INTERNET-DRAFT   Jeroen Houttuin
RARE WG-MSG   Klaus Hansen
                                                          Serge Aumont
ver. 2.2.   May 1993

Address mapping functions and authorities

Abstract

This document defines the responsibilities and authorities for
defining, collecting and distributing RFC 1327 address mapping
rules. It clearly defines the items: mapping function, addressing
authority, administrative equivalence as well as a mechanism for
registering mapping authorities and administrative equivalence. This
mechanism is based on an extension of RFC 1327 mapping rules (during
the collection distribution process). No changes to already
installed gateway software are required.

Status of this Memo

NB. The reader is assumed to have a solid understanding of X.400,
RFC 822 and RFC 1327. This document is produced by the RARE WG-MSG
Task Force on Mapping Authorities. Comments can be sent to the
authors, tf-mapauth@cosine-mhs.switch.ch , or to wg-msg@rare.nl.
Before sending comments, please read the 'About this document'
section.

This document is an Internet Draft. Internet Drafts are working
documents of the Internet Engineering Task Force (IETF), its Areas,
and its Working Groups. Note that other groups may also distribute
working documents as Internet Drafts.

Internet Drafts are draft documents valid for a maximum of six
months. Internet Drafts may be updated, replaced, or obsoleted by
other documents at any time. It is not appropriate to use Internet
Drafts as reference material or to cite them other than as a
"working draft" or "work in progress."

Please check the I-D abstract listing contained in each Internet
Draft directory to learn the current status of this or any other
Internet Draft.

Distribution of this memo is unlimited.

About   this   document

This chapter explains the goals and reasons for choices made in the
remainder of the document.

– Goals

There are a number of problems this document is targeted to solve.
Some of these targets are by nature conflicting, so the presented
solution will have to be a compromise. We are aware that the
proposed solution will not fully satisfy all parties. However, we
believe that we present in this document a reasonable, pragmatic and
feasible approach. Our goals are:

  – Agreement on the **mapping function** (see chapter 2). RFC 1327
    defines the address mapping function by describing it in text,
    which leaves room for some ambiguities. We need global agreement
    on a precise function definition.

  – Agreement on **mapping authorities**. So far there has not been a
    global consensus on the definitions of 'mapping authority',
    'administrative equivalence' etc. Since these terms must be well
    agreed upon before responsibilities and duties of parties
    involved in the mapping process can be described, we need clear
    definitions.

  – Although most experts agree that **local mappings** are an evil,
    it must also be recognised that they cannot be abandoned for the
    time being. Ignoring them is not realistic and will only create
    more difficulties. Locally mapped addresses can at any time leak
    out to the global level through so-called third-party-routed
    mails. Recognising this fact, we believe that we should at least
    try to gain more control over local mappings, whilst at the same
    time strongly discouraging their use. Our approach treats the
    registration of local mappings the same as this of normal
    mappings, thus enabling the automatic refusal of local mappings
    in case of conflicts.

  – We need clear algorithms and procedures to **solve conflicts** in
    mapping rules.

  – The algorithms and procedures must be **easy to automate** and not
    require adaptation of the installed gateway products.

  – If a transition in the collection and redistribution process of

mapping rules is needed, it must be a **smooth transition** from
the currently used procedures in the Internet and GO-MHS
community.

– Considerations

One of the main issues to be addressed was that of local mappings.
The considerations in favour of our approach as opposed to the
current situation are listed below:

Now

- There exist no global rules for which local mappings are
  allowed. Every gateway manager can add local mappings to the
  distributed tables, even if they conflict with other rules.
  Since it is not feasible to completely abandon local
  mappings some time soon, this leads to anarchistic use of
  local mappings.

- Every gateway manager must check his own local mappings for
  conflicts before merging them with the international tables.
  This may sound like a welcome punishment for those who
  insist on the use of local mappings; in practice however, it
  is a well known source of errors.

- No algorithm for deciding in case of conflicts.

+ Smaller international tables (R2X)

Proposal

- Larger international tables (only R2X, which will however
  not become larger than X2R)

+ One agreed set of rules. If local mappings are collected,
  verified and redistributed for use under certain well
  defined conditions, control is gained over local mappings.
  Invalid local mappings will be automatically overruled. If
  invalid local rules are still being used in a gateway, the
  gateway manager can easily be blamed of violating written
  Internet requirements.

Another choice that has to be explained is the tagging of individual
mapping rules. The disadvantages of this approach are:

- Larger mapping tables during collection and redistribution.

- Extra steps during collection and distribution. Note however

that extra steps will be necessary in any solution for checking
mapping authorities.

It has been said that a better choice would be to tag mapping rules
in bunches. The disadvantages of such an approach however would be:

- It is unlikely that a set of rules will have exactly the same
  definer and exactly the same path of registries. The combination
  of all this information is needed however to be able to decide
  the administrative equivalence and validity of every single
  rule. And in case of a mapping rule rejection, the source of a
  single mapping rule must be traceable.

