

Meetings

- [London meeting](#) (January 2000)
- [Utrecht meeting](#) (January 1999)
- [Brussels meeting](#) (June 1996)
- [Salford meeting](#) (March 1995)
- [London meeting](#) (September 1995)

The NameFLOW-Paradise Managers meetings are open and free, without any charges.

- [LDAP pilot meeting](#) at London Heathrow (May 1998)
- Indexing: [Index meeting at the University of Brunel](#). Trip notes of a first orientation.
- Managing the Root Context: [minutes](#) of the DANTE meeting. (June 1996)
- Managing the Root Context: Report on [Geneva](#) X.500 ISO/ITU-T Meeting. (April 1996)

Nameflow Meeting January 2000

London, 20 January 2000

Time:

10h00 - 17h00

Where:

Brunel University

- [Invitation](#)
- [Agenda](#)

Move to [[DANTE](#) | [FLOW Services](#) | [NameFLOW](#)]

Contact: Konstantin.Chuguev@dante.org.uk

Nameflow Meeting January 1999

Utrecht, 18 January 1999

Time:

09h00 - 17h00

Where:

Surfnet Office

Radboudburcht kantoorgebouw

Hoog Catharijne

Postbus 19035

3501 DA Utrecht

Tel: (030) 2 305 305

Fax: (030) 2 305 329

- [Invitation](#)
- [Agenda](#)

Move to [[DANTE|FLOW Services|NameFLOW](#)]

Contact: Konstantin.Chuguev@dante.org.uk

NameFLOW-Paradise Directory Managers Meeting

Brussels, 11 June 1996

Time:

9h00 - 12h00

Where:

Sheraton Brussels

Place Rogier 3

B-1210 Brussels, Belgium

tel + 32 2 224 3111

fax + 32 2 224 3456

The minutes of this meeting: click [here](#)

- [Agenda](#)
- [Registration Form](#)
- [Minutes London Meeting](#)
- [Minutes Closed Customers Meeting \(London/ULCC\)](#)

Move to [[DANTE](#) | [FLOW Services](#) | [NameFLOW](#)]

Contact: Konstantin.Chuguev@dante.org.uk

March 1995, Salford University

- [1995 Service Specification](#)
- [Minutes \[ps\]](#)

September 1995, ULCC - London

- [Minutes \[ps\]](#)
- [Minutes of the \(closed\) customer meeting](#)

Nameflow Meeting on the LDAP Pilot

London 29 May 1998

Time:

10h30 - 16h00

Where:

Novotel London Heathrow

Junction 4 M4

Cherry Lane

West Drayton

Middlesex

UB7 9HB

United Kingdom

tel +44 01895 431 431

fax +44 01895 431 221

- [Invitation](#)
- [Agenda](#)
- [Minutes](#)

Move to [[DANTE](#)][FLOW Services](#)[NameFLOW](#)]

Contact: Konstantin.Chuguev@dante.org.uk

Trip note: Index Meeting at Brunel University

Date: 21 October 1996

Time: 10h00 - 16h00

Place: University of Brunel, United Kingdom

Participants:

AF Andrew Findlay
DM Damanjit Mahl
JF John Farell
NM Nils Meulemans
PJ Peter Jurg
RH Roland Hedberg
VB Vincent Berkhout
OD Other Dude from Brunel who's name I forgot to write down. sorry.

Introduction

Introduction of people around the table.

VB gave a short introduction why we need indexing, firstly as an X.500 performance improvement and secondly as a potential solution for the Root Context problem as described by David Chadwick (see DANTE in PRINT #18). The possibilities for indexing have been addressed by DANTE several times and one of the initiatives (another paper written by David Chadwick) is called Index.500 (see for more information DIP #13 and DIP #19). The paper deals with indexing and sets out to be a true X.500 based solution. VB suggests to separate the following three sub-topics

1. Index generation
 - a. none
 - b. do-it-yourself
 - c. being indexed
2. Index storing ("Granularity" and centroid construction)
3. Query resolving

RH continued after this. RH is co-chair of the FIND working group at the IETF. David Chadwick describes a solution and he actually implemented something before writing it down. RH has a strong Whois++ background and also built a Ph->X500 gateway. His index server works for all three directory types and is based on a SLAPD server. He build a special web client for index querying. He uses SLAPD to look at the performance with about 90.000 -100.000 entries (indices). The current source is in PERL but he has plans to rewrite it in C.

The Swedish government is forcing "service providers" to set up a public directory and if they fail there will be legislation for this. They believe that index servers can be used as "common solution" between different directory types. This could also be a solution for the Multiple Service Provider problem (see DIP #13).

NM gave a presentation during the last NameFLOW-Paradise managers meeting. There he explained that his model is based on the principle of Whois++ and the way it builds centroids. The principle is building centroids from centroids and store them in the Directory. The centroid builder uses a base (e.g. c=BE) and does a sub tree search. The index builder goes to the leaves and collapses every level into a centroid. These centroids subsequently collapse while going up the levels until the base is reached. It typically generates a single index for instance for surnames. This means that if a query is fired it actually uses "search space pruning" rather than the "look there for the answer" solution. As reference to the original entry the seeAlso attribute is used. One problem with this approach is that one has to store the index one level up. Nils did this for a great number of countries (32?) and used a dedicated DSA to solve this problem. The problem is now that it takes very long to start the QUIPU DSA because of the large number of entries.

JF has three years experience implementing DSAs and knows that X.500 has its flaws. He would like to introduce an index feature to X.500 DSAs.

DM has experience with building DUAs (WLU) for the University of Brunel. They have some experience with the abdux project dealing with indexing of books doing "grey paper search".

RH continued as people were interested what he had actually built. A rather technical discussion followed:

each index has a dSI, a dataSetIdentifier. Within the dSI there is a baseURI with a URL referring to a Whois++/Ph/LDAP server.

dSI= Index for Sweden

baseURI= ldap://ldap.umea.se/c=SE

The index is stored in a SLAPD server. Underneath the country node there is one entry called dSI=Index for Sweden. Under this node there is a complete one-level list with all indices. A typical index would be:

idx = "Roland Hedberg+user+name"

dSI = "LDAP"

dSI = "Whois++"

baseURI = "ldap://some.server/some/base/reference"

description = language:swedish

Each index like this is one entry.

The problems so far is the need for a common schema for the different services and people want to modify entries. Furthermore, he wants to add a weight to the idx, e.g. this server/index has 29 John Smiths. Another problem is the updating, adding and modifying is relatively straightforward but for deletions the index has to be regenerated from scratch. Indices are generated by RH automatically or send via mail (e.g. for a Ph date base).

As LDAP and Whois++ are not yet supported as URI (read browser protocols) RH runs a proxy to all different directory service types.

The open discussion.

Adding an objectClass to the index for the 'name + user' would be better than overloading the name. Another issue is that there is only one index for Sweden which could become very big. Having one index is not really a truly distributed solution. Loose queries can generate (too) many replies. Longer names are preferred over the "tokenised" names. Longer names produce less incorrect hits and hence reduces the search space. A perceived problem is replication, there is the need for some kind of agreement for public replication.

For sub tree searches put a strong hint in dSI and add OID for type of service. Character sets are important but not yet, how would one map soundex characters between multiple languages.

It seems that LDAP is going to be the common protocol to use and gateways between LDAP and others will be supported. For the generation of indices LDAP can be used too, similar to the Centipede approach. The generation of indices will be done bottom up.

The indices can be made for people, organisations and OUs. To connect them (countries) together some kind of root is needed.

What next?

VB: Contact UKERNA to see if they are interested in participation.

VB: Distribute URLs for documents (<http://www.dante.net/pubs/dip.html>)

RH: Work on Internet Draft for indexing and this will be used for further discussion in this group.

PJ. needs a search engine as aprt of the DESIRE project and will invite Roland to work on this (?)

VB: set up mailing list for this group?

Thanks to University of Brunel and Andrew in particular for hosting the meeting.

Notes of DANTE meeting to discuss Managing the Root Naming Context. 18 June 1996.

Present

Vincent Berkhout (chair)	DANTE
David Chadwick	JTM Consultancy
John Farrel	ISODE Consortium
Andrew Palk	Digital
Keith Richardson	ICL

Introduction

Vincent started the meeting by saying that currently 770 DSAs copy the root context information, and so an automated solution based on 1993 protocols is essential to the Internet community.

The meeting used David's report of the Geneva X.500 ISO/ITU-T Meeting [1] as a framework for the discussions. This document listed four tasks that needed to be completed in order for a technical solution to be finalized. These were:

- I) Add access control information to subordinate references in the DISP to allow the List operation to work securely in a shadow DSA;
- II) Enhance the DISP to allow one level Search operations to work in a shadow DSA;
- III) Add clarifying text to the 1993 and 1997 editions of the Standard to describe how HOBs carry subordinate context prefix information to the superior DSA, and to state how this information should be held in operational attributes;
- IV) Enhance the EWOS Shadowing profiles to allow a single entry copy to be made.

After an initial technical discussion of all of the issues, and with the suppliers saying how easy or difficult and how long it would be for them to incorporate all of the above changes into their products, it was agreed that a two track approach would be pragmatically the best way forward. Track One (the fast track), would support knowledge distribution of the root context, but only the LIST operation of the root's subordinates. This should be available before the end of 1996. Track Two (the slower track) will in addition support one level Searches of the root. Manufacturers did not realistically expect this to be in their products much sooner than two years from now, given the long development cycles and that the changes would only be in the standard (at the best) before the end of this year.

The meeting then went on to discuss each of the four items in more detail, and to design technical solutions where appropriate.

I) ACI for List

The meeting agreed that a defect report should be submitted ASAP, and David agreed to do this based on the output from the group. A draft of this is appended as Appendix 1.

II) Enhance DISP for 1 Level Searches

It was agreed that we need a new type of shadowing agreement to signal that additional information is now to be shadowed (subordinate entries rather than subordinate references), but we do not need a new version of the DISP protocol (the ASN.1 for tree walking will remain the same). The new type of agreement will be set up by DOP, and shadowing will be rejected at this stage if it is not supported. In implementations where the DOP is not supported, humans will have to configure up their DSAs to support the new type of agreement (or not, as the case may be).

The meeting discussed the details of the technical changes to the shadowing agreement, and David agreed to write them up as a defect report. This is appended as Appendix 2.

III) Clarifying text for HOBs

The meeting agreed that this text was low priority, as it was believed that hardly any supplier supported the DOP as yet. No further technical work was done in this area. David agreed to take a low priority action to write this up.

