Network Working Group
INTERNET-DRAFT

P. Barker and S.E. Hardcastle-Kille
University College London

April 1992

Naming Guidelines for Directory Pilots
(OSI-DS 12 (v5))

Status of this Memo

Deployment of a Directory will benefit from following certain guidelines. This document defines a number of naming guidelines. Alignment to these guidelines is recommended for directory pilots.

This draft document will be submitted to the RFC editor as an informational document. Distribution of this memo is unlimited. Please send comments to the authors or to the discussion group `<osi-ds@CS.UCL.AC.UK>`.

1 Introduction

As a pre-requisite to this document, it is assumed that the COSINE and Internet X.500 Schema should is followed [BHK91].

2 DIT structure

The majority of this document is concerned with DIT structure and naming for organisations, organisational units and personal entries. This section briefly notes three other key issues.

2.1 The top level of the DIT

The following information will be present at the top level of the DIT:

Participating Countries The entries should contain suitable values of the "Friendly Country" attribute.

International Organisations An international organisation is an organisation, such as the United Nations, which inherently has a brief and scope covering many nations. Such organisations might be considered to be supra-national and this, indeed, is the raison-d'etre of such organisations. Such organisations will almost all be governmental or quasi-governmental.

A multi-national organisation is an organisation which operates in more than one country, but is not supra-national. This classification includes the large commercial organisations whose production and sales are spread throughout a large number of countries.

International organisations, may be registered at the top level. This will not be done for multi-national organisations. The only international organisation registered so far is: Internet. This is not a formal registration, but is adopted for the Internet Directory Service.

Localities A few localities will be registered under the root. The chief purpose of these locality entries is to provide a "natural" parent node for organisations which are supra-national, and yet which do not have global authority in their particular field. Such organisations will usually be governmental or quasi-governmental. Example localities might include: Europe, Africa, West Indies. Example organisations within Europe might include: European Court of Justice, European Space Agency, European Commission.

DSA Information Some information on DSAs may be needed at the top level. This should be kept to a minimum.

The only directory information for which there is a recognised top level registration authority is countries. Registration of other information at the top level may potentially cause problems. At this stage, it is argued that the benefits of additional top level registration outweighs these problems. However, this potential problem should be noted by anyone making use of such a registration.

2.2 The DNS within the DIT

The rules for the DNS parts of the DIT are defined in [HK91]. One modification to this is that the DNS tree will be rooted under "O=Internet", rather than at the root of the DIT.

2.3 Access control

An entry's object class attribute, and any attribute(s) used for naming an entry are of special significance and may be considered to be "structural". Any inability to access these attributes will often militate against successful querying of the Directory. For example, user interfaces typically limit the scope of their searches by searching for entries of a particular type, where the type of entry is indicated by its object class. Thus, unless the intention is to bar public access to an entry or set of entries, the object class and naming attributes should be publicly readable.

3 Naming Style

The first goal of naming is to provide unique identifiers for entries. Once this is achieve, the next major goal in naming entries should be to facilitate querying of the Directory. In particular, support for a naming structure which facilitates use of user friendly naming is desirable. Other considerations, such as accurately reflecting the organisational structure of an organisation, should be disregarded if this has an adverse effect on *normal* querying. Early experience in the pilot has shown that a consistent approach to structure and naming is an aid to querying using a wide range of user interfaces, as interfaces are often optimised for DIT structures which appear prevalent.

Naming is dependent on a number of factors and these are now considered in turn.

3.1 National Guidelines

Where naming is being done in a country which has established guidelines for naming, these guidelines should in general be followed. These guidelines might be based on an established registration authority, or may make use use of an existing registration mechanism (e.g., company name registration).

Where an organisation has a name which is nationally registered in an existing registry, this name is likely to be appropriate for use in the Directory, even in cases where there are no national guidelines.

3.2 Structure Rules

A DIT structure is *suggested* in Annex B of X.521, and it is recommended that Directory Pilots should follow a slightly modified form of these guidelines. The rules should be extended for handling DNS [HK91]. Some simple restrictions should be applied, as described below.