- This solution would assume a semantics in the order in which
  mapping rules are listed in a table, thus any implementation in
  a non-serial database (DNS, X.500) structure would become more
  complicated.

Finally, proposals were made to not use the mapping rules themselves
for storing the authority information, but to use DNS or X.500
instead. Our main consideration in taking the inline registration
method was the following. It cannot be expected that every mapping
collection/redistribution point has access to X.500 or DNS. Many
gateways exist on islands, e.g. address gateways, which will only
allow end users to address persons in the other mail world in the
address format of that other world, but do not perform the actual
gatewaying themselves. The least common denominator of all involved
parties is access to the mapping rules, regardless whether these are
distributed to them per e-mail, ftp, X.500 or by any other method.
Since every involved party needs to get the mapping rules anyway,
this will even minimise overhead that would be created otherwise by
extra X.500/DNS querying for every single mapping rule.

Contents

1.   Introduction

RFC 1327 defines a mapping between X.400(84/88) and RFC 822
addresses. The requirement for co-ordinated mapping and gateway
tables is included in the RFC to ensure smooth interworking. This
document describes the co-ordination procedures to be used for RFC
1327 gateways connected to the Internet and the GO-MHS community. It
is highly desirable that also other networks using RFC 1327 gateways
use these guidelines.

Note that for brevity this document does not always follow the
normal conventions for representing X.400 domains (see [JHtut]). If
needed, the slash separated notation is used while omitting keywords
of the standard attributes, e.g. /S=plork/dom/pre/amade/nl/ instead
of /S=plork/O=dom/PRMD=pre/ADMD=amade/C=nl/

1.1 Address trees

Addresses may span up a structured or an unstructured address space.
Only the former case is relevant in this document. Structured
addresses may be hierarchical, and thus shown as a path in a tree,
where the tree shows all possible addresses in a given domain. The
authority for registering an address (or a part of an address) may

be associated with the address tree, although it is conceptually a
separate tree.

1.2 Authorities, rights, and responsibilities

An authority gets its authorisation from a higher authority, i.e. an
authority on a higher level, except when it is itself the highest
(or root) authority.

An authority may assume authority in certain circumstances, although
it has not formally got it as described. This may be due to

-   it is unclear who has higher authority
-   no higher authority has yet been set up
-   the higher authority itself is assumed

An assumed authority is temporary in nature, and registration rules
and register may be changed at a later point in time.

An authority has normally some or all of following responsibilities:

a. establishment of registrar and rules for registration
b. delegation of authority
c. establishment of rules for use
d. acting as primary source for validated data on registered items

An assumed authority will be limited in establishing the rules in a
and c.

An authority has the right to revoke registration according to the
set rules.

1.3 The registration process

An item from a certain domain (e.g. name, address, mapping rule) is
registered by a registrar appointed by the authority in question,
according to rules defined at that time. The rules may be simple and
unwritten, or formal (even defined by a national standard). Often
the following steps are taken:

a.The application for registering a certain item is validated,
    to avoid conflicting claims to a certain item, or because
    legal or technical conditions may have to be met. The
    registering may cover single items or whole subtrees.

b.If the conditions are met, the item is filed in the register
    together with information about the applicant. In some cases

part of this information is passed to higher authorities. At
the same time it is decided which rights and responsibilities
is attached to the items.

c.If the conditions for the use of an item are not being met,
or the registrant does not need the item anymore, the item is
remove from the register. It is up to the authority to decide
if the item may be used again.

1.4 Addressing authorities

The difference between names and addresses, and thus between naming
authorities and addressing authorities is insignificant in the
context of this document.

An addressing authority will create an addressing concept with
addressing guidelines that must be followed in (parts of) the
subtree. Underlying addressing authorities can then add their own
addressing concept etc.

1.4.1. Internet

The Internet contains several addressing domains, e.g. RFC 822
addresses, IP addresses, Ethernet addresses, host names. Only RFC
822 addresses are relevant in this document. An RFC 822 address has
the following structure:

        localpart@...sdom(2).sdom(1).tldom

where "sdom" stands for "subdomain", "tldom" stands for "top level
domain", and a hierarchy of addressing authorities is considered to
be growing from left to right:

        localpart < sdom(n) < sdom(n-1) < ... < tldom

Only the domainpart will be considered here (the localpart is - as
one would have expected - mainly a local matter).

The root authority for the RFC 822 address tree resides at SRI-NIC,
and the top level branches have addresses that are either ISO 3166
two-letter country codes or are taken from a small table of domains
(e.g. net, edu, gov, com).

Authority for top level addresses reside at a national organisation;
however a significant number of countries have no authority (and
whether authority then is at the root is an open question).