IV) New Shadowing Profile

The "context prefix only" profile is needed for 1 level Searches, for copying country entries up to the root DSA. For root knowledge collection, all that is needed is for the subordinate reference of the first level DSA to be configured into root DSA. A telephone call is enough to achieve this. For List operations, ACI information is also needed by the root DSA, but it was generally thought that country entries would be publicly readable, and so in practice no access control information would be associated with them (other than "public"). So this new profile is only needed to be implemented in the longer term.

EWOS already have a more complex profile (EWOS profile ADY53, subset C ??) that allows a naming context to be chopped into arbitrary pieces, and so the context prefix entry is clearly a subset of the EWOS profile. Therefore it was argued, the EWOS profile would be appropriate for shadowing a single entry using the chop specification. The counter argument was that suppliers might be very slow to implement the full complexity of this profile, therefore a new much simpler profile is needed now. However Keith, who is the editor of the EWOS profiles, stated that the profiles were quite stable now, and so he was reluctant to add a new one at this late stage. If suppliers did turn out to be

slow to implement this profile then country managers would have to manually (or otherwise) configure their country entry information into the root DSA. The meeting decided not to ask EWOS to add an additional profile to its document.

Note. David was not too happy with this decision, and so has written an Internet-Draft that lists the shadowing profiles that are important to the Internet X.500 community. This draft document is appended as Appendix 3, and contains the new context prefix profile.

Whilst both country and root DSAs need to enter into a shadowing agreement to copy the root naming context downwards, existing implementations already support the profile (full naming context) that is appropriate for this. Minor modifications may be needed to deal with shadowing the root context comprising as it does of only a single root entry, but the suppliers did not think that this would be a big problem for any of them to support in the short term.

OneLevelSearch vs. Indexes

The meeting spend some time discussing the relative merits of 1 level searches and indexing, given that 1 level searches were on the slower track. One level Search can be fixed up in the standard in a reasonable time-scale (hopefully by the end of this year). However, it was agreed that indexes are not appropriate for ISO standardisation at this point in time due to insufficient operational experience of them.

It was agreed that whilst 1 level Searches provided significant operational benefits, they were not the complete solution. Indexes could prove to be more beneficial. If it came down to a choice between implementing indexes or 1 level searches the group would choose indexes. However it is seen that to get indexes to work could be more difficult than getting one level searches to work, since the latter have been working successfully in Quipu implementations for many years. The group thought that early operational experience of using indexes would be beneficial.

References

[1] Chadwick, D.W. "Managing the Root Context - Report on Geneva X.500 ISO/ITU-T Meeting, April 1996". Circulated to the IDS list at the end of April.

Appendix 1. Defect Report for Adding full ACIs to DISP for Subordinate References, so that List Operation can be performed in Shadow DSAs

DEFECT REPORT FORM

1. Defect Report Number:

2. Source: U.K. (BSI)/Internet ?

3. Addressed to: ISO/IEC/JTC1/SC21/WG8 and ITU-T Study Group VII Editor Group on the Directory

4. (a) WG Secretariat:

(b) ITU-T WP: WP4

5. Date Circulated by WG Secretariat:

6. Deadline for Response from Editor:

7. Defect Report Concerning:

(number and title of IS or DIS final text/ITU-T Recommendation)

X.525/ ISO 9594-9 (1993) The Directory - Replication

8. Qualifier: (e.g: error, omission, clarification required)

Omission

9. References in Document: (e.g.: page, clause/section, figure, and/or table numbers)

7.2.2.3 and 9.2.4.1

10. Nature of Defect: (complete, concise explanation of the perceived problem)

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.

11. Solution Proposed by the Source: (optional)

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."

Insert a new third paragraph:

"If subordinate knowledge (SDSE of type **subr**) points to an administrative point, then the SDSE shall be additionally of type **admPoint** and the **administrativeRole** attribute shall be present. If the **administrativeRole** attribute is present, then any subentries containing **PrescriptiveACI** shall be supplied, but other subentries need not be.

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

III) 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 subordinate knowledge is the access control information which governs access to the RDN of the subordinate knowledge. When the subordinate entry is an administrative point in another DSA, then part of this access control information may be held in **prescriptiveACI** subentries beneath the subordinate knowledge."

IV) 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."

12. Editor's Response:

(any material proposed for processing as an erratum to, an amendment to, or a commentary on the IS or DIS final text/CCITT Recommendation or Draft Recommendation is attached separately to this completed report).

Appendix 2. Defect Report Proposing an Enhancement to the Shadowing Agreement in order to support 1 Level Searches in Shadow DSAs.

DEFECT REPORT FORM

1. Defect Report Number:

2. Source: U.K. (BSI)/ Internet ?

3. Addressed to: ISO/IEC/JTC1/SC21/WG8 and ITU-T Study Group VII Editor Group on the Directory

4. (a) WG Secretariat:

(b) ITU-T WP: WP4

5. Date Circulated by WG Secretariat:

6. Deadline for Response from Editor:

7. Defect Report Concerning:

(number and title of IS or DIS final text/ITU-T Recommendation)

X.525/ ISO 9594-9 (1997) The Directory - Replication (Final text)

8. Qualifier: (e.g: error, omission, clarification required)

Omission

9. References in Document: (e.g.: page, clause/section, figure, and/or table numbers)

7.2.2.3 and 9.2

10. Nature of Defect: (complete, concise explanation of the perceived problem)

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 master DSA. This defect report proposes the additional text for Part 9.

11. Solution Proposed by the Source: (optional)

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**. In addition, if a subordinate entry is an administrative point, then the SDSE shall also be of type **admPoint**, and the **administrativeRole** attribute shall be present. All appropriate subentries below the administrative point shall also be supplied. The subordinate entries will contain attributes according to the attribute selection.

12. Editor's Response:

(any material proposed for processing as an erratum to, an amendment to, or a commentary on the IS or DIS final text/CCITT Recommendation or Draft Recommendation is attached separately to this completed report).

Appendix 3. New "Shadowing Profiles" Internet Draft

IDS Working Group David Chadwick

Internet-Draft University of Salford

DANTE IN PRINT n June 28 1996

draft- ietf- ids- x500- shadprof- 00.txt Expires: Dec 28 1996

X.500 Shadowing Profiles

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 1id- 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

This document describes several shadowing profiles for X.525|ISO/IEC 9594-9 The Directory: Replication. These profiles are needed in order for the Internet operational X.500 service to migrate from Quipu replication to ISO standard replication.

1.Introduction

The Internet Operational X.500 Service is currently based on Quipu replication [RFC 1276]. It is the intention to move to products conforming to the 1993 edition of the X.500 standard [X.500 93] as soon as practicable. Whilst it is recognised that the 1993 shadowing protocol is very comprehensive and fulfils all the operational requirements of the Internet X.500 directory service, it is also recognised that manufacturers are only gradually implementing the full complexity of the protocol in their products. Profiling is a recognised method of determining which parts of a standard to implement first and which parts to leave until later. The profiles specified in this document are a guide to X.500 product suppliers to indicate which parts of the shadowing protocol are most important to the Internet community.

Shadowing profiles are being produced by EWOS [EWOS], that should eventually be published as ISPs. The relevant EWOS profiles are referenced in this document, and are not reproduced here.

2 The Profiles

2.1 Full Naming Context

The unit of replication is the full naming context, with no attribute selection and full subordinate knowledge. This is equivalent to EWOS profile A. (check!)

Rationale. This is the most useful unit of replication in a public directory service.

2.2 Context Prefix Entry

The unit of replication consists of the single context prefix entry and all its subentries. (Q. Can we relax the "all subentries" requirement.) There is no attribute selection and no knowledge selection.

Rationale. This is needed for performing 1 level Searches in the superior DSA e.g. when managing the root naming context.

Note. There is no EWOS profile exactly the same as this, but profile F (check!) in its simplest form can produce the same unit of replication. However, profile F is much more general and consequently more complex to implement. It may thus take much longer for it to appear in products.

2.3 Single Entry (Spot Shadowing)

The unit of replication is a single entry from anywhere within a naming context. There is no attribute selection and no knowledge selection. Relevant administrative point information and subentries shall also be supplied.

Rationale. This is needed to maintain backwards compatibility with Quipu replication, and may still be useful for example, for shadowing aliased objects.

Q. Does the rationale still hold true?

Note. This is a more general case of profile 2.2, but is still much simpler than EWOS profile F. (check!)

2.4 Refined Naming Context

The Unit of Replication is the full naming context, but subtree refinement based on the selection of one or more object classes may be performed.

Rationale. Subtree refinement allows the consumer to select specific object classes e.g. organisationalPerson or index entries. In this way unwanted entries can be filtered out.

2.5 Other Profiles

Need to add a paragraph describing the other EWOS profiles, and saying that with these, all the remaining needs of the Internet community will be satisfied.

3 Security Considerations

Security considerations are not discussed in this memo.

4 Acknowledgements

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

5 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

[EWOS] EWOS Replication Profiles ADY53

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

Managing the Root Context

Report on Geneva X.500 ISO/ITU-T Meeting, April 1996

Although Managing the Root Context was not officially on the agenda for this meeting, the relatively large gathering of experts from many countries did give me the opportunity to unofficially discuss the problems that the Internet community are facing in this area. I took along the latest copy of the Internet-Draft [Managing the Root Context](#) for the delegates to read.

The response was generally very sympathetic and encouraging. The ISO/ITU-T working group is becoming much more pragmatic in its approach to X.500, and if there is a real problem to solve then they are generally willing to take a look at it.

Having discussed the root context problem with the ISO rapporteur (US), the ITU-T rapporteur (Australia), the Canada rep, the German rep and the UK head of delegation, the following strategy was agreed upon as a sensible way forward. The result should be the incorporation of sufficient functionality into the 1997 edition of the X.500 standard to make the Internet RFC for Managing the Root Context little more than a profile of how to use the standard features at the root level of the DIT.

There are a number of separate pieces of functionality that need to be agreed upon and incorporated into the X.500 Standard.

1. Add access control information to subordinate references in the DISP to allow the List operation to work in a shadowed DSA. (The ACI will need to be in both the subr DSE and its subentry.) This can be approached in one of two ways, viz: by either submitting a defect report on the 93 standard, or by submitting a ballot comment on the Operational Security addendum to the 1997 standard.
2. Enhance the DISP to allow one level Search operations to work in the shadowed DSA. This can be achieved by submitting a defect report on the 1997 edition of the standard, and if the defect can be agreed at the Nov/Dec 1996 ISO editing meeting, then the text can be incorporated into the final text of the 1997 edition of the standard before the text is sent for publication. Version 2 DISP, which will support enhanced 1-level Searches in the shadowing DSAs, will need to be negotiated between the supplier and consumer DSAs, and we have to decide how best to do this.
3. Add clarifying text to the 1993 and 1997 editions of the Standard to describe how HOBs carry subordinate context prefix information to the superior DSA, and to state how this information should be held in operational attributes. This will need to be made as defect reports on both standards (for the List operation in the case of the 93 standard and the Search operation in the case of the 97 standard.)
4. A separate work item that needs to be undertaken, is to enhance the EWOS Shadowing profile

to allow a single entry copy to be made. (This is needed when DISP replaces HOBs for carrying first level information to the root DSA.) At the moment the profile does not allow this. I have already contacted the EWOS editor to see how this might be achieved.