For most countries pilots, the following simple structure should suffice. The country entry will appear immediately beneath the root of the tree. Organisations which have national significance should have entries immediately beneath their respective country entries. Smaller organisations which are only known in a particular locality should be placed underneath locality entries representing states or similar geographical divisions. Large organisations will probably need to be sub-divided by organisational units to help in the disambiguation of entries for people with common names. Entries for people and roles will be stored beneath organisations or organisational units. An example plan evolving for the US is the work of the North American Directory Forum [For91].

As noted above, there will be a few exceptions to this basic structure. International organisations will be stored immediately under the root of the tree. Multi-national organisations will be stored within the framework outlined, but with some use of aliases and attributes such as *seeAlso* to help bind together the constituent parts of these organisations. This is discussed in more detail later.

3.3 Depth of tree

The broad recommendation is that the DIT should be as flat as possible. A flat tree means that Directory names will be relatively short, and probably somewhat similar in length and component structure to paper mail addresses. A deep DIT would imply long Directory names, with somewhat arbitrary component parts, with a result which it is argued seems less *natural*. Any artificiality in the choice of names militates against successful querying.

A presumption behind this style of naming is that most querying will be supported by the user specifying convenient strings of characters which will be mapped onto powerful search operations. The alternative approach of the user browsing their way down the tree and selecting names from large numbers of possibilities may be more appropriate in some cases, and a deeper tree facilitates this. However, these guidelines recommend a shallow tree, and implicitly a search oriented approach.

It may be considered that there are two determinants of DIT depth: first, how far down the DIT an organisation is placed; second, the structure of the DIT within organisations.

The structure of the upper levels of the tree will be determined in due course by various registration authorities, and the pilot will have to work within the given structure. However, it is important that the various pilots are cognisant of what the structures are likely to be, and move early to adopt these structures.

The other principal determinant of DIT depth is whether an organisation splits its entries over a number of organisational units, and if so, the number of levels. The recommendation here is that this sub-division of organisations is kept to a minimum. A maximum of two levels of organisational unit should suffice even for large organisations. Organisations with only a few tens or hundreds of employees should strongly consider not using organisational units at all. It is noted that there may be some problems with choice of unique RDNs when using a flat DIT structure. Multiple value RDNs can alleviate this problem. The standard recommends that an organizationalUnitName attribute can also be used as a naming attribute to disambiguate entries. Further disambiguation may be achieved by the use of a personalTitle attribute in the RDN.

3.4 Organisation and Organisational Unit Names

The naming of organisations in the Directory will ultimately come under the jurisdiction of official naming authorities. In the interim, it is recommended that pilots and organisations follow these guidelines. An organisation's RDN should usually be the full name of the organisation, rather than just a set of initials. This means that *University College London* should be preferred over *UCL*. An example of the problems which a short name might cause is given by the proposed registration of *AA* for the Automobile Association. This seems reasonable at first glance, as the Automobile Association is well known by this acronym. However, it seems less reasonable in a broader perspective when you consider organisations such as *Alcoholics Anonymous* and *American Airlines* which use the same acronym. Just as initials should usually be avoided for organisational RDNs, so should formal names which, for example, exist only on official charters and are not generally well known. There are two reasons for this approach:

1. The names should be meaningful.
2. The names should uniquely identify the organisation, and be a name which is unlikely to be challenged in an open registration process. For example, *UCL* might well be challenged by United Carriers Ltd.

The same arguments on naming style can be applied with even greater force to the choice of RDNs for organisational units. While abbreviated names will be in common parlance within an organisation, they will almost always be meaningless outside of that organisation. While many people in academic computing habitually refer to *CS* when thinking of *Computer Science*, CS may be given several different interpretations. It could equally be interpreted as *Computing Services, Cognitive Science, Clinical Science* or even *Counselling Services.*

For both organisations and organisational units, extra naming information should be stored in the directory as alternative values of the naming attribute. Thus, for University College London, *UCL* should be stored as an alternative organizationName attribute value. Similarly *CS* could be stored as an alternative organizationalUnitName value for Computer Science and any of the other departments cited earlier. In general, entries will be located by searching, and so it is not essential to have names which are either memorable or guessable. Minimising of typing may be achieved by use of carefully selected alternate values.

