#18

# Managing the X.500 Root Naming Context

David Chadwick

This paper was written by David Chadwick (University of Salford), in his role as consultant to DANTE on NameFLOW-Paradise directory development. It has been submitted as an Internet-Draft (January 1996).

*DANTE IN PRINT* is a track record of papers and articles published by, or on behalf of DANTE. HTML and Postscript versions are available from: http://www.dante.net/pubs/dip

For more information about DANTE or *DANTE IN PRINT* please contact:

DANTE
Francis House
112 Hills Road
Cambridge CB2 1PQ
United Kingdom

Tel: +44 1223 302 992
Fax: +44 1223 303 005
E-mail: dante@dante.org.uk

# Managing the X.500 Root Naming Context

David Chadwick

## Abstract

The X.500 Standard [X.500 93] has the concept of first level DSAs, whose administrators must collectively manage the root naming context through bi-lateral agreements or other private means which are outside the scope of the Standard.

The NameFLOW-Paradise X.500 service has an established procedure for managing the root naming context, which currently uses Quipu proprietary replication mechanisms and a root DSA. The benefits that derive from this are twofold:

- firstly it is much easier to co-ordinate the management of the root context information, when there is a central point of administration,
- secondly the performance of one-level Search operations is greatly improved because the Quipu distribution and replication mechanism does not have a restriction that exists in the 1988 and 1993 Standard.

The NameFLOW-Paradise service is moving towards 1993 Standard replication protocols and wants to standardise the protocol and procedure for managing the root naming context which will be based on 1993 Standard protocols. Such a protocol and procedure will be useful to private X.500 domains as well as to the Internet X.500 public domain. It is imperative that overall system performance is not degraded by this transition.

This document describes the use of 1993 Standard protocols for managing the root context. Whilst the ASN.1 is compatible with that of the Standard, the actual settings of the parameters are supplementary to that of the Standard.

This paper results from the work done by the author as a consultant to DANTE for NameFLOW-Paradise directory development. His e-mail address is: D.W.Chadwick@iti.salford.ac.uk

## 1. Introduction

The NameFLOW-Paradise service has a proprietary way of managing the set of first level DSAs and the root naming context. There is a single root DSA (Giant Tortoise) which holds all of the country entries, and the country entries are then replicated to every country (first level) DSA by Quipu replication [RFC 1276] from the root DSA. The root DSA is not a feature of the X.500 Standard [X.500 93]. It was introduced because of the non- standard nature of the original Quipu knowledge model (also described in RFC 1276). However, it does have significant advantages both in managing the root naming context and in the performance of one-level Searches of the root. Performance is increased because each country DSA holds all the entry information of every country.

By comparison, the 1988 Standard root context which is replicated to all the country DSAs, only holds knowledge information and a boolean (to say if the entry is an alias or not) for each country entry. This is sufficient to perform a List operation, but not a one-level Search operation. When access controls were added to the 1993 Standard, the root context information was increased (erroneously as it happens - this is the subject of defect report 140 - see Annex 1) to hold the access controls for each country entry, but a note in the Standard restricted its use to the List operation, in order to remain compatible with the 1988 edition of the Standard.

## 2. Migration Plan

The NameFLOW-Paradise service is now migrating to 1993 Standard [X.500 93] conforming products, and it is essential to replace the Quipu replication protocol with the 1993 shadowing and operational binding protocols, but without losing the performance improvement that has been gained for one-level Searches.

It is still the intention of the NameFLOW-Paradise service to have one master root DSA. This

root DSA will not support user Directory operations via the DAP or the DSP, but each country (first level) DSA will be able to shadow the root context from this root DSA, using the DISP. Each first level DSA then only needs to have one bi lateral agreement, between itself and the root DSA. This agreement will ensure that the first level DSA keeps the root DSA up to date with its country level information, and in turn, that the root DSA keeps the first level DSA up to date with the complete root naming context. When a new first level DSA comes on line, it only needs to establish a bi-lateral agreement with the root DSA, in order to obtain the complete root context.

This is a much easier configuration to manage than simply a set of first level DSAs without a root DSA, as suggested in the Standard. In this case each first level DSA must have bi-lateral agreements with all of the other first level DSAs. When a new first level DSA comes on line, it must establish agreements with all the existing first level DSAs. As the number of first level DSAs grows, the process becomes unmanageable.

However, it is also important to increase the amount of information that is held about every country entry, so that a one-level Search operation can be performed in each first level DSA, without it needing to chain or refer the operation to all the other first level DSAs (as is currently the case with a Standard conformant system.)