1.4.2. X.400

The root authority lies in the standard (a sort of "virtual authority"). The first-level authority is more or less implicitly delegated to the various countries according to ISO 3166. Exactly who has the national authority is a national matter, and a "natural" authority depends on whether CCITT X.400 or the equivalent ISO 10021 is to be followed. Some countries assumes that delegation of authority is to the national tele-administration, others define the authority in a national standard covering ADMD names, and in some cases also PRMD names.

The X.400 hierarchy has the following levels:

        PN < OU* < .. < O < PRMD < ADMD < C

                PN   Personal Name
                OU   Organisational Unit(s)
                O   Organisation
                PRMD   Private Management Domain
                ADMD   Administration Management Domain
                C   Country

Other attributes, such as Domain Defined Attributes (DDAs), may be a part of an address, but are not unambiguously hierarchical in nature.

According to a national decision, authority over PRMD names is either delegated to the ADMD level (each ADMD having a complete addressing subtree) or kept at the national level. In the latter case ADMD names and PRMD names may possibly all be taken from the same set, thus making it possible to use a registered name as an ADMD name, a PRMD name, or both.

1.5 Pruned subtrees

An addressing authority does not automatically have control over all branches of a sub-tree. Consider the following example:

```
---------------------------+---------------------------
      X.400                 |          RFC 822
---------------------------+---------------------------
CH:                         | .ch:
addr. authority: Swiss PTT  | addr. authority: SWITCH
---------------------------+---------------------------
 root                       |          root
  /\                        |          /\
 /   /CH/                    |         / .ch
/     /\                     |        /    /\
   /   \                     |       /   \
  /    /CH/ARCOM/            | .arcom.ch     \
  /\       /\               |        /\     /\
 /  \   /   \               |       /  \   /  \
/\   /\/\  /CH/ARCOM/SWITCH/ |
          /\                 |
         /   \               |
---------------------------+---------------------------
/CH/ARCOM/SWITCH/:          | .arcom.ch:
addr. authority: SWITCH     | addr. authority: Swiss PTT
---------------------------+---------------------------
```

The example also shows that a mapping registration tree will in
general not coincide with either of the complete addressing trees.

2. Mapping functions

2.1 Introduction

RFC 1327 describes the address mapping function, but at the same
time leaves room for some controversies. This chapter aims to
unambiguously define the address mapping function.

Note that RFC 1327 defines 4 types of mapping rules:

```
    RFC 822 -> X.400
    X.400 -> RFC 822
    RFC 822 -> gateway X.400 domain
    X.400 -> gateway RFC 822 domain
```

Since the last type of rule is not being used in an internationally
co-ordinated way, use of such rules is considered a local matter and
is discarded in this document. If one absolutely wants to use such
rules, it is easy to extend the X.400 -> RFC 822 algorithm.

2.2 Function definition

The mapping algorithm to be used by a gateway assumes the existence
of three tables, X2R, R2X and GW, which associate RFC 822 and X.400
domains. Left hand sides are unique in X2R and in the concatenation
of R2X and GW. The algorithm is defined as follows in pseudo code
(for a more comprehensive description, see [JHtut] and [1327]):

RFC 822 -> X.400

```
LHS encoded X.400 address?
y: unpack; GOTO END                                        [a]
n:
    map2 entry?
    y: use SA's of map2 entry; follow hierarchy for other SAs
       localpart regular?
       y: map localpart -> PN
          GOTO END                                          [b]
       n: GOTO DDA                                           [c]
    n: gate entry?
       y: use SAs of gate entry                              [d]
       n: use SAs of local gateway                           [e]
:DDA: encode complete address in a DD.RFC-822
:END:
```

X.400 -> RFC 822

```
Address contains DD.RFC-822?
y: unpack DDA; GOTO END                                    [A]
n:
    map1 entry?
    y: use domains of map1 entry
       other attributes regular?
       y: follow hierarchy for other subdomains;
          map PN-> localpart
          GOTO END                                          [B]
       n: follow hierarchy for other subdomains as
          far as possible;
          GOTO LHS with rest of attributes                  [C]
    n:                                                      [D]
:LHS: Left hand side encoding
:END:
```

Examples:

Consider a gateway in country 'Z' which is known in the RFC 822
world as 'gw.z' and in the X.400 world as '/GW/Z/', and suppose only
the following mapping rules are used in this gateway:

```
X2R:
    C$A#a#
R2X:
    a#C$A#
GW:
    c.a#ADMD$D.PRMD$E.C$A#
    b.c#ADMD$B.C$C#
```