It has been suggested that we convene a 1-day meeting, probably to be hosted by DANTE in the UK in late May or June, to work out technical solutions and text for all of the above. We should invite all X.500 vendors/implementors/interested parties to this meeting, so that an amicable technical solution can be agreed upon. The solution can then be presented to the IETF in June. Once the IETF has agreed upon it, it can be presented to the ISO in the fall.

David Chadwick is a consultant for the [NameFLOW-Paradise](#) international directory services

22 April 1996

SCHEDULE 1 : Technical Annex

The service specification is split into the following parts, namely (1) operational services, (2) future development, and (3) other services.

1 Operational Services

The following are the core services needed to maintain the status quo, and needed to provide any kind of operational pan-European and Global directory service.

1a Operation of a Central Directory Services Agent

As a result of removing tasks from the central DSA the operation of this DSA becomes less critical. The central DSA will be run on a best effort basis and will be regularly backed up. The central DSA will be primarily accessible via Internet. X.25 access is provided only for a limited number of national networks with special requirements; this number will only be increased if specific requests are made and this service may be charged separately. Customers will be informed via the managers mailing list (see 1c) when an occasional scheduled downtime is anticipated. The operation of the central DSA is subcontracted to the University of London Computer Centre.

1b Operation of a Other Central Facilities

The well-used Central Directory User Agent is still provided at virtually no extra cost and will be maintained on a best effort basis. Availability and network connections are as for the DSA. In addition dial-up access is provided. As DANTE does not see that the operation of a central DUA being within the scope of expected services this service will be terminated before the end of 1995; the precise date will depend on an analysis of usage.

Organisations that have no means of managing directory information themselves will be referred to the national service provider. The national service provider will provide storage space and assistance on getting a location in the DIT. In exceptional cases DANTE might provide this service but it will be subject to charging.

Additionally, a list of public DUAs and LDAP servers may be maintained, on a best effort basis.

1c Mailing List Management

The following three mailing lists will be maintained:

<*Managers@nameflow.dante.net*> is an open list intended for national DSA managers for technical discussion on operational matters.

<*Forum@nameflow.dante.net*> is an open list intended for general directory services discussion.

<Customers@nameflow.dante.net> is intended for NameFLOW customers to discuss the service DANTE is providing. For the time being the old mailing lists will remain in operation. Other mailing lists will be provided if needed.

1d Information Server

DANTE will operate an information server containing documents relevant to the offered service. Gopher, FTP and e-mail access will be provided. A World Wide Web server <<http://www.dante.net/nameflow.html>>, will be set up on a best effort basis.

1e Help Desk

DANTE will provide help desk facilities for directory services during normal English office hours. The help desk will have two functions: primarily to assist national directory managers in solving international problems and to provide support for national directory managers in joining the global service, and, secondarily, it will act as a last resort for end users. When a query from an end user is received he will normally be pointed to the appropriate national help desk.

Additionally, for more technical issues a list of references to existing national directory help desks will be maintained on a best effort basis.

1f Support for New Customers

DANTE will provide new customers, subject to separate charging, with support while joining the service. This support reduces the load new participants put on the established customers while they join in. A small amount of support to newcomers will be a part of the help desk activity (1e), while more extensive support such as on site software installation falls under this item. To enable "easy access" for new organisations a list of national contact persons will be maintained.

1g Forum for Directory Managers

Organisation of a forum for Directory service managers with at least one meeting for the period April 1995 to December 1995 to discuss operational issues. This is likely to be organised by splitting the meeting into two parts: one part open for all interested parties, and the other part restricted to NameFLOW-Paradise customers.

1h Management of Pan-European Node

As an interim service DANTE will continue to manage the pan-European node, also known as "locality = Europe", of the DIT. Responsibility of this locality and subordinate entries will be transferred to a body that can reasonably exercise authority over locality Europe if such a body shows up and is willing to take the responsibility.

2 Future Development

DANTE will work in the following areas related to the future of the directory service:

2a Transition to X.500(93)

One of the reasons for the current status quo with respect to development of the service is the fact that the infrastructure is mainly based on Quipu software. A step to a more open and standard infrastructure will occur when the infrastructure is based on the 1993 edition of the X.500 standard this will allow non-Quipu implementations to be incorporated.

The study of the 1993 migration will be translated into a transition plan. This transition plan will have two phases: the first phase involves opening of the root context to allow non-Quipu implementations to act as a First Level DSA, and the second phase will replace the existing root DSA with a 1993 conformant DSA.

The first phase concerns opening up the root context which consists of the root DSA and First Level DSAs to enable non-Quipu implementations to connect. As an operational aspect of the second phase DANTE will investigate the possibility of replacing the root DSA with a knowledge reference server on the top level, a so-called DISP engine (DISP = Directory Information Shadowing Protocol). This DISP engine service will be made available to DANTE's customers. The migration to a X.500(93) system should start in 1995. If funds are available and customers are willing to participate a small scale pilot project will be initiated to test the 1993 edition of the standard. Depending on the outcome of the first pilot a larger pilot will be initiated which will be the basis for a complete 1993 based directory service which is planned for 1996.

Both phases will address replication issues (esp. shadowing first level DSAs), and if necessary relaying between different protocol stacks (in particular TCP/IP and X.25).

2b Index Servers

Depending on resources available, DANTE will continue the work on index servers. The minimal outcome of this will be a study report. This report could result in an implementation, depending on external funding.

2c Non-X.500 Directories

As X.500 is not the only directory service, DANTE will examine opportunities for coexisting and inter working with other solutions, such as WHOIS++. This will include following the development work done at IETF.

2d Locality Europe

It is unclear who should manage locality Europe, and who is allowed to register under it. The management of the node locality = Europe will depend on the outcome of the EEMA TOPOL-project. If the solution proposed by EEMA is acceptable to DANTE, DANTE will conform to this solution.

2e Service Level Agreements

One of the goals of DANTE is to improve the quality of the directory service offered. As part of a directory management framework Service Level Agreements between participants are needed as the reference for the required quality of service. One of the deficiencies is that currently the measured quality level is based on DSA availability, being only a part of the total directory system. Adequate management of the directory service requires monitoring and controlling the quality of the total system. At this moment (May 1995) the requirements for monitoring the quality of the service are defined. Measuring the system against these requirements requires tools. DANTE will produce an initial software specification for development of these QoS-tools and consider further development.

3 Other Services

Other activities, such as liaison with other directory pilots, are presented here.

3a International Co-ordination and Liaisons

DANTE will co-operate regarding directory and white pages services and development with relevant international bodies, including IETF and EEMA. DANTE will, subject to the agreement of the NADF, participate (on behalf of the European research network community) in the activities of NADF. Contacts with Eurescom will be maintained and, if possible, interconnection issues will be addressed with them. Co-operation with TERENA WG-NAP is considered a part of service development. International co-ordination is obtained through attending meetings and working in various bodies, whenever appropriate.

3b Promotion and Publicity

Depending on the availability of sponsors and the quality of the annual report (94/95), a glossy version of the annual report may be produced.

Other promotional activities will be targeted to the research community. DANTE will attempt to maximise the impact of directory promotion by co-ordinating activities with other organisations, for example EEMA.

3c Reporting

DANTE will provide quarterly reports dealing with service performance and provide an overview of the service and its development.

3d Interworking Problems

A log of all known interworking problems will be maintained on a best effort basis, and will be made available to customers.

Move to [[DANTE](#)][FLOW Services](#)[[NameFLOW](#)]

Contact: Konstantin.Chuguev@dante.org.uk

NameFLOW open meeting 20 January 2000

London

Invitation

DANTE intends to hold an open meeting on directory services and particularly [NameFLOW](#).

The following topics for discussion are proposed:

1. Country file (update from national initiatives).
2. Current NameFLOW state (DANTE, national managers).
3. DESIRE project:
 - o LDAPv3 client for Windows (Brunel University);
 - o LDAPv3 index server (SURFnet).
4. X.500 '93 migration (DANTE, Brunel University).
5. LDAP interoperability (SURFnet, DANTE, Brunel University).
 - o DIRECT Project (SURFnet)
6. AOB.

DANTE would like to ask you for participation. Other directory related topics are welcome. Online registration form is [here](#).



Time:
10am - 5pm

Location:
Room M128, Maths Building, Brunel University
Uxbridge Campus.
[\[How to get\]](#) [\[Uxbridge map \(PDF\)\]](#)

Lunch:
Local pub

Registration deadline:
January 7.

List of participants

Name	Organisation	Remarks
Henny Bekker	SURFnet ExpertiseCentrum	
Vincent Berknout	DANTE	
Konstantin Chuguev	DANTE	
Andrew Findlay	Brunel University	
Peter Gietz	DFN	
Bertold Kolics	MTA-SzTAKI	
Peter Lazzari	European Commission, JRC	
Thomas Lenggenhager	SWITCH	
Anders Lund	UNINETT	
Damanjit Mahl	Brunel University	
Luuk Oostenbrink	SURFnet	
Markus Sauer	DFN	
Rodney Tillotson	UKERNA	
Peter Valkenburg	SURFnet	
Jean-Marc Verbergt	BELNET	
Benjamin Zwitternig	ARNES	Early departure

Last updated: Tuesday, 18-Jan-2000 15:51:10 GMT

Move to [[DANTE](#)][FLOW Services](#)[[NameFLOW](#)]

Contact: Konstantin.Chuguev@dante.org.uk

NameFLOW January 2000 meeting

London 20 January 2000

Agenda

- 10:00 Introductions (VB)
- 10:15 Update from national initiatives
 - DFN, current state (MS)
 - DFN, future plans (PG)
 - SWITCH (TL)
 - BELNET (JV)
 - Short reports from other NRNs
- 11:15 Coffee break
- 11:30 NameFLOW Status Report
 - Agreement with MessagingDirect (VB)
 - RootDSA clean-up and '93 transfer (KC)
 - '93 FLDSA Replication (DM)
- 12:30 Lunch
- 13:30 LDAP interoperability
 - DIRECT project, SURFnet (LO)
 - DANTE activities (KC)
- 14:15 DESIRE Project
 - LDAPv3 client for Windows (DM)
 - LDAPv3 index server (PV)
- 15:00 Coffee break
- 15:15 Discussions
- 16:30 Any other business
- 17:00 Close

Abbreviations:

AF Andrew Findlay

AL Anders Lund

BK Bertold Kolics

BZ Benjamin Zwitternig

DM Damanjit Mahl

HB Henny Bekker

JV Jean Marc Verbergt

KC Konstantin Chuguev

LO Luuk Oostenbrink

MS Markus Sauer

PG Peter Gietz

PL Peter Lazzari
PV Peter Valkenburg
RT Rodney Tillotson
TL Thomas Lengennhager
VB Vincent Berkhout

Move to [[DANTE](#)][FLOW Services](#)[[NameFLOW](#)]

Contact: Konstantin.Chuguev@dante.org.uk

NameFLOW meeting January 1999

Utrecht 18 January 1999

Invitation

DANTE about to define new NameFLOW services that will replace the current X.500 Quipu based infrastructure. Secondly DANTE has taken over the maintainance of the current Root DSA service.