## 3. Technical Solution

*3.1* The solution is three-fold. Firstly, create a root DSA, and establish hierarchical operational bindings between it and each master first level DSA (3.2). Secondly, the Standard is enhanced to allow extra information to be carried to the root DSA via the HOB, and for this information to be used for one-level Search operations (3.3). Thirdly, each master first level DSA enters into a shadowing agreement with the root DSA, to shadow the enlarged root context information. In this way each first level DSA is then capable of independently performing List and one level Search operations, and name resolving to all other first level DSAs (3.4).

(Note 1. It is strongly recommended that non-specific subordinate references should not be allowed in the root context for efficiency reasons. This is directed by the European functional standard [ENV 41215] and the NADF standing document [NADF 7]. It is also preferred by the International Standardized Profile [ISP 10615-6].)

(Note 2. It is recognised that manufacturers are taking a phased approach to implementing the features of the 1993 Standard, and are usually implementing the DISP prior to the DOP. For this reason, section 4 details an interim solution that relies entirely on the DISP for populating the root DSA.)

*3.2* Each master first level DSA will have a hierarchical operational binding with the root DSA of the domain. Each master first level DSA will master one or more first level entries. The hierarchical operational binding will keep the appropriate subordinate reference(s) (of category shadow and master) up to date, as well as the other entry information that is needed for one-level Search operations (such as access controls, and attributes used in filtering).

Whilst hierarchical agreements are standardised, this particular novel use of a HOB is not specifically recognised in the Standard. Although the ASN.1 will support it, there is no supporting text in the Standard. The following text supplements that in the Standard, and describes how a first level DSA may have a hierarchical operational binding with the root DSA of its domain.

"Clause 24 of ISO / IEC 9594-4:1993 shall also apply when a first level DSA is a subordinate DSA, and the root DSA of the domain is the superior DSA. The naming context held by the superior (root) DSA is the root naming context (or root context - the terms are synonymous) of the domain. The root context consists of the root entry of the DIT (which is empty) plus a complete set of subordinate DSEs, one for each first level naming context in the domain. The subordinate DSEs will contain, in addition to specific knowledge attribute values of category master and shadow, sufficient attributes, including access control information, to allow List and one-level Search operations to be performed on them.

In clause 24.1.2, the DistinguishedName of the immediateSuperior component of HierarchicalAgreement shall be null."

*3.3* The ASN.1 of hierarchical operational bindings already allows any attributes to be passed from the subordinate DSA to the superior DSA (SubordinateToSuperior parameter in clause 24.1.4.2 of X.518). However, a note in the Standard limits this to those which are required to perform a List operation. The UK submitted a ballot comment to the PDAM on Minor Extensions to

OSI Directory Service to support User Requirements, to remove this restriction, so that the attributes may also be used for a one-level Search operation. This amendment has been accepted, and the restriction has been removed in the current Draft Amendment to the 1996 version of the Standard [DAM User].

1993 implementations of X.500 conforming to this RFC, shall also remove this restriction.

*3.4* Each master first level DSA will enter into a shadowing agreement with the root DSA, for the purpose of shadowing the root naming context.

The 1993 edition of the Standard explicitly recognises that there can be master and shadow first level DSAs (X.501 Section 18.5). (The 1988 edition of the standard does not explicitly recognise this, since it does not recognise shadowing.) A shadow first level DSA holds a copy of the root context, provided by a master first level DSA. In addition it holds shadow copies of the (one or more) country entries that the master first level DSA holds. There is currently an outstanding defect report [UK 142] on the 1993 Standard to clarify how a shadowing agreement is established between first level DSAs. Once this has been ratified, the only additional text needed in order to establish a shadowing agreement between the root DSA and a master first level DSA is as follows:

"When clause 9.2 of ISO/IEC 9594-9:1993 is applied to the shadowing of the root context by a first level DSA from the root DSA of a domain, then UnitOfReplication shall be set as follows:

contextPrefix of AreaSpecification shall be null,

replicationArea of AreaSpecification shall be set to
SEQUENCE {
specificExclusions [1] SET OF {
chopBefore [0] FirstLevelEntry},
maximum [3] 1 }

where FirstLevelEntry is the RDN of a first level entry (e.g. country, locality or international organisation) held by the master first level DSA. specificExclusions shall contain one FirstLevelEntry for each first level entry mastered by this DSA, attributes of UnitofReplication shall be an empty SET OF SEQUENCE, knowledge of UnitofReplication shall be set to both (shadow and master).

In other words, the information that will be replicated will be an empty root entry plus all the attributes of the complete set of subordinate DSEs of the root, excluding the DSEs that the first level DSA already masters."

Note that the maximum component of replicationArea, although not strictly necessary, is there for pragmatic reasons, for example, where a community of users wish to use the root DSA to hold some country specific entries.