Then this gateway could perform the following mappings (the examples
follow the order of the pseudo code):

```
[a] /S=jan/ADMD=amade/C=xy/@gw.z        /S=jan/ADMD=amade/C=xy/
[a] /S=jan/ADMD=amade/C=xy/@gw.y        /S=jan/ADMD=amade/C=xy/
[b] jan@c.b.a                                         /S=jan/C/B/A/
[b] jan@b.c.a                                         /S=jan/B/C/A/
[c] j_h@b.c.a                        /DD.RFC-822=j(u)h(a)b.c.a/B/C/A/
[d] jan@a.b.c                        /DD.RFC-822=jan(a)a.b.c/A/B/C/
[e] jan@d.b                          /DD.RFC-822=jan(a)d.b/GW/Z/

[A] /DD.RFC-822=jan(a)xx.yy/GW/Z/                        jan@xx.yy
[A] /DD.RFC-822=jan(a)xx.yy/GW/Y/                        jan@xx.yy
[B] /S=jan/C/B/A/                                        jan@c.b.a
[C] /S=jan/GQ=jr/C/B/A/                      /S=jan/GQ=jr/@c.b.a
[C] /S=jan/D C/B/A/                          "/S=jan/D C/"@b.a
[D] /S=jan/B/C/                              /S=jan/B/C/@gw.z
```

Note that the sole fact that the gateway could perform a mapping
doesn't force it to do so. This depends on the preferred routing,
e.g. a gateway may choose not to map back (and gateway) a DDA mapped
address which contains the SAs of a remote gateway, but rather route
this message over X.400 to the addressed gateway which will then
have to perform the gatewaying. This is normal practice in most
algorithms for source routing, for which left-hand side encoding and
DDA mappings can also be used.

2.3 Ideal situation

In the set of co-operative RFC 1327 gateways on the planet 'GW-manager-
utopia', there exist no mapping rules. Every address is mapped with LHS
encoding and DDA mapping. Although this configuration may ease the life
of gateway managers, it also creates many problems, mainly for the
users (see [JHtut] Chapter 3.3.2.).

In the set of co-operative RFC 1327 gateways on the planet 'user-
walhalla', mapping of RFC 822 addresses and X.400 O/R addresses is
simple and divided in:

- a set of RFC 822 and X.400 addresses (domains) with
  administrative equivalence and bijective mappings between them.
- a set of O/R addresses which are RFC 822 visible by using left-
  hand side encoding.
- a set of RFC 822 addresses which are X.400 visible by using DDA
  mapping.

This should be quite simple and each solution should be exclusive.
Therefore an address should have only one representation in each
address space.

2.4 Reality

On the planet earth, the experience shows several cases where the
mapping between RFC 822 and X.400 domains is not bijective
(asymmetric mappings) and that such asymmetry is perhaps still
indispensable. It is useful to list the most common cases.

- Fading out old address forms.

If, for instance, a domain changed from one ADMD connection to
another, it may choose to support the old mapping for a certain
period. Since two X.400 domains are now associated with one and the
same RFC 822 domain, asymmetry is introduced.

- An address tree is not always a real tree

For instance, a PRMD may subscribe to several ADMDs and then one
mailbox can be identified by different O/R addresses. The following
mapping rules show an example:

```
    PRMD$blabla.ADMD$ .C$ch#blabla.ch#          [1]
    PRMD$blabla.ADMD$switch.C$ch#blabla.ch#     [2]
    PRMD$blabla.ADMD$eunet.C$ch#blabla.ch#      [3]

    blabla.ch#PRMD$blabla.ADMD$ .C$ch#          [4]
    blabla.ch#PRMD$blabla.ADMD$eunet.C$ch#      [5]
    blabla.ch#PRMD$blabla.ADMD$switch.C$ch#     [6]
```

Since rules [4] [5] and [6] all have administrative equivalence (see
chapter 3.1), it is perfectly legal for EUnet to use mapping rule [5],
probably to ensure that subsequent mails will be routed over their
ADMD. Since these mapped address forms can leak out to the rest of the
world (third party problem), all reverse rules must be global. This
also results in asymmetry.

The gateway's choice of which of the rules [4] [5] and [6] to use can
either be made according to local routing considerations or the

"blabla.ch" addressing authority can claim its preferred O/R address.
Not only routing, but also mapping depends on where you are (the most
trivial examples being the default left hand side encoding and DDA
mapping, but they can be mapped back without the use of mapping rules).

- **Subtrees without addressing authority**

The domains ".uucp" or ".bitnet" are used and usually well routed in
Internet networks but no addressing authority has ever registered them.
This implies that nobody can define an official mapping for those
domains. Therefore there are many ways to map such addresses, none of
which can be considered more valid or invalid than any of the others.
Some examples are:

R2X mappings :

```
bitnet#PRMD$bitnet.ADMD$ada.C$at#                  [1]
bitnet#O$bitnet.PRMD$switch.ADMD$arcom.C$CH#       [2]
bitnet#PRMD$bitnet.ADMD$dbp.C$de#                  [3]
```

GW rules :

```
bitnet#PRMD$bitnet.ADMD$0.C$FR#                    [4]
```

All those rules show different ways to relay X.400 messages to the
bitnet network. Rules [1], [2] and [3] have different semantics than
the domain association rules R2X. They are used to force the address of
a gateway into an O/R address. As such, they can be considered as the
RFC 987 equivalent of gateway rules. The usage of such rules is limited
to a specific area and every gateway has to choose which one to use.
Since these mapped address forms can leak out to the rest of the world
(third party problem), all reverse rules must be global. This also
results in asymmetry.