3.5 Naming human users

A reasonably consistent approach to naming people is particularly critical as a large percentage of directory usage will be looking up information about people. User interfaces will be better able to assist users if entries have names conforming to a common format, or small group of formats. It is suggested that the RDN should follow such a format. Alternative values of the common name attribute should be used to store extra naming information. It seems sensible to try to ensure that the RDN commonName value is genuinely the most common name for a person as it is likely that user interfaces may choose to place greater weight on matches on the RDN than on matches on one of the alternative names. It is proposed that pilots should ignore the standard's recommendations on storing personal titles, and letters indicating academic and professional qualifications within the commonName attribute, as this overloads the commonName attribute. A personalTitle attribute has already been specified in the COSINE and Internet Schema, and another attribute could be specified for information about qualifications.

Furthermore, the common name attribute should not be used to hold other attribute information such as telephone numbers, room numbers, or local codes. Such information should be stored within the appropriate attributes as defined in the COSINE and Internet X.500 Schema. If such attributes have to be used to disambiguate entries, multi-valued RDNs should be used, such that other attribute(s) be used for naming in addition to a common name.

The choice of RDN for humans will be influenced by cultural considerations. In many countries the best choice will be of the form *familiar-first-name surname.* Thus, *Steve Hardcastle-Kille* is preferred as the RDN choice for one of this document's co-authors, while *Stephen E. Hardcastle-Kille* is stored as an alternative commonName value. Sets of initials should not be concatenated into a single "word", but be separated by spaces and/or "." characters. Pragmatic choices will have to be made for other cultures.

3.6 Application Entities

The guidelines of X.521 should be followed, in that the application entity should always be named relative to an Organisation or Organisational Unit. The application process will often correspond to a system or host. In this case, the application entities should be named by Common Names which identify the service (e.g., "FTAM Service"). In cases where there is no useful distinction between application process and application entity, the application process may be omitted (This is often done for DSAs in the current pilot).

4 Multinational Organisations

The standard says that only *international* organisations may be placed under the root of the DIT. This implies that multi-national organisations must be represented as a number of separate entries underneath country or locality entries. This structure makes it more awkward to use X.500 within a multi-national to provide an internal organisational directory, as the data is now spread widely throughout the DIT, rather than all being grouped within a single sub-tree. Many people have expressed the view that this restriction is a severe limitation of X.500, and argue that the intentions of the standard should be ignored in this respect. This note argues, though, that the standard should be followed.

No attempt to precisely define *multinational organisation* is essayed here. Instead, the observation is made that the term is applied to a variety of organisational structures, where an organisation operates in more than one country. This suggests that a variety of DIT structures may be appropriate to accommodate these different organisational structures. This document suggests three approaches, and notes some of the characteristics associated with each of these approaches.

Before considering the approaches, it is worth bearing in mind again that a major aim in the choice of a DIT structure is to facilitate querying, and that approaches which militate against this should be avoided wherever possible.

4.1 The multi-national as a single entity

In many cases, a multi-national organisation will operate with a highly centralised structure. While the organisation may have large operations in a number of countries, the organisation is strongly controlled from the centre