## 4. Interim Solution

*4.1* The interim solution may be of use to systems which do not yet support the DOP for managing hierarchical operational bindings.

The interim solution comprises of two replacement steps for HOB establishment between the root DSA and master first level DSAs. Step one (4.2) allows the root DSA to shadow first level entries from a master first level DSA. Step two (4.3) requires either the root DSA administrator or the root DSA implementation to massage the shadow first level entries so that they appear to have been created by a HOB. Managing the root context then continues as in 3.4 above.

*4.2* The hierarchical operational binding between the root DSA and a master first level DSA can be replaced by a set of "spot" shadowing agreements, in which the first level DSA acts as the supplier, and the root DSA as the consumer. Each "spot" shadowing agreement replicates a first level entry which is mastered by the first level DSA. The UnitOfReplication shall be set as follows:

contextPrefix of AreaSpecification shall be FirstLevelEntry,

replicationArea of AreaSpecification shall be set to
SEQUENCE {
specificExclusions [1] SET OF {
chopAfter [1] {null} } }

where FirstLevelEntry is the Distinguished Name of a first level entry (e.g. country, locality or international organisation) held by the master first level DSA.

attributes of UnitofReplication shall be an empty SET OF SEQUENCE,

knowledge of UnitofReplication shall be absent.

*4.3* The root DSA administrator, or the root DSA implementation (suitably tailored) must then administratively update each shadowed first level entry, so that they appear to have been created by a HOB, i.e. it is necessary to add a subordinate

reference to each one of them. The subordinate reference will point to the respective master first level DSA, and will comprise of a specific knowledge attribute, and the DSE bit of type subr being set. The contents of the specific knowledge attribute can be created from the contents of the supplier knowledge attribute already present in the first level entry and created by the "spot" shadowing agreement.

*Appendix 1 Solution Text of Defect Reports submitted to ISO/ITU-T by the UK*

### Defect Report 140

Nature of Defect

In section 24.1.4.2 it is defined that the SubordinateToSuperior parameter of a HOB can pass an entryInfo parameter. This should contain entryACI which may be used in the resolution of the List operation.

This is not correct as the prescriptive ACI from the relevant subentries is also required in the superior DSA.

Solution Proposed by Source

It is proposed that the following is added to the SubordinateToSuperior SEQUENCE of section 24.1.4.2 of X.518:

```
subentries  [2] SET OF SubentryInfo OPTIONAL
```

This is used to pass the relevant subentries from the subordinate to the superior. This is similar to the way subentry information is passed in the SuperiorToSubordinate parameter defined in 24.1.4.1.

### Defect Report 142

Nature of Defect

The text which describes AreaSpecification in clause 9.2 of X.525 is completely general. However, for the special case of replicating first level knowledge references between first level DSAs, a clarifying sentence should be added.

Solution Proposed by Source

In Section 9.2, under the ASN.1, after the description of area, and before the description of SubtreeSpecification, add the sentence:

"For the case where a DSA is shadowing first level knowledge from a first level DSA, the contextPrefix component is empty."

*Acknowledgement*

The author would like to thank Nexor, who reviewed the document in detail and provided valuable comments, and who first suggested the Interim Solution as a stop-gap measure until the DOP is widely implemented.

*References*

[RFC 1276] Kille, S., "Replication and Distributed Operations extensions to provide an Internet Directory using X.500", UCL, November 1991.

[X.500 93]

| | |
|---|---|
| X.500 \| 9594.Part 1 | Overview of Concepts, Models and Services |
| X.501 \| 9594.Part 2 | Models |
| X.511 \| 9594.Part 3 | Abstract Service Definition |
| X.518 \| 9594.Part 4 | Procedures for Distributed Operations |
| X.519 \| 9594.Part 5 | Protocol Specifications |
| X.520 \| 9594.Part 6 | Selected Attribute Types |
| X.521 \| 9594.Part 7 | Selected Object Classes |
| X.509 \| 9594.Part 8 | Authentication Framework |
| X.525 \| 9594.Part 9 | Replication |

[ENV 41215] "Behaviour of DSAs for Distributed Operations", European Pre-Standard, Dec 1992

[ISP 10615-6] "DSA Support of Distributed Operations", 5th draft pDISP, Oct 1994

[NADF 7] SD-7 "Mapping the North American DIT onto Directory Management Domains", North American Directory Forum, V 8.0, Jan 1993

[UK 142] Defect report number 142, submitted by the UK to ISO, March 1995. (Proposed solution text included in Annex 1)

[DAM User] Draft Amendments on Minor Extensions to OSI Directory Service to support User Requirements, August 1995.