These mappings are referenced as "local mapping" in the document
[table-creation-tutorial].

3. Mapping authorities

This chapter defines the parties involved in the process of
defining, collecting and distributing mapping rules within a
community. Note that a party may at any time choose to have certain
responsibilities and authorities represented by automated processes.
Another important generalisation is that the intuitively centralised
approach need not be followed strictly. If mapping rules are
maintained in a distributed way, distributed tools may become
necessary to enforce the responsibilities and authorities described

in this document. However, since conflicts must normally be solved
on an inter-personal basis, the defined parties must be clearly
defined: a central contact point for every involved party must be
available.

3.1 Administrative equivalence (AE)

A mapping rule establishes a one-way correspondence between
addresses from two different domains. A mapping rule is used e.g. in
a gateway between mailing systems for transformation of addresses. A
mapping rule may map one address only (as in diku.dk --map-->
C=dk;A=dk400;P=minerva;O=diku), but it is more common to define
general rules for mapping a whole tree.

A mapping rule is based on the existence of administrative
equivalence. This means that to define a valid mapping rule, one
must have authority over the relevant addresses in both addressing
domains. AE is defined as follows:

A mapping rule has AE if and only if both sides of the rule have the
same addressing authority (or they agree on the rule), or all
mapping rules implied by this rule span up two subtrees that have AE
in every (at least one) corresponding pair of nodes.

Examples

Danish mapping rules (internal):

        RFC 822 addresses:

                id   domain   authority?
                --   ------   ----------

                 A   teldk.dk   assumed
                 B   y-net.dk   assumed
                 C   ooo.dk   yes
                 D   .dddd   assumed

        X.400 addresses:

                id domain   authority?
                -- ------   ----------

                I   C=dk;A=teldk   assumed
                II   C=dk;A=dk400;P=y-net   yes
                III C=dk;A=dk400;P=minerva;O=ooo yes
                IV   C=dk;A=dk400;P=inet;O=dddd   yes

```
   Mapping RFC 822 <--> X.400

      I ---> A    (ass to ass)      AE
      B ---> II   (ass to auth)     Agreement exists
      C ---> III  (auth to auth)    AE
      IV --> D    (auth to ass)     Local rule, no AE
```


3.2 Mapping registries (MRs)

The following parties and corresponding responsibilities are
defined:

Mapping rule originator

  define mappings

Mapping registry (MR):

  Designate subordinate MRs per subdomain (822/X.400)
  Collect mappings from subordinate MRs
  Inform subordinate MRs about rejected mappings
  Register mappings with next higher MR
  Redistribute mappings received from next higher MR

Ideally, there will b onmly one MR per community or branch
(subdomain) of the address tree.

A top level MR (initially the MHS Co-ordination Service) is
responsible for the collection and redistribution of the complete
gateway and mapping tables.

3.3 A mechanism for claiming and tracing authorities

In order to check for AE, and see who is responsible for certain
mapping rules, a formal mechanism for registering the relevant
authority information is needed. A commonly used strategy is to
merge this authorisaty information with the information that is to
be authorised (e.g. authority records in DNS). This approach has the
advantage that authorisaty information is automatically available at
the moment it is needed. Therefore, to each mapping rule some extra
fields are added. The first extar field indicates the AE:

  com#....C$it#n#
  ch#PRMD$switch.ADMD$arcom.C$ch#y#

Before registering a mapping rule at the next higher MR, a mapping
rule originator or MR adds an extra field to the mapping rule,

indicating its identity relative to this next higher MR. The top
level mapping registry could thus receive the following mapping
rule:

    ciba.ch#O$ciba.PRMD$eunet.ADMD$arcom.C$ch#y#eunet#switch

More formally:

Appendix F of RFC 1327 gives a syntax definitions of three kinds of
mapping rules:

rules defined in RFC 1327   named in this document

mapping from O/R address to
Internet domains   X2R
mapping from Internet domains
to OR-addresses using standard
attributes   R2X
mapping from Internet domains to
RFC 822 using domain defined
attributes OR-addresses   GW

The general form is :

R2X and GW: domain-syntax"#"dmn-or-address"#"

X2R: dmn-or-address"#"domain-syntax"#"

There a need to extend each rule with the following information :

    - Is there a administrative equivalence of both domains ?
    - which addressing authority submit this rule
    - who collect this rule

The left side and right hand side of the rules are unchanged. Existing
conformant RFC 1327 gateways do not need any change.