We intend to hold a meeting at the [SURFnet Office, Utrecht](#) on the 18th of January 1999 to discuss these and other ongoing directory activities in Europe. ([How to get there](#)) The venue is direct opposite to the central station of Utrecht. From Amsterdam, Schiphol airport it takes 40 minutes to arrive there. You can get timetable information of the trains at <http://www.ns.nl/reisinfo/>

There will be a "get together" on the 17th of January at 20h00 in the Djakarta, an Indonesian restaurant opposite the City Theatre. The address is:

Indonesisch Restaurant Djakarta
Lucasbolwerk 19
3512EH Utrecht
phone +31 (0)30-2319101

The participants hotel is:

Smits Hotel
Vredenburg 14
3511 BA UTRECHT
T: +31 30 2331232
F: +31 30 2328451

This invitation is primarily aimed at the representatives of DANTE's NameFLOW customers, although other interested parties may be allowed to attend if considered appropriate by DANTE and/or it's customers.

Last updated: Tuesday, 12-Jan-1999 17:34:00 GMT

Move to [[DANTE|FLOW Services|NameFLOW](#)]

Contact: Peter.Gietz@dante.org.uk

NameFLOW January 1999 meeting

Utrecht 18 January 1999

Agenda

- 9:00 Introductions (VB)
- 9:15 NameFLOW Status Report (PG)
 - Usage statistics
 - RootDSA transfer
 - PGP directory initiative
 - Hybrid solution
 - LDAP index work in DESIRE II
- 10:15 Report on the start of TISDAG (RH)
- 11:00 Coffee break
- 11:15 ITU work for X.500 (2000) (DC)
- 11:45 Results of the Dirconnect 3 at IDC (CB)
- 12:30 - 13:30 Lunch
- 13:30 Status report on Directory implementation work from those present
 - Belnet (JMV)
 - DFN (Project AMBIX) (RS)
 - SURFnet (TV)
 - SUNET (RH)
 - SWITCH (TL)
 - UKERNA (RT)
 - GRnet/GNDS (DP)
- 14:30 Coffee break
- 14:45 Report on IETF 43 (TL)
- 15:15 TWEB (Tuebinger Web to LDAP gateway) (KS)
- 15:30 Assignment of Actions, date of next meeting
- 15:45 Any other business
- 16:00 NameFLOW customers internal meeting
- 17:00 Close

Abbreviations:

CB Clive Betteridge

DC David Chadwick

DP Dina Papayannaki

HV Hans de Vries

JV Jean Marc Verbergt

KS Kurt Spanier

MK Michalis Konstantopoulos

PG Peter Gietz
RH Roland Hedberg
RS Ralf Schneider
RT Rodney Tillotson
TL Thomas Lengennhager
TV Ton Verschuren
VB Vincent Berkhout

Move to [[DANTE](#)|[FLOW Services](#)|[NameFLOW](#)]

Contact: Konstantin.Chuguev@dante.org.uk

NameFLOW-Paradise Directory Managers Meeting

BRUSSELS - 11 JUNE 1996

AGENDA

Provisional

Proposed Items and Presentations:

- Country File
An update/introduction by the attendees.
 - The Dan Net Success Story by Jens Ramsboel
The presentation will deal with:
 - the political setup in Denmark,
 - the technical setup by Dan Net - X.500 and interfaces, search and update facilities
 - the traffic, use and types of inclusions,
 - what did/do we do to promote the directory,
 - future aspects
 - SOLO by Ascan Woermann
Ascan will expand on his presentation given at the IETF. A short discussion of the future of SOLO should follow, "Is SOLO in the nick of time or too late?"
 - Indexing (by David Chadwick/Roland Hedberg/Nils Meulemans)
David just finished an improved version of an indexing strategy for X.500 and LDAP. Roland did a first experiment using SLAPD. Nils Meulemans has some experience using the WHOis+ + centroid approach.
 - Piloting X.500(93) and the Root Context by Vincent
In February X.500(93) was tested by several organisations. A short results paper will be available, hopefully the results of the expert committee how to proceed with the internet-draft and we will discuss whether people want to participate in a follow-up.
 - T.61 Naming in Poland by Tomasz Wolniewicz
n RfC 1617 it is proposed to use US-ASCII for DN's and it is felt that this restriction is no longer desirable. Tomasz will give a presentation on their approach, including:
 - a) special local attributes scheme duplicating the standard attributes but in correct spelling
 - b) DUA for displaying the stuff with discussion of the efficiency problems in name resolution. The "real" working version of the DUA is the WWW gateway based on web500gw.
 - MAITS
?
-

EEMA Annual Conference

The NP Directory Managers meeting is co-located with the EEMA Annual conference in Brussels (for detailed information <http://www.eema.org/eemahq>) from 10 to 14 June 1996. There are at least three interesting topics at the EEMA conference, next to numerous interesting presentations:

1. a) Open discussion of Top Level Naming in Europe
2. b) The EEMA Directory Committee
3. c) Exhibition

a) TOPOL - Free to all interested parties

Monday 10 June 1996 14h30 - 17h00

The latest news on TOPOL and please attend if you are interested: "It is an opportunity to draw in a larger group into the debate and to kick off the iterative, consultative second phase of the project which is intended to provide recommendations. Presentations should not take up more than half the session in order to give ample time for discussion. The presentations should encompass the following:

- a review of the TOPOL I document;
- consequences of recent LDAP, DNS and other Internet initiatives announcements;
- the issues around the use of c=WW (SITA to be invited);
- consequences of the recent NADF revival;
- strawman on recommendations for top level and several European countries:
- WW
- UK (DWC to be invited)
- France
- Switzerland (Thomas Lenggenhager to be invited)

EEMA Directory Committee - Free to all interested parties

Thursday 13 June 1996, 12h30 - 16h00

The Exhibition - Free to all interested parties

Tuesday 11 June 10-18h

Wednesday 12 June 10-18h

Thursday 13 June 10-16h

and from the promo-flyer: The extensive exhibition incorporates demonstrations of the latest messaging an electronic commerce technology including:

- o X.500 Directory applications
- o X400 88 Challenge - reliable File Transfer
- o Computer Fax technology

- o Free consultancy
- o Internet Cafe

If your interested in the rest of the Annual Conference, there is an awareness seminar on Monday 10 June morning and a conference session wednesday 12 June 9h00 to 10h45 on "Directories - showing the way forward". Note that there are registration fees for both the morning seminar and conference session.

If you have any questions:
nameflow@dante.org.uk

Minutes of the NameFLOW meeting at ULCC

21 September 1995

Opening and Administration

Welcome and introduction of the group. The previous minutes were accepted and the agenda for this meeting was presented. Agenda for the meeting:

Thursday: 21 September 1995

10:15 Agenda and Administration

11:00 REDACS NAMMAN (Paddy Boocock, Unisys)

13:30 Country updates

14:00 Root Context (David Chadwick)

14:30 EurOSInet (Chris Ridd, NHS)

16:00 PGP Keys in the Directory. (Roland Hedberg)

17:00 PARTY: PGP signing party (Roland and Vincent)

Friday: 22 September 1995

10:00 Whois++ (Alan Emtage, Bunyip)

13:00 Customer meeting / NP-93 pilot

REDACS NAMMAN

PB explained the REDACS NAMMAN (RN) project. RN is a X.500 based directory service which provides a management information base for the Royal Navy. It is used to store information on users, profiles and systems. The system performs access management based on this data. The system is used by 5-7000 users (total number entries up to 100.000) and operated by 8 people.

(Lunch)

Country Reports

Belgium

NM gave a short presentation on the LIRN project which uses X.500 to store information on libraries in BE, UK (GB) and PT.

LIRN holds a multilingual thesaurus, a list of resources (library name, collections etc.) and other Information (contacts, opening hours etc.). LIRN can be found at:

http://sun7.iihe.ac.be/index_uk.html (uk could be replaced with fr, pt or nl)

http://lirn.viscount.org.uk/index_uk.html

http://diana.ci.ua.pt/index_uk.html

The size of the Belgian directory is growing. The University of Brussels proposed a Yellow Pages service based on X.500 which could be the subject of an "already existing" X.400/X.500 project they have with Belgacom. They are now awaiting Belgacom's reaction.

Switzerland:

The size of the directory is stable, the administration of the federal government plans to connect with its Quipu based DSAs from SSE into the Swiss DIT by the end of the year. There are still some legal issues to be solved in respect to the data protection law. Switch plans on testing the IC '93 release in the near future (29 September 1995).

Germany

The size of the directory is stable, the use of the public DUA is decreasing in favour of the Web-500 gateway. Birko introduced their new Web-500 gateway of Frank Richter. <Frank.Richter@hrz.tu-chemtiz.de>. There are also some legal issues which need to be resolved.

Netherlands

SURFnet introduces their new booklet "Introducing a Directory Service". The size of the directory is stable and there are also some legal issues which need to be resolved. SURFnet is testing NEXOR's X.500 '93 implementation (MDS).

Finland

Pekka presents the Directory services in Finland, at present there are 4 operators: 3 using X.500 and 1 using Whois++. They see a strong requirement for X.500/Whois++ gateway and plan to work on this.

UK

The size of the directory is still growing, both for universities and companies. UKERNA reports less legal problems due to British legislation (which does not conform to European legislation).

Hungary

The size of the directory is stable. The government is starting X.500 services as well.

The Hungarian PTT is planning to set up an X.500 service.

