Network Working Group
INTERNET–DRAFT

S.E. Kille
University College London
March 1991

DSA Naming

Status of this Memo

This INTERNET–DRAFT describes a few problems with DSA Naming as currently deployed in pilot exercises, and suggests a new approach. This approach is suggested for use in the Internet Directory Pilot.

I believe this note to be an important step forward, as it begins to evolve a clear model of a Directory Management Domain.

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

1 Problem Statement

There are a number of difficulties with the approach being used in QUIPU based pilots, which is described in the original QUIPU Design and in the QUIPU documentation. The problems are:

1. It is hard to achieve very high replication of knowledge information, as this is very widely spread. This is particularly true of secondary DSAs, which are named inside an organisation. Due to the sensitivity of sibling information, replication of this data outside the organisation is less likely.
2. There is a need to give DSAs more reasonable names. The current names are too high up the tree to indicate the role of the DSA. This can lead to management confusion.
3. There is too much DIT clutter — more endangered South American animals than organisations!
4. There is no real concept of DMD (Directory Management Domain) authority. This needs to be introduced, in order to allow DSAs to be named by the authority that runs them.

The following are good properties of the current approach, which must not be lost in any changes.

1. The directory should be used to look up the presentation addresses of DSAs as described in [Kil91] and not require them to be manually associated with all knowledge references as implied in X.500.
2. The approach must be deadlock free. This will not be the case if DSAs named in a careless manner, and 1. is followed.
3. There should be no manual replication of information, which is implied by the presentation address configuration approach of the standard.

2 What is a DMD

A DMD is a special type of Organisation, which has two roles.

1. Managing directory configuration information.
2. Operating one or more DSAs (A DSA belongs to the DMD that operates it)

A common case is for an Organisation (or Organisational Unit) to operate its own DSAs. In this case the Organisation has a dual role as both a DMD

and an Organisation. Such a DMD is primarily responsible for operating a number of DSAs, and some minimal configuration information associated with those DSAs. For most organisations, the number of DSAs will be small, although a large organisation may run quite a large number.

Another common case is for for a DMD to be primarily concerned with managing configuration information between a large number of DSAs, and perhaps operating a small number of DSAs in order to achieve this. A research network pilot is an example of such a DMD. Such a DMD will:

- Collect and distribute configuration information
- Collect and distribute user data (replication)

An Administrative DMD (as distinct from a Private DMD) is one which systematically maintains knowledge of the root configuration within the DMD. This will only be a useful distinction if this knowledge is restricted (it is not in the current pilot environment). For now it should be assumed that this information is widely replicated, as adequate performance of the global directory depends on it.

A DMD may have a peer relationship with other DMDs. There may also be a hierarchical relationship. This is appropriate for a DMD which wishes to obtain all of its configuration information and replicated from a parent DMD. A hierarchical relationship does not preclude other relationships between DMDs (e.g., to mutually replicate data).

A DMD has the following responsibilities:

- It must operate at least one DSA, to manage the information associated with the DMD.
- It may operate DSAs to store information for Organisations supported by the DMD. These DSAs are named within the DMD
- Allocate sub-DMDs. The EDB for such a DMD should be mastered in the DMD DSA
- Manage the replication of knowledge information within the DMD, and make it available for external replication.

3 Naming DMDs and DSAs

An example is given in Figure 1.

A basic decision is to name DSAs as entries beneath the associated DMD. DMDs are arranged in a separate hierarchy.

Deadlock due to needing to locate a DSA to determine its own presentation address is prevented, because of the following conditions:

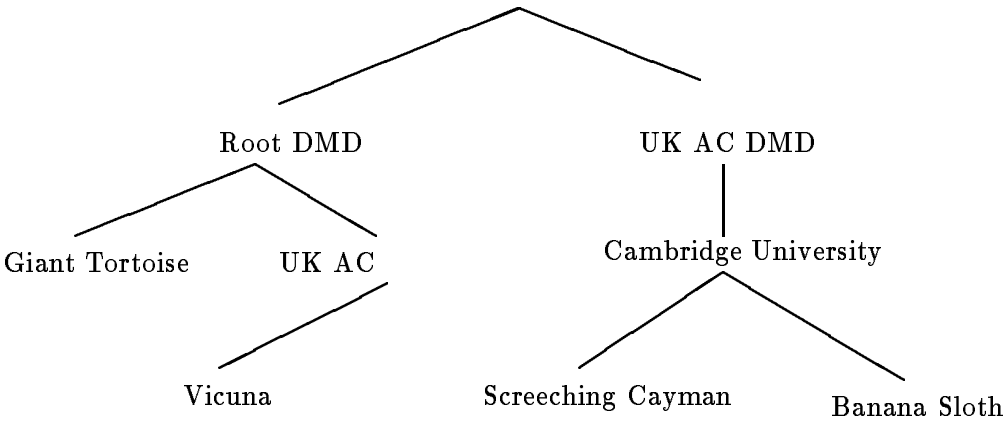

Figure 1: Example DMD Hierarchy

- DSAs are named in a separate DMD hierarchy. This groups configuration information in one area, which will facilitate replicating it widely.
- For information not in the DMD hierarchy, this separation prevents deadlock
- Within the DMD hierarchy, special rules apply to prevent deadlock. This often leads to configuration information being mastered remotely.

DMDs will be subclasses of Organisation or Organisational Unit (the latter will be used when the parent of a DMD is another DMD). This will be defined formally by making the DMD object class subclass of top, as neither Organisation or Organisational Unit is required. There is an optional attribute to indicate organisations associated with the DMD. Object Identifiers will be defined in the COSINE and Internet Naming Architecture [BK90].

```
AssociatedOrganisation ATTRIBUTE
WITH ATTRIBUTE-SYNTAX DistinguishedNameSyntax

AssociatedDMD ATTRIBUTE
WITH ATTRIBUTE-SYNTAX DistinguishedNameSyntax

DMD OBJECT-CLASS
SUBCLASS OF top
MAY CONTAIN {associatedOrganisation,
   associatedDMD}
```

There will be a number of top level DMDs. In the longer term, there will be a small number of nationally allocated DMDs, probably allocated under the country level. This will lead to many DMD trees, and not just one. This will raise more problems for the support of multinational organisation. For

now, we will avoid this complexity, and resist the temptation to introduce it too early. To avoid deadlock in looking up DSAs, the DSA which masters an Entry Data Block (as defined in [Kil91]) should have a name which has the same number of components as or less components than the EDB name, or the DSA should be in the Root DMD.

There will be a special top level DMD, with the name "Root DMD". This DMD will contain the names of DSAs which master top level DMDs. The well known DSA: "Giant Tortoise, Root DMD" will hold the master EDBs for the root and for "Root DMD". This will allow a way in to the system, in the same manner as at present.

Typical top level DMDs will correspond to national pilots. That is, there will be a small number for each country. For example: "Internet Pilot DMD'; "UK AC Pilot DMD"; "Paradise DMD". Each node of the DIT, and in particular each country, will need to be mastered by a single DMD[1].

Where a DMD allocates (a large number) of sub-DMDs, it is important that this is done in a separate branch of the tree[2]. Such a DMD will have two entries in the DIT. The part of the tree which contains the DSAs will either be a sub-DMD of the root DMD or have the string ␣DSAs at the end. The two parts of the DMD will be linked by the AssociatedDMD Attribute. For top level DMDs, the part of the DMD which names its DSAs should be under the Root DMD.

An example use of the hierarchy might be to name a DSA:

> Screeching Cayman, Cambridge University DMD,
> UK AC Pilot DMD

Note that although DMDs are named hierarchically, individual DMDs are autonomous. The naming of DSAs within a DMD is up to the DMD. The QUIPU guideline of using South American Wildlife will remain (for the QUIPU implementation), but this may be inappropriate for some DMDs.

In many cases, there will be a 1:1 mapping between an Organisation and DMD. In this case, it is recommended that the DMD is named by appending ␣DMD to the organisation name if it is a top level DMD or by using the organisation name if it is a sub-DMD.

References

[BK90] P. Barker and S.E. Kille. The COSINE and Internet X.500 nam-

[1] The introduction of real NSSRs would be a way around this, but there would need to be significant additional machinery in place.

[2] This is so that a DMD can manage is own sub-DMDs in one of its own DSAs. This separation is forced by the approach taken to replication. Another consequence is that a DMDs (usually small number of) DSAs will have their entries mastered in the DSA of another DMD. They will still be managed by the DMD which owns them.

ing architecture, December 1990. Internet Draft: draft-ietf-osids-cosinex500-02.txt.

[Kil91] S.E. Kille. Replication requirement to provide an internet directory using X.500, January 1991. Internet Draft: draft-ietf-osids-replsoln-01.txt, ps.

4 Security Considerations

Security issues are not considered in this INTERNET–DRAFT .

5 Author's Address

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

Phone: +44-71-380-7294

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