Mapping authority fields are added:

R2X and GW:

domain-syntax"#"dmn-or-address"#"AE"#"originator\
"#"registry"#"*(registry"#")

X2R:

dmn-or-address"#"domain-syntax"#"AE"#"originator\
"#"registry"#"*(registry"#")

   AE = "Y" / "N"

     - Administrative equivalence is Y / N . If it's "Y" it means the
       two domains are under control of the same addressing authority
       or it exist a agreement between the two different addressing
       authorities that guarantee this equivalence.

   originator= < non empty string without "#" >

     - if the administrative equivalence is "Y", this is the name of
       the addressing authority over both domains, otherwise it is the
       name of the authority that submitted this rule.

   Registry = < non empty string without "#" >

     - Registry : each registry are the named of the mapping registry
       who accept this mapping and relay it to upper registry.

   Examples:

   glvt.fr#O$@.PRMD$GLVT.ADMD$atlas.C$FR#Y#glvt-cnrs#uniren1#PT#

   This rule shows administrative equivalence, it has been submitted by
   glvt-cnrs via the University of Rennes (French registry) and the
   COSINE MHS Project Team (PT), which is initially designated as the
   GO-MHS community registry.

   bitnet#PRMD$bitnet.ADMD$atlas.C$fr#N#uniren1#uniren1#PT#

   This gateway rule is issued by the University of Rennes and without
   administrative equivalence.

3.4 Using the extended mapping rules

   Mapping collection from subordinate registries

   Each registry will accept or refuse rules. Accepted rules are stamped
   at the end with the registry name according to the following process :

   For each rule received : Does this rule conflict with other rules?
      [important note : R2X and GW rules set are considered as a
      single set when checking for conflicts]
   No: accept (I.e. trust! This will be the default, to save us work)
   Yes:
     Does at least one rule claim AE ?
     No: accept
     Yes: For each of the conflicting mapping rules: check AE

                AE: accept
                No AE: Refuse

     Mapping rule conflicts are classified as follows:

             - Pure conflicts: two mapping rules have exactly the same left-
               hand side.
             - Exception conflicts: An exception rule without AE tries to
               overrule a more general rule with AE.

     Subordinate registries are informed about rejected rules. This
     subordinate registry must propagate the rejection notifications to
     the appropriate subordinate registry or originator. This is
     necessary because conflicts may not occur until a later stage in the
     collection process.

     All accepted mappings are stamped with the registry identification
     and registered with the next higher registry.

     When mappings arrive at the top-level registry, the redistribution
     process starts.

     This collection process does not guarantee unique left-hand sides,
     but ensures that in a remaining set of rules with the same left-hand
     side, either each rule has AE, or not one rule has AE. The first
     case can occur because the addressing tree is not a real tree (ADMD
     = ; ADMD=0; other aliases); the second case covers the traditional
     'local mappings'.

     Redistribution of mappings

     The collection of all mappings that were accepted by the top level
     registry is distributed in three tables, X2R, R2X and GW, using the
     same channel of mapping registries (actually, this is a supertree of
     the rigistry tree, see chapter 3.4.c.). Before distributing the mapping
     rules, a registry may tailor the tables for a certain domain according
     to the algorithm described below. It is recommended though that this
     tailoring is done as close to the actual gateways as possible, i.e.
     ideally within the gateway.

     Use of mapping rules in a gateway

     Here is a description of what can be a pre-processor for the 3 sets
     of mapping rules. The goal is to select the best set of rules and
     remove the extension before using the rules in the gateway. The GW
     software could be adapted to use the extended rules directly, or a
     script can generate the traditional 1327 tables from the distributed

new-format tables, just before installing them in the gateway. The
mapping of a domain remains as described in chapter 2, except there
may be non-unique left-hand sides (again, R2X and GW are considered
one set of rules here). In this case, a gateway must use the rule
which is closest to the gateway (the metric being the distance in
the tree of registries).

This algorithm ensures that in every gateway, for each domain,
exactly one mapping rule is left. It also allows different domains
to automatically use different versions of equivalent rules,
depending on their location in the authority tree.

4 Registering Authorities

4.1 Top level authority registration

The following steps should be taken when a new mapping domain is
introduced on the top level.

   a. The domain (country) finds a top-level mapping collecting
      organisation, which can act as the representative for the
      domain, and introduces the organisation to the other top-level
      mapping registries. The following information is given for the
      organisation:

        – Mapping domains (e.g. RFC 822 to X.400)

        – Mapping table designator for the organisation

        – Name, address, telephone and fax numbers for the
          organisation

        – Name and e-mail address for the responsible person(s). The
          e-mail address is used to verify that the source of a given
          mapping is an authorised person.

        – Domain, for which this organisation has responsibility for
          collection of mappings.

   b. The collection organisation inside the domain is set up. It may
      be done in the following way.

        – In case a mapping authority for a certain address tree pair
          has been established, the relevant authority decides whether
          to rely on a general (default) rule or to define own mapping
          rules. In the former case, a collecting path from the

authority up to the top-level registry is established, and
an organisation designator (an abbreviated identifier for
it) is defined for each part of the path. In the latter case
the decision is propagated to the relevant registry.

