

draft-ietf-ids-root-naming-01.txt Expires: March 20 1997

| | |
|---|------------------------|
| IDS Working Group | David Chadwick |
| Internet-Draft | University of Salford |
| DANTE IN PRINT 18 | September 20 1996 |
| draft-ietf-ids-root-naming-01.txt | Expires: March 20 1996 |

Managing the X.500 Root Naming Context

Status of this Memo

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".

To learn the current status of any Internet-Draft, please check the lid-abstracts.txt listing contained in the Internet-Drafts Shadow Directories on ds.internic.net, nic.nordu.net, ftp.nisc.sri.com, or munnari.oz.au.

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 project is moving towards 1993 ISO 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 ISO 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.

Table of Contents

| | | |
|--------------------|---|----|
| 1 | Introduction | 2 |
| 2 | Migration Plan | 3 |
| 3 | Technical Solutions | 4 |
| 4 | The Fast Track Solution | 5 |
| 5 | The Slower Track Solution | 6 |
| 6 | The Long Term Solution | 7 |
| 7 | Security Considerations | 8 |
| 8 | Acknowledgments | 8 |
| 9 | References | 9 |
| 10 | Author's Address | 10 |
| Annex 1 | Solution Text of Defect Reports submitted to ISO/ITU-T by the UK | 10 |
| Annex 2 | Defect Report on 1993 Standard for Adding full ACIs to DISP for Subordinate References, so that Secure List Operation can be performed in Shadow DSAs | 11 |
| Annex 3 | Defect Report on 1997 Standard Proposing an Enhancement to the Shadowing Agreement in order to support 1 Level Searches in Shadow DSAs. | 13 |

[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 and other DSAs by Quipu replication [RFC 1276] from the root DSA. In June 1996 there were 770 DSAs replicating this information over the Internet. 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 an insecure 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 LDAP, 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 ISO Standard. In the Standard 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 conforming system.)

3 Technical Solutions

0.1 The solution at first appears to be relatively straight forward, and involves two steps. Firstly, create a root DSA, and establish hierarchical operational bindings using the DOP, between it and each master first level DSA. Secondly, 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.

3.2 Unfortunately there are a number of complications that inhibit a quick implementation of this solution. Firstly, few DSA suppliers have implemented the DOP. Secondly there are several defects in the Standard that currently stop the above solution from working.

3.3 At a meeting chaired by DANTE in the UK on 18 June 1996[Mins], at which several DSA suppliers were present, the following pragmatic technical solution was proposed. This comprises a fast track partial solution and a slower track fuller solution. Both the fast and slower tracks use the shadowing protocol (DISP) for both steps of the solution, and do not rely on the DOP to establish HOBs. The fast track solution, described in [section 4](#), will support knowledge distribution of the root context, and the (insecure) List operation of the root's subordinates. The List operation will be insecure because access control information will not be present in the shadow DSEs. (However, since it is generally thought that first level entries, in particular country entries, are publicly accessible, this is not considered to be a serious problem.) Suppliers expect to have the fast track solution available before the end of 1996. The slower track solution, described in [section 5](#), will in addition support fully secure one level Search and List operations of the root (without the need to chain to the master DSAs). Suppliers at the DANTE meeting did not realistically expect this to be in their products much sooner than mid 1998.

3.4 The long term solution, which relies on the DOP to establish HOBs, is described in [section 6](#) of this document.

(Note. 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].)

4 The Fast Track Solution

4.1 The fast track solution provides root knowledge collection and insecure List operations for first level DSAs, and will be of use to systems which do not yet support the DOP for managing hierarchical operational bindings. The fast track solution relies upon the DISP with very few changes to the 1993 edition of the Standard.

4.2 Each master first level DSA administrator will make available to the administrator of the root DSA, sufficient information to allow the root DSA to configure a subordinate reference to their DSA. In the simplest case, this can be via a telephone call, and the information comprises the access point of their DSA and the RDNs of the first level entries that they master.

4.3 Each master first level DSA enters 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 that are held in the root DSA excluding the DSEs that the first level DSA already masters, plus a complete set of subordinate reference."

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.