ICE:

DC presented the "Infrastructure for Certification authorities in Europe" (ICE) project which involves setting up a certification infrastructure for verifying digital signatures. The project starts in November 1995 and will run for approx. 2 years. The project deals with the development of key and certification tools (applications) and services. The actual testing/implementation of the service should be done in the second phase; the TWICE project.

Root Naming Context:

DC presented his Internet Draft. The first experience learns that it would be better to allow DISP-DISP as well as DISP-DOP due to a lack of HOB implementations. This has the implication that the minimal root DSA would be running in a consumer and a producer role. DISP-DISP replication will be added to the ID and the ASN definition will need an update for the chop argument for retrieving the top level information without (excluding) overwriting its own entry. Ronan pointed out that we are rebuilding a 1993 infrastructure using the existing Quipu model as a basis.

EuroSInet: (Contents of sheets appended)

CRidd presents both the UK NHS Pilot and EuroSInet. EuroSInet works on demos and interconnection test of OSI products, however there is a move in focus to Systems Interconnection (not just OSI). Future co-operation has to be explored. A web page is available at <http://www.imc.exec.nhs.uk/eurosinet/>

Next meetings:

VB proposes to arrange the next NameFLOW meeting in close conjunction with the EEMA meeting (10 June 1996 in the "Pyramid", Brussels). Arrangements will be made to make this happen.

PGP Keys in the directory

RH started with an explanation of PGP, he has modified the German Web500 gateway to read the PGP data from the directory using a web browser. There are however some challenges with string termination (<nl><cr>)

Someone is going to write a proposal to set-up a PGP hierarchy which is similar to the DIT hierarchy. What is needed is a policy statement where we should go to set up a trusted infrastructure of people (or organisations) using the directory.

VB's PGP fingerprint

ID 1024/B348C7B1 20/9/95

E2 57 CB 6C 26 AD 5A CE

48 C0 F9 CF 96 59 96 E2

The idea is for VB to retrieve the keys and sign them and redistribute them, upon receipt they will be signed and returned by the other people.

22 September 1995

Whois++

ALM gave an introduction to WHOIS++. Alan is head (officially Vice President Technology) of the Whois++ development team at Bunyip (<ftp://ftp.bunyip.com> or <http://www.bunyip.com> holds more information). Currently there are about 70 servers. The first two RFC are published: RFCs 1834; Whois++ and Network Information Lookup Service and 1835; Whois++ architecture.

Discussed topics: Whois++ is a wrong name but this is due to historical reasons. It is based on a query and on an index protocol. It was acknowledged that there is the need for guidelines to prevent loops and to be able to do a "global search". The collapsing of the name space, (e.g. the entries John Smith and John Jones would collapse to three indexes: john, smith and Jones) is experienced in Sweden with 40% "collapse" on given names and 60% on last names. The advantage of Whois++ that it does not search the servers that not satisfy the conditions of the query, vs. X.500 where one has to traverse to a complete tree. There were two disadvantages at this moment: security (access control) on attribute level and scalability. These two topics are now being researched. There were questions whether it supports multiple languages (not yet solved).

Lunch

Closed Customer Meeting

(separate minutes, distributed via customers@nameflow.dante.net)

Attendee List

NAME , ORGANISATION <E-MAIL ADDRESS >

BERGT Birko, TU Chemnitz-Zwickau <bergt@hrz.tu-chemnitz.de>

BERKHOUT Vincent, DANTE <Vincent.Berkhout@dante.org.uk >

BOOCOCK Paddy, UNISYS <paddy@boocock.demon.co.uk>

CHADWICK David, JTM Consultancy <d.w.chadwick@iti.salford.ac.uk>

EMTAGE Alan, Bunyip <bajan@bunyip.com>

FLOOD Ronan , ULCC <R.Flood@noc.ulcc.ac.uk>

GAVINE Allan, DRA <allang@green.dra.hmg.gb>

HEDBERG Roland, SUNET <Roland.Hedberg@umdac.umu.se>

JARVELAI Pekka, FUNET <jarvelai@csc.fi>

JURG Peter, SURFnet <jurg@surfnet.nl>

KLUNDER Hans SURFnet <klunder@rc.tudelft.nl>

LENGGENHAGER Thomas, SWITCH <lenggenhager@switch.ch>

MEULEMANS Nils, Univ. of Brus. <meulemans@helios.iihe.rtt.be>

MORLEY Nicholas, DRA <nickm@green.dra.hmg.gb>

MOWITZ Enrico, TU Chemnitz-Zwickau <ds-manager@hrz.tu-chemnitz.de>

RICHARDSON Keith, ICL <k.richardson@man0523.wins.icl.co.uk>

RIDD Chris, NHS/EuroSInet, <C.Ridd@imc.exec.nhs.uk>

ROBBINS Colin, NEXOR <c.robbins@nexor.co.uk>

ROTTSCHAFER Marjo, SURFnet <rottschafer@surfnet.nl>

SEYMOUR John, ULCC <J.Seymour@noc.ulcc.ac.uk>

TETENYI Istvan, DANTE <Istvan.Tetenyi@iif.hu>

TILLOTSON Rodney, UKERNA, <R.Tillotson@ukerna.ac.uk>

Thanks.

I would like to thank Paddy, Nils, David, Birko, Chris, Roland and Alan for their presentations. Thanks to ULCC for hosting the meeting and providing lunches and other arrangements. A personal thanks to Hans for taking the minutes. -- Vinc&

8< - - - - - Sheets EuroSInet Presentation - - - - - >8

* NHS X.500 Pilot and EuroSInet (1)

by Chris Ridd, NHS IMC

* NHS X.500 Pilot (2)

Co-sponsored by IMC and NWND

Potential uses in NHS

Phase 1

* Doom and Gloom? (3)

Other sites dropped out

Only used one DSA

Huge data problems

* Yes, but... (4)

PARADISE connection

Several DUAs

Schema

Practical experience gained

* Phase 2 (5)

Starting January 1996

Aims

X.500 (1993)

* EuroSInet(6)

Founded in 1986

Initially demo'ing OSI in business scenarios

Now focused on testing

* Approach (7)

Test writing workshops

Interoperability Testing workshops

Broadening of scope

* Current state (8)

Current Members

- Alprange, Apple, AT&T GIS
- Audilog, Bull, Data Connection
- - Digital, EDS-Scicon, HP
- IBM, ICL, ISODE
- Marben, Microsoft, Net-Tel
- NEXOR, NHS IMC, Olivetti
- SNI, SITA, Tandem

* Test Suites (9)

Emphasis on User Requirements

- Test purposes
- Smaller suites - core functionality

* Test Workshops (10)

Collaborations with user groups

- EEMA

Pre-staging events to ensure successful demos

* Test Workshops (11)

Results are registered

Vendor neutral reports

* Sample FFS (12)

...[] stuff deleted []....

* X.500 Test Suite (13)

Currently 1988 based

3 test scenarios

Used at EEMA pre-staging workshop

* For Further Study (14)

NSSRs

First level DSA tests

1993 protocols in 1988 products

1993 tests

Dangling aliases

Looping

MHS use of Directory

NameFLOW Closed Customer Meeting - 22 September 1995

VB handed out the NameFLOW service specification, the Quarterly Report (April-June) and the Annual Report (1994-1995).

The central DUA service is going to be discontinued next year as it using a second machine. Over 2/3 of usage is coming from North America and another part from the UK. Both domains provide their own public DUA (DE) as an acceptable alternative.

1993 Pilot Experience

DANTE experienced some problems with MDS (NEXOR's 93 implementation) as conversion tools are unsupported and not fully functional. ULCC has problems with replication (chop?), the DISP part of MDS. SURFnet has problems with EDB conversion and replication in MDS;

SWITCH has not started with the IC '93 yet but promises to start testing as soon as the beta release is available (end September 1995).

VB suggested to perform a more focused test in a short "dedicated" time frame.

Testing of NP-93 will be done between 15 -19 January 1996, followed by a one-day meeting to discuss results/plans. The meeting will be held on Thursday 25 January 1996 at Schiphol Airport (alternative date suggested Wednesday 24 January 1996). The minutes state "arrangements by DANTE" but some co-operation will be much appreciated!

The Common Index Protocol is the way to go forward (indexing both Whois++ and X.500 directories) and it is suggested as agenda item for the next meeting.

DANTE will not co-ordinate Whois++ service in the short term (1995-1996), however will run a Whois++ server (as done by NEXOR). If however Whois++ servers are being deployed on a large scale, then DANTE will run a centroid, depending on the centroid policy available.

Monitoring Quality of Directory Services. The current paper of CJR will be used as input to the IETF IDS working group and DANTE will ask CJR to produce a first draft.

As the Annual report is being well appreciated DANTE will produce an Annual Report for 1995. Other more general remarks to a 1996 service specification.:

1e Help desk

Look into the possibility that queries from the UK are send to the wrong help desk?

2b Index Servers, and 2c Non-X.500 Directories

It was decided that points 2b (Index Servers) and 2c (Non-X.500 Directories) would be considered "points of interest".

3a Liaisons

Liaisons paragraph will be will use more general terms, without naming the actual organisations.

3d Interworking

Deleted as it is an integral part of the 93 transition.

The Distinguish Name used for directory managers vary at this moment because there is no proper guideline. There are two RFC's addressing this naming issue: 1803 and 1617. (Excerpts appended) Ronan has sent comments about this to the RFC editor when it was still in draft but proposed changes were not adopted .

A guideline for NP will be:

1. Each organisation has an entry "CN=Directory Manager".
2. The proposed generic e-mail address should be "directory-manager@foo.com".

VB will send a mail how to add organisational URL's to the directory. This topic must be addressed now as managers have the EDB file with the information local. Doing this after the 93 transition is almost impossible as **all** managers have to be contacted.

Next managers meeting is planned for the 10 June in co-operation with the EEMA in Brussels.

8< ----- RFC 1803 (begin excerpt) ----- >8

2b. Operational recommendations for the service provider

* Provide a generic e-mail address for the DSA manager (e.g., x500-manager@foo.com). More than one manager should be available to handle problems as they come up (i.e., the manager should be able to go on vacation!).

* E-mail to the manager of the master DSA must be answered in a timely fashion:

* All mail to the manager should be acknowledged as received

within one working day.

8< ----- RFC 1803 (end excerpt) ----- >8

8< ----- RFC 1617 (begin excerpt) ----- >8

2.4.1 Directory Manager, Postmaster & Secretary

Similar to messaging, where every domain has its postmaster address

it is highly recommended that each organisation in the X.500

Directory has two entries: Postmaster and Directory Manager. In

addition, Secretary entries for an organisation and its units should

be listed. If this guidance is followed, users will benefit because

it will be straightforward to find the right contacts for questions

or problems with the service.