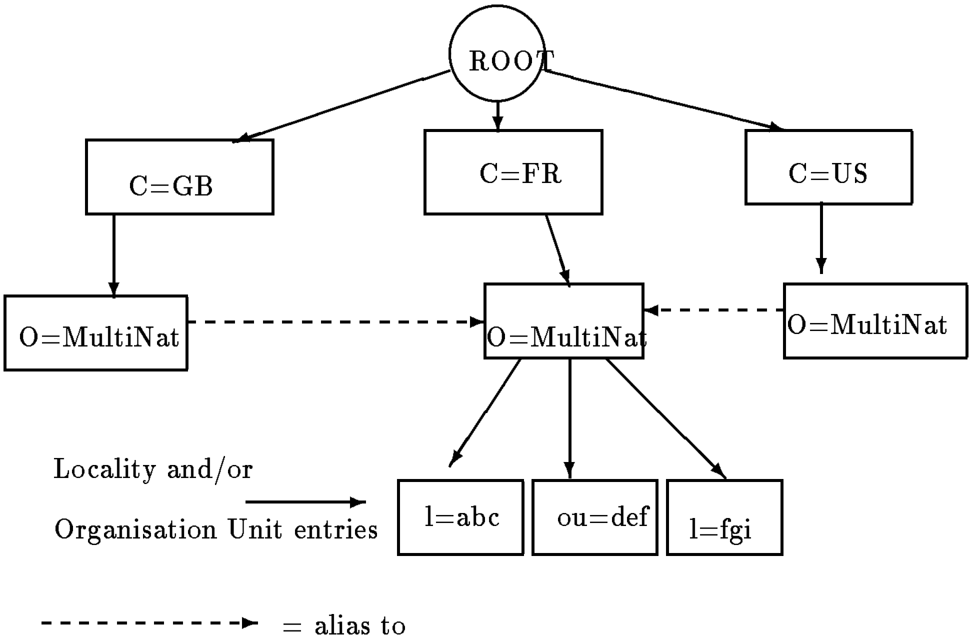

Figure 1: The multi-national as a single entity

and the disparate parts of the organisation exist only as limbs of the main organisation. In such a situation, the model shown in figure 1 may be the best choice. The organisation's entries all exist under a single sub-tree. The organisational structure beneath the organisation entry should reflect the perceived structure of the organisation, and so no recommendations on this matter can be made here. To assist the person querying the directory, alias entries should be created for all countries where the organisation operates.

4.2 The multi-national as a loose confederation

Another common model of organisational structure is that where a multi-national consists of a number of national entities, which are in large part independent of both sibling national entities, and of any central entity. In such cases, the model shown in Figure 2 may be a better choice. Organisational entries exist within each country, and only that country's localities and organisational units appear directly beneath the appropriate organisational entry. Some binding together of the various parts of the organisation can be achieved by the use of aliases for localities and organisational units, and this can be done in a highly flexible fashion. In some cases, the national view might not contain all branches of the company, as illustrated in Figure 2.

