") 'hhmm'

    <Transport-Protocol> ::= 'Name of a transport protocol'

    <UniquePersonKey> ::= (<X.400 address> | <DirectoryName> )

    <UniqueRELAY-MTAkey> ::= (([ "P=" 'PRMDname' "; " ] \
                        ["A=" 'ADMDname' "; " ] \
                        "C=" <Country-Code> "; " \
                        "MTAname=" 'MTAname')
                        | <DirectoryName> )

    <Update-info> ::= "Update: FORMAT=V3; DATE=" 'yymmdd' \
                        "; START=" 'yymmdd' \
                        ["; END=" 'yymmdd'] <CR>

    <window-size> ::= "RTS-window-size: " \
                        'window size' <CR>
```

```
      <X.400 address> ::= 'OR address, ISO 10021-2 Annex F'

      <X.400 routing coordination document set> ::= \
                              <COMMUNITY-document> \
                              { <RELAY-MTA-document> } \
                              { <DOMAIN-document> } \
                              { <PERSON-document> }
```

Security Considerations

Security issues are not discussed in this memo.

Author's Address

Urs Eppenberger
SWITCH Head Office
Limmatquai 138
CH-8001 Zurich
Switzerland

Phone: +41 1 261 8112
Fax: +41 1 261 8133

EMail: Eppenberger@switch.ch
S=Eppenberger; O=SWITCH; P=SWITCH; A=ARCOM; C=CH;

Comments to the document may also be sent to the distribution list wg-msg@rare.nl of the RARE Working Group on Mail and Messaging.