These entries should use the object class `organizationalRole` with the

`roleOccupant` attributes containing the distinguished names of the

persons in charge of this role. The values

CN=Directory Manager

CN=Postmaster

CN=Secretary

should be added as additional values whenever another language than

English is used for the name of the entries.

8< ----- RFC 1617 (end excerpt) ----- >8

Minutes NameFLOW-Paradise Meeting

14 & 15 March 1995, University of Salford, Manchester, United Kingdom

(This file is also in [Post script format](#) available.)

First day (14 March 1995)

Administration (9:30)

The next meeting will be held on the 21 and 22 September (week 38) at a location near Cambridge. Depending on the number of persons representing DANTE and the availability of a host an alternative location outside the United Kingdom. might be chosen. At this moment there are three mailing lists:

- [<forum@nameflow.dante.net>](mailto:forum@nameflow.dante.net), for general directory discussions
- [<managers@nameflow.dante.net>](mailto:managers@nameflow.dante.net), for national managers discussing technical issues and operational aspects.
- [<customers@nameflow.dante.net>](mailto:customers@nameflow.dante.net), for the customers to discuss the NameFLOW-Paradise service DANTE is providing. The list is not yet fully operational. There were some operational problems with the mailing lists configuration, but they have been provisionally resolved.

The previous minutes were accepted, which are available at [<ftp://ftp.nameflow.dante.net/meetings/minutes.nov94>](ftp://ftp.nameflow.dante.net/meetings/minutes.nov94).

Quality of Service & Service Level Agreements (10:00)

Colin Robbins discussed the Quality of Service (QoS) paper. Paragraph 4.2 mentions the bias of the directory: passive probing versus user requirements. Active probing is done but what does it actually mean? (or, what do the figures say?) The actual thing that needs to be measured is the service to the end-user: the "stuff" people want to find. The most important perspective is the user perspective. The advantage of passive probing is that it is not limited to the first level Directory System Agents (DSAs), but says more about the service as a whole and provides a general quality indicator of the Directory. Passive probing results in a set of figures representing the availability of one "country" (and relating information) versus another set representing all other countries. The current measurements are based on a few selected Directory User Agent (DUAs) probing the same data set and the question is raised if this has any importance at all.

David Chadwick raised the point that with respect to the first requirement reliability and response time should be distinguished. A nice example was: a bicycle is very reliable and takes me there in 12 hours, British Rail is "unreliable", but take me there in 30 minutes (good response).

The second requirement also comprises two different issues: accuracy and completeness. (It's better to have only a good phone number in the directory than retrieving a bad phone number). Accuracy can be determined statistically e. g. take 100 entries and see how many are correct. However completeness is harder to determine. Completeness should be separated into two issues: missing entries and missing attributes (per entry).

The Service Level Agreements (SLAs) are not yet that relevant to the service, but are to be used as a measurement tool e.g. for DANTE to service providers (DSA up and running), and for service providers and their customers. The 93 model needs bilateral agreements which can be based on SLAs. Some parameters e.g. regarding completeness could be stored in the directory itself.

The current entries in the Directory Information Tree (DIT) vary widely in quality. There could be two approaches for operating the service.

1. A free for all, with few (if any) rules for joining: here is a pool of information of variable quality and use whatever you need. Tools will be needed in order to measure the quality of the different parts of the DIT, so that users know which part to use when.

2. An imposed QoS: to join the NameFLOW-Paradise you must guarantee to provide a given QoS (possibly in accordance with a SLA).

The conclusion was that if strict rules for QoS are imposed on customers there would be **no customer at all**, unless we pay them.

What would be the next step with the paper (goal of the paper?) Does DANTE provide the tools to measure? The answer is that DANTE has the prime objective of co-ordinating the service and not of providing tools. The community replied that there could be alternative ways to provide these tools. The service providers could provide the tools, if DANTE is willing to provide the guidelines.

It was decided that we would postpone the official distribution of the QoS paper and test it first. There is a need for software specification to build the tools (e.g. using PERL). These tools should be tested in a small organisation.

There is a ratio of one external communication to 60 internal communications (e-mail). The usage of the Directory could be different (more external use than regular e-mail)

People were concerned that the SLAs are going to be signed, filed and never looked upon again. When the SLA is signed by the right person, this will not happen.

David Chadwick suggested that SLAs could be used for commercial reasons: if an organisation is having an "x" level of directory service at "y" costs, another service provider could deliver the same "x" level cheaper. Now the provided service level is not measurable.

The overall conclusion was that the SLAs are basically in order and other operational agreements like shadowing etc. could be added. The Internet Draft "Recommendations for an X.500 Production Directory Service" by Russ Wright should be consulted for any additions to the QoS paper. We need to look at guidelines with respect to e.g. DSA manager. The RFC should be changed if it does not support our requirements.

Indexed DSAs (12:00)

Paul Barker gave a clear presentation of his paper. The question posed was how does a DUA know which indexed DSA to use. Paul will incorporate security aspects to the final version of the paper. It was noted that a similar mechanism is used by CAN (?). It was questioned whether Indexed DSAs scale.

Locality=Europe

Currently the Locality=Europe issue is part of a proposed EEMA project named TOPOL. DANTE is a member of a consortium bidding for the project. If the TOPOL project not awarded to us, we have to make certain that our voice is heard.

** Action: Ronan, Vincent, and Colin will see to it that 4 obsolete nodes at the Locality=Europe node will be removed (4 of the total of 8 nodes).*

More research on DUAs algorithms (12:30)

Paul Barker gave a presentation on his latest research on DUA algorithms optimisation. The presented sheets are distributed on the forum@nameflow.dante.net mailing list, however they are not to be distributed any further. It was noted that Paul did his research on one specific set of data (queries). Paul said that his research (performance tests for several algorithms) would be resubmitted to a second set of test data.

Liaisons

- **Eurescom might connect at the end the year.**
- **NADF, Vincent will attend the next NADF meeting.**

- **APIA is working on directory synchronisation. There is another group (?) working on the Directory management.**
- **EEMA is reviewing the directory group itself. The meeting was too technical for most of the attending members. A problem is the coming and going of people within the Directory Group results in a very small core group. In the future the meeting will be split into one "liaison" part for information and another part focusing on user issues. Discussed issues were: the Directory implementations report, the DSA service (containing information about EEMA itself and members), the vendor demonstration at the Amsterdam annual conference (June) and the review period for the TOPOL project.**

Defect reports

Defect reports were submitted to the BSI (on behalf of DANTE) by David Chadwick and Colin Robbins (NEXOR). The reports are submitted because there is lack of clarity and defects in the standard:

- **How to shadow first level DSAs (One function of the root DSA is distributing the root EDB. In addition it replicates (shadows) all country EDBs.)**
- **How to perform list and one level search operations at the higher part of the DIT.**
- **Root context is no longer mentioned in the 93 standard (is rejected because the standard is better without the definition)**
- **How to operate a root DSA which masters first level entries and has HOBs with first level DSAs. David Chadwick was asked to write an RFC which described this, as it was said to be outside the scope of the intentions of the base standard (even though the protocol supports it).**

Country report

(UK/GB by Rodney Tillotson representing UKERNA)

There have been no major changes. The UK DIT has about 250 000 entries. It is hard for a service provider to move from a pilot to a real service. One of the problems is the lack of real service documentation. The physical part of the Directory Service provided by UKERNA is a collection of boxes (3 DSAs) run at ULCC.

After the country report there were some complaints about the "buggy" DOS version of the DUA (de-DOS). This will remain "as is" for the time being, as no-one was willing to fund the improvements.

Vendor reports

There were 3 vendors present at the meeting Digital, ISODE Consortium and NEXOR. They were all given a 10 minute time slot to present technical information of their

X.500 product, where after questions could be asked. There were questions concerning management tools, accounting (e.g. X.25 usage) and interfacing to database systems. Further details can be obtained via Nick Emery, Steve Kille and Colin Robbins. (See list of attendees).

Patrick Fantou attended the meeting and was allowed to give a brief update on the current X.500 developments at Siemens Nixdorf. DANTE will look at the possibility to give them a vendor slot at the next meeting.

Root DSA

The first thing that was discussed is that we should provide this service for the (NameFLOW-Paradise) community represented by our customers, not for the complete world.

The current format of the root EDB is too QUIPU specific and needs an OID-table. There are already scripts to convert the EDB to a knowledge reference file. Colin will send the script to ULCC and they'll run the script over the root EDB and generate a "knowledge reference file". DEC, NEXOR and Siemens Nixdorf are willing to test the ftp-retrieval (for both the EDB and knowledge reference formats).

It is recommended to remove audio and photo-attributes from the FLDSAs and country entries. However, leaving the audio and photo attributes "as is" will not break the system (DSA). (The software will detect that the attributes are invalid) The conversion script should filter out the references to audio and photo files.

** Action: provide three mechanism for root EDB distribution: root EDB via ftp, root EDB via the QUIPU getEDB protocol and a knowledge reference "root" via FTP. QUIPU replication for FLDSAs should be turned off at December 31, 1995.*

** Action: Thomas encountered a problem with installing the ftp'ed copy of the root EDB at his FLDSA. The problem could be resolved by removing edbINFO=#gt# attribute. Thomas will check the proposed changes.*

Why ftp of EDBs? ("out of band replication") When the DSA is pulling the root EDB it is very slow, it even stops for a short period after a getEDB. It was suggested that further replication (other than just the root EDB) could be optimised by ftp-ing all first level EDB information (preferably in tar format).

** It is decided that the Giant Tortoise (G.T. is the root DSA) will no longer be available as relay DSA. This would have impact on all FLDSAs (First Level DSAs) having no Internet access. The responsibility for relaying should be turned around: a "customer" should be looking for an organisation (DSA?) that is willing to relay on request. At this moment there are approximately 20 DSA using only x.25.*

** Action: The proper way to handle this is to contact the managers of pure X.25 DSAs via the NameFLOW managers list and osi-ds list. The message should state that relaying at the root level will be turned off and pure X.25 DSAs can be reconnected on request.*

(Note: As a comment to the meetings minutes Colin Robbins said that turning relaying around might need software modifications)

** There is communication with Tim Howes of NADF concerning the c=US node of the Paradise pilot (and investigating future co-operation?).*

** The stage two of phase one will be discussed at the next (September) NameFLOW-Paradise meeting. UKERNA and DANTE will look at the possibility for an alternative*

backup of the root DSA.

* The Giant Tortoise (root DSA) is referenced as parent-dsa by most QUIPU DSAs as this is part of the default set-up. If a DSA can not find any suitable reference it will fall back to the parent-DSA (in most cases G.T.).