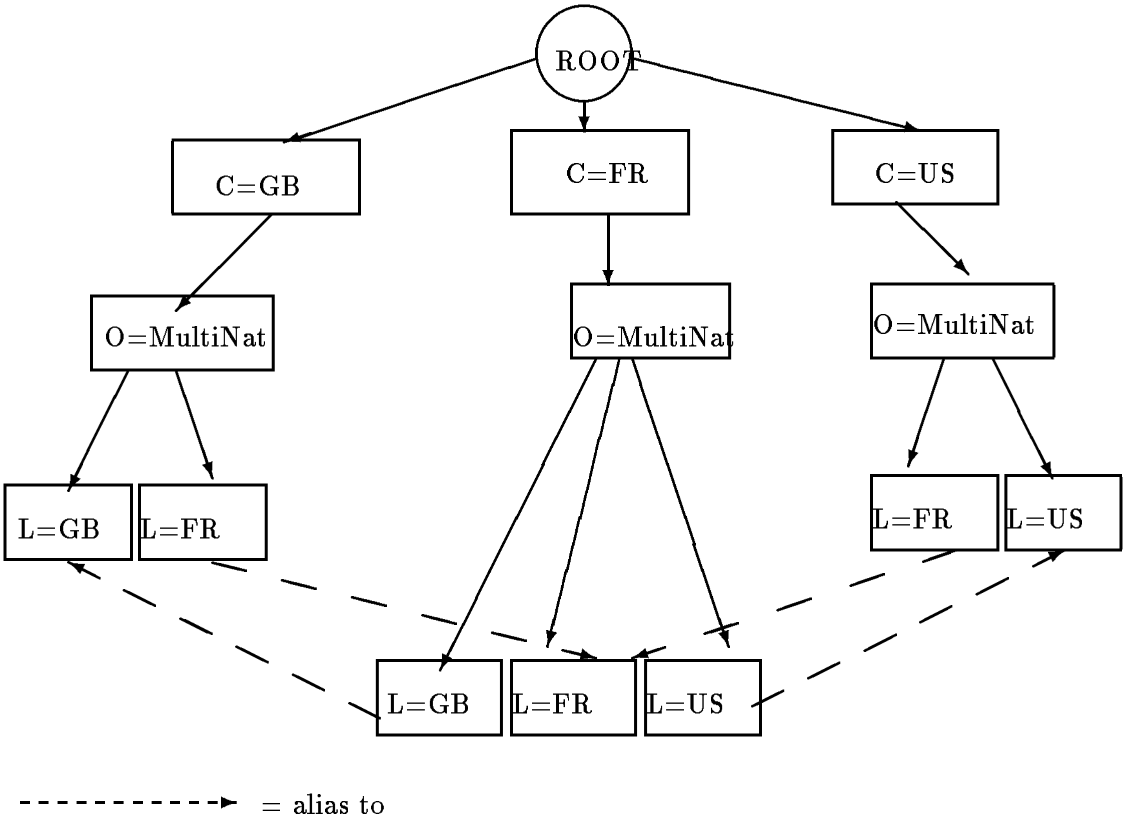

Figure 2: The multi-national as a loose confederation

4.3 Loosely linked DIT sub-trees

A third approach is to avoid aliasing altogether, and to use the looser binding provided by an attribute such as *seeAlso*. This approach treats all parts of an organisation as essentially separate. A unified view of the organisation can only be achieved by user interfaces choosing to follow the seeAlso links. This is a key difference with aliasing, where decisions to follow links may be specified within the protocol. (Note that it may be better to specify another attribute for this purpose, as seeAlso is likely to be used for a wide variety of purposes.)

4.4 Summary of advantages and disadvantages of the above approaches

Providing an internal directory All the above methods can be used to provide an internal directory. In the first two cases, the linkage to other parts of the organisation can be followed by the protocol and thus organisation-wide searches can be achieved by single X.500 operations. In the last case, interfaces would have to "know" to follow the soft links indicated by the seeAlso attribute.

Impact on naming In the single-entity model, all DNs within the organisation will be under one country. It could be argued that this will often result in rather "unnatural" naming. In the loose-confederation model, DNs are more natural, although the need to disambiguate between organisational units and localities on an international, rather than just a national, basis may have some impact on the choice of names. For example, it may be necessary to add in an extra level of organisational unit or locality information. In the loosely-linked model, there is no impact on naming at all.

Views of the organisation The first method provides a unique view of the organisation. The loose confederacy allows for a variety of views of the organisation. The view from the centre of the organisation may well be that all constituent organisations should be seen as part of the main organisation, whereas other parts of the organisation may only be interested in the organisation's centre and a few of its sibling organisations. The third model gives an equally flexible view of organisational structures.

Lookup performance All methods should perform reasonably well, providing information is held, or at least replicated, within a single DSA.

5 Miscellany

This section draws attention to two areas which frequently provoke questions, and where it is felt that a consistent approach will be useful.

5.1 Schema consistency of aliases

According to the letter of the standard, an alias may point at any entry. It is beneficial for aliases to be "schema consistent". The following two checks should be made:

1. The Relative Distinguished Name of the alias should be a valid Relative Distinguished Name of the entry.

2. If the entry (aliased object) were placed where the alias is, there should be no schema violation.

5.2 Organisational Units

There is a problem that many organisations can be either organisations *or* organisational units, dependent on the location in the DIT (with aliases giving the alternate names). For example, an organisation may be an independent national organisation and also an organisational unit of a parent organisation. To achieve this, it is important to allow an entry to be of *both* object class organisation and of object class organisational unit.

References

[BHK91] P. Barker and S.E. Hardcastle-Kille. The COSINE and Internet X.500 schema. Request for Comments RFC 1274, Department of Computer Science, University College London, November 1991.

[For91] The North American Directory Forum. A Naming Scheme for C=US, September 1991. Also NADF-175.

[HK91] S.E. Hardcastle-Kille. X.500 and domains. Request for Comments RFC 1279, Department of Computer Science, University College London, November 1991.

6 Security Considerations

Security considerations are not discussed in this INTERNET–DRAFT .

7 Author's Addresses

Paul Barker and Steve Hardcastle-Kille

Department of Computer Science
University College London
Gower Street
WC1E 6BT
England

Phone:+44-71-380-7366 (Barker)
Phone:+44-71-380-7294 (Kille)

EMail: P.Barker@CS.UCL.AC.UK
EMail: S.Kille@CS.UCL.AC.UK