- The top-level registry may be the authority for certain
  mappings, or act simply as a registry. It defines a general
  (default) mapping for the whole domain, unless otherwise
  decided by the domain.

- If authority cannot be established for certain mappings, it
  is marked in the mapping rule. A path is constructed as if
  authority had been established, from the relevant
  organisation to the top-level registry.

- Note that the collection path need not coincide with the
  path in the address tree, as addressing authorities and
  mapping rule originators will often have no interest in the
  collection and distribution process.

   c.Distribution of top-level mappings will in principle follow the
     reverse path of the collection process, but as potentially the
     set of sites using mappings may be significantly larger that the
     set of organisations defining mappings, the distribution could
     be entirely independent of the collection process. It is
     entirely up to the domain to organise distribution of mappings.

4.2 Authentication of mapping registries

The mapping registries on the top level (e.g. country level) need to
be identified by a designator, occurring in the mapping tables. To
be able to know which organisation has a given designator, and which
person (mail address) is authorised to publish mapping tables, some
information has to be collected. This is similar to the COSINE
documentation for the national networks. The final implementation of
this documentation could be either table-based, DNS based, X.500
based, or a combination of any of those. Although X.500 would seem a
natural choice, the same problems could occur as with the tagging
and storing of the mapping rules themselves: not every gateway will
have access to DNS or X.500 (e.g. address gateways). However, since
the registry documentation is only to be used by the registries
themselves, we would consider it a feasible requirement for each
registry to have access to X.500. For the time being, a simple table
based approach will be feasible, as conflicts will have to be solved
by humans anyway.

The information needed is:

        – Registry designator
        – Name, address, telephone and fax numbers for the organisation
        – Name and e-mail address for the responsible person(s). The e-
          mail address is used to verify that the source of a given
          mapping is an authorised person.
        – Domain, for which this organisation has responsibility for
          collection of mappings.

   Example:

   Registry: isi-dk
   Description: ISI-DK. A co-operation between Danish parties. Contact
   organisation is UNI-C, Bygning 305 DTH, DK-2800 Lyngby, Denmark.
   Telephone: +45 45 938355, Fax: +45 45 930220
   Responsible: Erik Lawaetz
   822: Erik.Lawaetz@uni-c.dk
   X.400: C=dk;A=dk400;P=minerva;O=uni-c;S=Lawaetz;G=Erik
   X.400-domains: C=dk;A=dk400;P=minerva
   822-domains: all sub-domains in ".dk" except "teldk.dk" and
   "dk400.dk"

   The exact syntax for this information will in a next version of this
   document be aligned with [UE93].

5. Guidelines

   It is recommended that whoever defines a mapping rule informs the
   mapped subtrees that an implicit mapping for their domains exists.

   Every mapping registry is recommended to have X.500 access.

   The use of local mappings is strongly discouraged. Instead, the
   definition of GW entries for such domains is encouraged. Please note
   that also GW entries without AE will be rejected if one GW rule with
   AE exists.

   The originator and registry strings to be used as tags in the
   mapping rules should be as short as possible.

   Justified by registry overhead, it is recommended to minimise the
   distance (in the registry tree) between the top level registry and
   the originators cq. gateways. This means that the introduction of
   any not strictly necessary MR is discouraged.

A. Glossary

   - Assumed authority: Temporary authority assumed in the absence of
     a higher authority, or because a higher authority has only
     assumed authority.

   - Authority: A combination of rights and responsibilities in a
     given context. Examples are the right to delegate authority, the
     right to define a mapping or override downwards mappings, and
     the responsibility to validate mappings.

   - Confirmed authority: Authority established either by natural
     rights, or by delegation from a higher authority.

   - Delegation: Attributes may be propagated downwards (to branching
     points in the tree further away from the root) by delegation.
     Attributes are e.g. the right to add new subtrees, or the right
     to use the names or addresses formed by the subtrees.

   - Downwards: Away from the root.

   - Explicit mapping rule: rule that is stated in a mapping table

   - Implicit mapping rule: mapping for subtrees implied by an
     explicit mapping rule.

   - Pruned subtree: A tree with one or more complete branches
     removed.

   - Registration: The process of registering e.g. names, addresses
     or mapping rules, according to rules defined by an authority,
     for a given domain (set of names, set of addresses, mapping
     context). Registration covers e.g. filing the item in question
     and related information about who registered, and (depending on
     the rules) validation of the application. Inquiries as to the
     contents of the register may or may not be allowed, depending on
     the rules.

   - Scope: Authorities, registration, trees, and mappings are only
     valid in a certain context, called its scope.

   - Subtree: A complete tree that is a part of another tree, i.e.
     where the root is a branching point of in a tree.