Action: As most FL managers did not reply to the previous message it should be sent to osi-ds list saying we are planning to remove the GT and that the use "parent-dsa = G.T." will not longer be valid.

* A problem not recognised previously is that country entries need to be changed by DAP. Changing the country entry is not possible via DSP: a DUA needs to contact the DSA directly because of authentication (security) reasons. Two solutions were proposed:

1. change country (c=xx) information by hand when needed (on request).
2. all country managers will be added to the DAP access list for the Giant Tortoise (quiputailor). This is an one time effort enabling country managers to change their country entry.

The second approach was preferred as it is a more automated solution.

Action: Tell the first level managers why, what and how it is being changed. If this is done successfully DAP access to the G.T. can be blocked (restricted) as suggested.

* When should we start a standard service with a root DSA that is available on a "best effort basis", rather than a guaranteed 7*24 hour. ("Best effort basis" means that "whenever the root DSA is down, it should be running within a reasonable time during office hours"). The transition to a standard service is proposed on the 1st of May 1995.

The following dates/actions were set:

ASAP: turn off DAP access.

1 May 1995: lower availability/ service level of root DSA.

1 May 1995: turn of relaying.

End of the year: end of QUIPU replication. (31 of December 1995)

Second Day (15 March 1995)

93 Migration (9:30)

All the current QUIPU versions, including XT-QUIPU 8.2, are based on the QUIPU model. Colin Robbins suggested two approaches for a 93 migration

1. The QUIPU-88-93 migration.
The first approach is to replace the top level (root DSA) with a DSA that supports

both standard 88 DSP and QuipuDSP. Subsequently the FLDSAs could then be replaced with similar "bi-model" DSAs. If all the FLDSAs are in place the first level could be based on a pure 88 system. The migration to a 93 system is easy when an 88 system is in place. The same approach is then used for the next level (country and organisation) approach and is to be repeated for the next level and so on.

2. The QUIPU-93 migration

The second approach is to skip the intermediate 88 DSA and replace the QUIPU DSAs with 93 DSAs.

The advantage of the first approach is that it takes more steps but they are less complex. The disadvantage of the first approach is that it will take two transitions. A disadvantage of the second approach is that it has "grey areas" which have to be explored. At this moment there is a QUIPU/88 implementation available and no 93/QUIPU implementation. The number of `88 implementations is not that big and software vendors providing an 88 system also provide a 93 system. It is expected that a wide range of 93 implementations will become available in the near future. It is best for those who do not have an 88 system to go directly to a 93 system.

The problem of the transition has to be tackled top down. The changing of all leave nodes and work to the top is unfeasible and will take too much time.

* What is the major difference between the 88 and 93 standard? In addition to the 88 standard, the 93 standard allows replication. The difference between QUIPU and the 88 model is that the information of country information is stored at a different level. As a consequence the 88 model supports LIST and does not support SEARCH operations.

The transition could be done using a shadowing strategy, that is running a QUIPU system and a `93 system in parallel. The drawback is that this will need additional resources, like memory and so, but no additional boxes as two DSAs can run on one box. A good point raised was that the 93 standard has some defects with lists and search operations (see defect reports). It is believed that these defects are real defects and these will be changed in the standard. This could take some time.

* Clear transition guidelines have to be prepared so that people know what we are planning to do.

* The root DSA will be replaced with a "*DISP engine*", this is an 93 conform DSA implementing only the DISP protocol. This could run on the same machine as the current root DSA as it will be a small process.

* What is the number of DSAs currently used? There are about 600 DSAs running now: 10% non-QUIPU, 20% based on IC software and about 12-14% based on NEXOR software. That leaves us with over 300 DSAs which are "not supported" versions of QUIPU.

* To gain experience with 93 software, it should be tested on an organisation level, on a country level and at the root level (the DISP engine).

Multiple Service Providers

Paul Barker suggest that most of the logic to resolve multiple services should be put

into DSAs rather than DUAs. Why the need for private DITs? Organisations want to maintain their own information. It is a user interface problem, whereby the user interface should select the appropriate DIT. It is a "private" issue: a certain application will require a certain DIT. If organisations want their information to be available for their customers, they must be willing to combine their efforts (like phone companies combining the phone numbers).

David Chadwick concluded that there are three distinct cases: none of the discussed examples was outside the scope of these three cases.

1. Unconnected DITs where the solution is completely solved by DUAs.
2. Different DNs for the same objects.
3. Same DNs for the same object

We might end up with a hybrid solution comprising all three cases.

Should there be one big "grand daddy" Directory tree? This "master" DIT could have naming links to information in other private DITs. The solution could be the other way around, where private DITs have links to the "master" DIT. The naming links can be replaced with attribute links. (One idea is to extend naming links to indicate which attribute the remote DSA holds) This needs DUA enhancement to follow links. (This is already supported by some applications, e.g. Eudora).

The suggestion is to shift the name link responsibility to the connecting networks, as we have the biggest DIT at the moment. This will result in a DIT with lots of pointers to other (private) DITs or vice versa.

Thomas Lenggenhager gave an example of the SWITCH approach, where DNs are registered via legislation. This is not going to be the case for all countries. Rodney Tillotson (UK) made an X.400 comparison with " " (space) option. The space option implies that ADMDs within the UK have agreed to route information to all other ADMDs (multi-lateral agreements).

It seemed that the unique DN is the preferred solution, with multiple private DITs. If people do not support the unique DN the information could not be retrieved.

Closed Meeting (Afternoon)

Feedback on Previous Service Specification. (14:00)

The previous service specification was evaluated and no substantial comments were made. The brochures of NameFLOW should be distributed at meetings and conferences. DANTE should be present at the next EEMA meeting. It was agreed that the quarterly reports for the period 5/94 - 4/95 will be replaced with an annual report. Instead of putting all daily/monthly statistics in the report it will contain a URL referencing to the ftp-server containing the complete statistics.

What about the low availability of FLDSAs? (Will this be solved with SLAs?) ULCC will take a look at them and see what can be done about it. (Note: ULCC has already taken action)

New Service Specification (14:30)

The new service specification should get rid of "free services" (DUA, LDAP, etc.) and X.25 connectivity. Back-up root DSA: what does it precisely do, do we still need it, for how long could the root DSA be down (max.), ...? With respect to the '93 migration it needs to mention that the root DSA will be replaced with a "DISP engine". Add a justification paragraph (why do we need to migrate?). Charging people for using DISP (i.e. only connect those who pay, unless they are peers, e.g. NADF). This has to be announced now, if we want to implement it 1/1/96. About SLAs: should we have SLAs in place between DANTE and customers? Should we make a specifications for software that would monitor the service? SLA should be used for a 93 standard service.

* DANTE should provide more services to customers only.

* To improve the visibility of DANTE it was suggested to join the "national X.500 day" organised by SURFnet in co-operation with EEMA in Amsterdam coming June. The "national X.500 day" is postponed by SURFnet until September. In addition they are producing a brochure called "Introducing a directory service", which will be ready in a month or so.

Piloting 93 Transition (15:15)

Colin gave a summary on the four transition steps. Thomas agreed to run both QUIPU and 93 parallel, depending on the availability of the IC 93 implementation. For the pilot NEXOR will make 93 implementations available for beta testing during the pilot phase. This might be on a one year license basis. ULCC, UKERNA and DANTE (as organisations) volunteered to do testing of organisations. There will be parallel testing of QUIPU and 93 implementations on country level by SWITCH and SURFnet. This piloting should start as soon as NEXOR makes the software available. The experiences and results should be presented at the next September meeting.

* Piloting discussions should go to the customers-only mailing list, at least initially. It could have negative impact if the first pilot experiences would go to a broader audience. The discussions could also take place on a separate (but restricted) mailing list. If messages are relevant to a broader audience they will be forwarded to the appropriate mailing lists.

* What is needed is a clear mission statement. The key question is: why change to 93 as the current system works? What are the benefits? We could use the EMA document on why to adapt to X.500.

* Make a press release on our plans on piloting 93? "*Meeting in Salford University agreed on worlds first international 93 pilot.*" Could do this together with those vendors that are joining the project. The release should contain the mission statement describing why to migrate.

Appendix 1

List of attendees :

NAME	REPRESENTING	E-MAIL ADDRESS
2-Barker Paul	UCL	P.Barker@cs.ucl.ac.uk
3-Berkhout Vincent	DANTE	V.Berkhout@dante.org.uk
3-Chadwick David	Salford University	D.W.Chadwick@iti.salford.ac.uk
2-Fantou Patrick	Siemens-Nixdorf	Patrick.Fantou@mch.sni.de
3-Flood Ronan	ULCC	R.Flood@Noc.ulcc.ac.uk
2-Hedberg Roland	SUNET	Roland.Hedberg@umdac.umu.se
3-Jurg Peter	SURFnet	Peter.Jurg@surfnet.nl
3-Kaittola Marko	DANTE	M.Kaittola@dante.org.uk
3-Lenggenhager Thomas	SWITCH	Lenggenhager@switch.ch
2-Meulemans Nils	University of Brussels	meulemans@helios.iihe.rtt.be
3-Pasternakowsky Silvio	DFN	DS-Manager@TU-Chemnitz.DE
3-Pinto Fernando	University of Minho	fernando@uminho.pt
3-Robbins Colin	NEXOR	c.robbins@nexor.co.uk
1-Stiefer Marc	RESTENA	MSTIEFER@restena.lu
2-Tomas Celestino	RedIRIS	celestino.tomas@noc.rediris.es
3-Klunder Hans	University of Delft	Hans.Klunder@RC.TUDelft.NL
3-Tillotson Rodney	UKERNA	Rodney.Tillotson@UKERNA.ac.uk
1-Kille Steve	ISODE Consortium	S.Kille@isode.com
1-Emery Nick	DEC	Emery@emery.reo.dec.com

And the numbers mean:

- 1- attended the first day
- 2- attended both days, but not the customers meeting
- 3- attended both days and the customers meeting

This page was written by Vincent.Berkhout@dante.org.uk

NameFLOW LDAP Pilot meeting

London 29 May 1998

Invitation

DANTE is trying to run an LDAP Pilot as described in "Draft Plan for a NameFLOW LDAP Pilot" [ref. JH(98)010], available from:

<http://www.dante.net/np/LDAP-Pilot-Plan-01.txt>

We intend to hold a meeting at the [NOVOTEL, London Heathrow](#) on the 29th of May 1998 to discuss the Pilot and how it will be executed. ([How to get there](#)) There is a shuttle service operated by Speedlink from Heathrow Airport to Novotel which runs from 05.30am until 11.30am every 10-15 minutes. A single ticket cost £2.00 and a return ticket costs £3.50 and the shuttle departs from the arrival level of all terminals.