5 The Slower Track Solution

5.1 The slower track solution provides support for fully secure one level Search and List operations of the root in first level DSAs, and comprises of two steps for HOB establishment between the root DSA and master first level DSAs, using the DISP instead of the DOP. Step one, described in 5.3, allows the root DSA to shadow first level entries from a master first level DSA. Step two, described in 5.4, 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 4.3 above.

5.2 This solution requires two significant defects in the ISO Standard to be corrected. Firstly, access control information needs to be added to subordinate references in the DISP to allow the List operation to work securely in a shadowed DSA. (The ACI are held in both the subr DSE and in its subentry.) This requires a defect report on the 93 standard to be submitted. The text of this defect report (that has been submitted to ISO) is given in Annex 2.

Secondly, a new type of shadowing agreement will need to be established between the supplier and consumer DSAs, to copy subordinate entries rather than simply subordinate references, so that one level Search operations can work in the shadowing DSA. This procedure should have been part of the 1997 edition of the standard, but due to an omission is not. Consequently a defect report on the 1997 Standard has been submitted. The text of this defect report is given in Annex 3.

5.3 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.

5.4 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.

6 The Long Term Solution

6.1 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 (i.e. first level DSEs), one for each first level naming context in the domain, and their corresponding subentries. The first level DSEs and their subentries will contain, in addition to specific knowledge attribute values of category master and shadow, sufficient attributes and collective 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."

6.2 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 1993 edition of the Standard limits this to those which are required to perform a List operation. In the 1997 edition of the Standard [DAM User] this restriction has been removed, so that the attributes may also be used for a one-level Search operation.

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

7 Security Considerations

Security considerations are discussed in this memo in relation to List and one-level Search operations. Each DSE has access control information associated with it, and these must be adhered to when the operations are performed.

8 Acknowledgments

The author would like to thank DANTE, without whose funding this work would not have been possible.

The author would also like to thank Nexor, who reviewed the first version of this document in detail and provided valuable comments, and who first suggested the use of the DISP as a pragmatic

solution for HOB establishment until the DOP becomes widely implemented.

The author would also like to thank John Farrell from the ISODE

Consortium, Andrew Palk from Digital and Keith Richardson from ICL who attended the DANTE meeting, and contributed to the technical contents of the defect reports in Annexes 2 and 3.

9 References

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

[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

[Mins] "Notes of DANTE meeting to discuss Managing the Root Naming Context. 18 June 1996." D W Chadwick, circulated to IDS mailing list

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

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

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

[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

Page 9

[draft-ietf-ids-root-naming-01.txt](#)

Expires: March 20 1997

10 Author's Address

D W Chadwick
IT Institute
University of Salford
Salford
M5 4WT
England

Phone: +44 161 745 5351
Fax: +44 161 745 8169
E-mail: D.W.Chadwick@iti.salford.ac.uk

Annex 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.

Page 10

[draft-ietf-ids-root-naming-01.txt](#)

Expires: March 20 1997

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."

Annex 2 Defect Report on 1993 Standard for Adding full ACIs to DISP for Subordinate References, so that Secure List Operation can be performed in Shadow DSAs

Nature of Defect:

The List operation may be carried out in a superior DSA using subordinate reference information, providing that the fromEntry flag is set to false in the response. However, in order to do this securely, complete access control information is needed for the RDN of the subordinate entry. The existing text assumes that this is held in entry ACI (e.g. see 9.2.4.1 c) or in prescriptive ACI held in subentries above the DSE (e.g. see 9.2.4.1 b). In the case of a subordinate reference, the prescriptive ACI may be held below the DSE, if the subordinate reference points to a new administrative point. The shadowing document needs to make it clear that this can be the case, and needs to allow for this additional access control information to be shadowed.

A related defect report (140) has already suggested that this same omission should be added to operational bindings.

Solution Proposed by the Source:

All the following changes are to X.525|ISO 9594-9.

I) Insert the following text into 7.2.2.3, at the end of both the second paragraph and the first sentence of the third paragraph (after "appropriate knowledge"):
"and access control information."