   - Tree: The classical computer science tree, with the root
     upwards. In this context, each arc (branch) bears a part of an
     address or name. A path from root to a leaf thus defines a
     complete name or address.

   – Upwards: Towards the root.

B. Initial top level mapping registries

B.1. X.400 to RFC 822

```
   Domain/
   Country  Org.          Organisation & responsible design.
   =======  ====          ================================
   AT       aconet        ACOnet Christian Panigl
   BE       iihe          ULB/Helios-B group, Eftimios Tsigros
   BR       dfn           DFN, Peter Kaufmann
   CA       cdn           CDNnet, Dave Brent
   CH       switch        SWITCH, Felix Kugler
   CN       dfn           DFN, Volkmar Kobelt
   DE       dfn           DFN, Volkmar Kobelt
   DK       isi-dk        ISI-DK, Erik Lawaetz, Klaus Hansen
   ES       iris          redIRIS/FUNDESCO, Ignacio Martinez
   FI       funet         FUNET, Marko Kaittola, Teemu Kurki
   FR       uniren1       CRI/Universite de Rennes1, Serge Aumont
   GB       janet         JANET
   GR       ariadne       ARIADNE Network, Yannis Corovesis
   IE       ucd           University College Dublin, Niall O'Reilly
   IN       ernet         ERNET, ???
   IS       isanet        ISAnet, Marius Olaffsen
   IT       infn          CNAF INFN, Claudio Allocchio
   LT       litnet        LITNET, Petras Sulcas
   LU       restena       RESTENA, Alain Frieden
   NL       surfnet       SURFNET, Aad Boer
   NO       uninett       SINTEF, Harald Eikrem, Harald Alvestrand
   PT       inesc         INESC, Henrique Silva
   SE       sunet         Chalmers, Per Andersson
   SI       si-ac         Jozef Stefan Institute, Avgust Jauk
   US       xnren         UW-Madison, Allan Cargille
   YU       yunac         ???, Avgust Jauk, Marko Bonac
   ZA       uninet        UNINET, Rob Brain
```


B.2. RFC 822 to X.400

   Identical to table for X.400 to RFC 822, except for he following
   entries

```
   Domain/
   Country  Org.          Organisation and responsible design.
   =======  ====          ================================
```

```
CH        switch      SWITCH, Felix Kugler
CH-EUNET ch-eu        CH EUNET, Simon Poole
US        xnren       UW-Madison, Allan Cargille
US-ES     esnet       ESnet, Tony Genovese
```

Table: RFC 822 domains for which national authorities assume local
responsibility:

```
  COM
  EDU
  GOV
  MIL
  NET
  ORG
  NATO
  ARPA
  BITNET
  UUCP
```


C. Bibliography

821   Jonathan B. Postel, "Simple Mail Transfer Protocol",
             RFC 821, University of Southern California, August
             1982

822   Crocker, D., "Standard of the Format of ARPA
             Internet Text Messages", RFC 822, UDEL, August 1982.

1327   Hardcastle-Kille, S., "Mapping between X.400(1988) /
             ISO 10021 and RFC 822", RFC 1327 & RARE Technical
             Report 2, UCL, May 1992.

JHtut   Houttuin, J., "A tutorial on RFC 822 <-> X.400
             gatewaying", Internet-Draft draft-houttuin-rfc1327-
             tutor-02.txt, draft-houttuin-rfc1327-tutor-02.ps,
             RARE Secretariat, February 1993.

UE93   Eppenberger, U., "Routing coordination for X.400 MHS
             services within a multi protocol / multi network
             environment", Internet-Draft draft-ietf-x400ops-
             service-coordination-03.txt, SWITCH, December 1992.

X.4xx(84)   CCITT Recommendations X.400 - X.430. Data
             Communication Networks: Message Handling Systems.
             CCITT Red Book, Vol. VIII - Fasc. VIII.7, Malaga-
             Torremolinos 1984

     X.4xx(88)   CCITT Recommendations X.400 - X.420. Data
                  Communication Networks: Message Handling Systems.
                  CCITT Blue Book, Vol. VIII - Fasc. VIII.7, Melbourne
                  1988

D. Table pre-processor

   Example perl script for table pre-processing to be written.

E. Authors' addresses

   Jeroen Houttuin
   RARE Secretariat
   Singel 466-468
   NL-1017 AW Amsterdam
   Phone +31 20 639 1131
   Fax   +31 20 639 3289
   Mail houttuin@rare.nl

   Klaus Hansen
   Department of Computer Science (DIKU)
   University of Copenhagen
   Universitetsparken 1
   DK-2100 Copenhagen
   Phone +45 35 32 18 18
   Fax   +45 52 32 14 01
   Mail khan@diku.dk

   Serge Aumont
   CRI/Universite de Rennes1
   Av. du General Leclercs
   F-35042 Rennes Cedex
   Phone +33 99 84 71 00
   Mail Serge.Aumont@univ-rennes1.fr