II) Insert a new third paragraph into 7.2.2.3:
"If subordinate knowledge is supplied, and the supplying DSE (of type subr) is also of type admPoint, then the SDSE shall additionally be of type admPoint and the administrativeRole attribute shall be supplied. If such a DSE has any immediately subordinate subentries containing PrescriptiveACI relating to the administrative point, then they shall also be supplied as SDSEs in the shadowed information."

Page 11

[draft-ietf-ids-root-naming-01.txt](#)


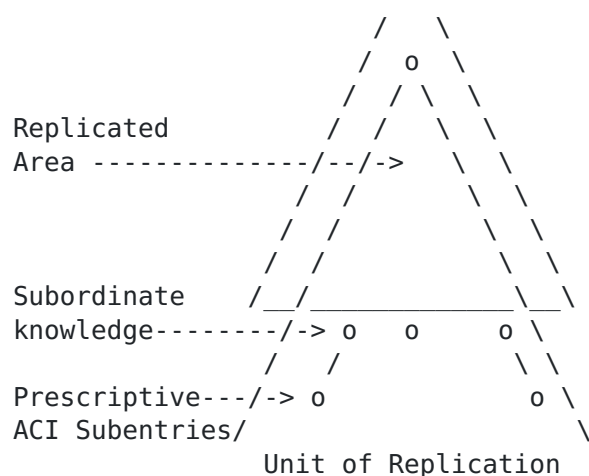
Expires: March 20 1997

Note. A DSE can be of type subr and admPoint in a superior DSA, when the naming context in the subordinate DSA is the start of a new administrative area."

III) Update figure 3 to show a subentry immediately below a subordinate reference. The subentry contains prescriptiveACI and is part of the shadowed information.

Etc.

·
/ \



ADDITIONS TO FIGURE 3, [SECTION 7.2](#), X.525

IV) Add supporting text to [section 7.2](#) in the paragraph after Figure 3. Insert after the sentence "Subordinate knowledge may also be replicated" the following sentences "Implicit in the Add supporting text to [section 7.2](#) in the paragraph after Figure 3. Insert after the sentence subordinate knowledge is the access control information which governs access to the RDN of the subordinate knowledge. When the subordinate entry is an

Page 12

[draft-ietf-ids-root-naming-01.txt](#)

Expires: March 20 1997

administrative point in another DSA, then part of this access control information may be held in prescriptiveACI subentries beneath the subordinate knowledge."

v) Add a new point d) to 9.2.4.1:

"if subordinate knowledge (not extended knowledge) is shadowed then any prescriptiveACI in subordinate subentries shall also be copied."

Annex 3 Defect Report on 1997 Standard Proposing an Enhancement to

the Shadowing Agreement in order to support 1 Level Searches in Shadow DSAs.

Nature of Defect:

The 1997 edition of the standard has allowed, for reasons of operational efficiency, one level Searches to be carried out in the superior DSA, when the actual entries are context prefixes in subordinate DSAs. The HOBs have been extended to allow this entry information to be carried up to the superior DSA. Unfortunately, we forgot to add the corresponding text to Part 9, so that shadow DSAs are able to copy this additional information from the supplier DSA. This defect report proposes the additional text for Part 9.

Solution Proposed by the Source:

All the following changes are to X.525|ISO 9594-9.

I) [Section 9.2](#), add a new subordinates parameter to UnitOfReplication, viz:

```
UnitOfReplication ::= SEQUENCE{  
  area                AreaSpecification,  
  attributes           AttributeSelection,  
  knowledge            Knowledge OPTIONAL,  
  subordinates        BOOLEAN DEFAULT FALSE }
```

subordinates is used to indicate that subordinate entries, rather than simply subordinate references, are to be copied to the consumer DSA. subordinates may only be TRUE if knowledge is requested and extendedKnowledge is FALSE.

II) Insert a new fourth paragraph (assuming previous defect for List was accepted) into 7.2.2.3:

"If subordinates is specified, then the supplier shall send subordinate entries rather than subordinate references, and the SDSEs will be of type subr, entry and cp. The subordinate entries will contain attributes according to the attribute selection.

In addition, if the supplying DSE is of type admPoint, then the SDSE shall additionally be of type admPoint and the administrativeRole attribute shall be supplied. All appropriate subentries below the admPoint DSE shall also be supplied as SDSEs in the shadowed information."