This invitation is primarily aimed at the representatives of DANTE's NameFLOW customers, although other interested parties maybe allowed to attend if considered appropriate by DANTE and/or it's customers.

Last updated: Wednesday, 10-Jun-1998 00:00:00 BST

Move to [[DANTE|FLOW Services|NameFLOW](#)]

Contact: Peter.Gietz@dante.org.uk

NameFLOW LDAP Pilot meeting

London 29 May 1998

Agenda

- 10:30 Introductions (VB)
- 10:45 NameFLOW Status Report (VB/JH/PG)
 - Usage statistics
 - X.500(93) Pilot
 - Year 2000 problem in QUIPU
 - DESIRE II
- 11:30 Status report on Directory development
 - IETF (PG)
 - ISSS/WG-Dir (OP)
 - SURFnet (ML)
 - SUNET (RH)
 - SWITCH (TL)
 - GUNET (PA)
 - Belnet (OP)
 - others
- 12:10 Presentation on the Internet Directory Forum (CB)
- 12:30-13:30 Lunch
- 13:30 Presentation of the LDAP Pilot (JH)
- 13:45 Detailed discussions on refining the Plan
- 15:00 Production of a Provisional Time-table for the Pilot
- 15:30 Assignment of Actions
- 15:45 Any other business
- 16:00 Date of next meeting and close

Abbreviations:

CB Clive Betteridge
JH John Horton
ML Marjon Luites
OP Olivier Paridaens
PA Panayotis Astithas
PG Peter Gietz
RH Roland Hedberg
TL Thomas Lenggenhager
VB Vincent Berkhout

Move to [\[DANTE|FLOW Services|NameFLOW\]](#)

Contact: Konstantin.Chuguev@dante.org.uk

NameFLOW LDAP Pilot meeting

London 29 May 1998

Minutes

Attendees

Andrew Findlay (AF), Brunel University
Clive Betteridge (CB), Internet Directory Consortium
Colin Robbins (CR), NEXOR
Damy Mahl (DM), Brunel University
David Chadwick (DC), University of Salford
John Horton (JH), DANTE
Marjon J. Luites (ML), SURFnet
Olivier Paridaens (OP), BELNET
Paniotis Astithas (PA), GUNET
Peter Gietz (PG), DANTE
Rodney Tillotson (RT), UKERNA
Roland Hedberg (RH), SUNET
Ronan Flood (RF), ULCC
Thomas Lenggenhager (TL), SWITCH
Vincent Berkhout (VB), DANTE

1. Opening

VB opened the meeting and welcomed the participants.

2. NameFLOW Status Report

2.1 Usage statistics

VB gave an insight to the latest NameFLOW statistics. There is still an increase in the number of DSAs as well as in the total number of organisation entries. VB showed some web to X.500 gateway statistics indicating that the number of successful hits is relatively increasing although the number of DSA unavailable marginally increased. CB asked if VB thought that the statistics were optimistic. VB answered that there is still some growth of X.500 and the only missing component is probably LDAP deployment.

2.2 X.500(93) Pilot

JH reported on the results of the X.500(93) Pilot. Although DAP, LDAP and DSP worked fairly well,

there were severe problems in DISP in terms of multiple vendor interoperability, as well as in ACI, being too complex to administer. CR remarked that the software tested was "early deployment software" and the market had moved on significantly since then.

2.3 Year 2000 problem in QUIPU

PG summarised the problem being all QUIPU-DSAs will fall over on 1.1.2000, especially in terms of replication. To fix the bug would take at least three man/months for programming and testing. The general agreement was that this would be too much an effort to be spent on an anachronistic technology. PG suggested to use the Y2K compliant version 3.1 of Isode, if it is affordable. AF and DM noted that universities would not buy the software, but for the NRNs the purchase wouldn't be a problem. A work around could be to make a consistently performed replication on file level 1.1.2000 at 0:01.

2.4 DESIRE II

VB gave an overview of the DESIRE II project including the financial budget as well as the working packages of this EC sponsored project. The working package partners next to DANTE is SURFnet as main contractor with the ACCU doing implementation. Brunel University will do work for DANTE and Lund University will implement CIP. The key issue is Directory indexing, not LDAP testing. This will require protocol support of LDAP and CIP.

3. Status reports on Directory development

3.1 IETF-work

PG gave a presentation on the LDAP work going on in numerous IETF WGs, showing the wide range of possible fields of LDAP deployment. The work of the LSD-WG seems to be most valuable for DANTE's future plans. CR made the remark that the chaotic situation of LDAP Internet Drafts could be a process initiated by Microsoft, to delay LDAP standardising allowing them to sell their upcoming Active Directory. Some vendors try to counter Microsoft by making LDAP client software freely available in source code and possibly in Q3-98 server software as well.

3.2 ISSS/WG-Dir

OP reported shortly on directory work going on in the ISSS/WG-Dir, which is a "continuation" of the EWOS/EG DIR directory profiling efforts. The currently proposed LDAP profiles are publicly available and comments are welcome. For more information see <http://www.cenorm.be/iss/Workshop/DIR/Default.htm>

3.3 BELNET

OP gave a short report on directory work in BELNET on behalf of Jean-Marc Verbergt. BELNET uses a combination of Isode (v3.1) DSAs and SLAPDs reachable via the Critical Angel (now

Innosoft) X.500 enabler. They experienced problems with Isode (v4.0) in respect to connecting the DSAs to SLAPDs and would like to hear from others about their experience with 4.0.

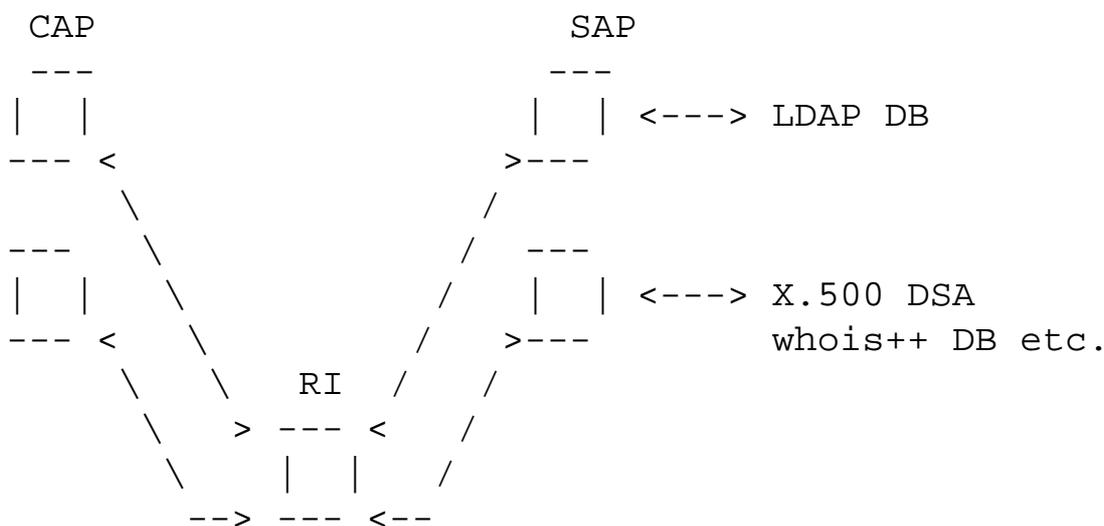
3.4 SURFnet

ML gave a presentation on the current situation of SURFnet activities. QUIPU DSAs are operating in combination with X.500(93) and LDAP v2+3 servers reachable via the enabler gateway.

Additionally, there is an index server, see <http://search.surfnet.nl/naam/index.html>. The future plans are to install a country-level LDAP referral server, to do interoperability testing of different LDAPv3 products and to develop a Dutch LDAP backbone. ML interpreted the instructions from SURFnet to go for a LDAP only solution, however whether this means that X.500 is completely going to be phased out, is open to discussion.

3.5 SUNET

RH reported on the Swedish situation, introducing his Web-LDAP-Gateway WIXI (<http://wixi.umu.se:1400/c=Se>), which conforms to [draft-IETF-lsd-nandi-00.txt](#). He then presented his work on the TISDAG project for a protocol independent referral index (RI), serving different protocols (e.g., X.500, whois++, LDAP v2+3, email interface, etc.). The solution for this are Client and Server Access Points (CAP and SAP), which are protocol specific interfaces to the internal "Hedberg-Protocol" used by the RI:



A Pilot to test this implementation will start in Sept. 1998.

3.6 SWITCH

TL gave an insight to the Swiss situation, where not much changed in the last 12 months. One Digital DSA and 7 Isode based (v1.1 - v3.1) DSAs are in operation. For four organisations the data gets centrally bulk loaded to X.500. Isode 4.0 is currently being tested. The handout of TL about a test load of the domain name registration info of CH and LI loaded into Isode 4.0 was not discussed due to lack of time.

3.7 GUNET/GRNET

Finally PA reported on the ongoing plans in Greece to build an LDAP-only directory service for Greece by connecting multiple organisational LDAP servers to one national LDAP server. These plans were based on the assumption that X.500 would be disconnected all together. After the discussion in the meeting it became clear that some X.500 infrastructure will remain as part of the NameFLOW service. PA nevertheless sees no real prospect of a national Greek DSA.

4. The Internet Directory Consortium (IDC)

CB presented his initiative to start a new global directory group after the close of the EuroSInet group and the withdrawal of the Internet Mail Consortium (IMC) from directory work. The IDC is intended for vendors, developers and other directory interested parties. The main aims are representing directory industry needs, developing test suites, interoperability testing and schema registration. Up to now five companies have signed up for membership, DANTE being one of them. Before starting the group a critical mass of at least ten members has to be found. Four additional companies seem to be interested. The response of the group was that as long as membership costs are > \$2,500 it is outside the interest of the academic community. Alternative arrangements should be considered for this particular interestgroup.

5. The LDAP Pilot

The last agenda item was the presentation of the LDAP Pilot by JH, a detailed description of which can be found at <http://www.dante.net/np/LDAP-Pilot-Plan-01.txt>.

In summary this pilot is the testing a LDAP only solution for a distributed directory service. LDAP fulfills the requirements of the DANTE customers, which are openness, cheapness, easy management, interoperability and extensibility. The proposal includes an infrastructure for interconnecting LDAP servers world wide through a backbone LDAP server operated by DANTE. This server gains its information on the connected LDAP servers via discovery robots. The referrals are located in national index servers provided by the respective NRN organisation or by DANTE. The service should support geographical naming as defined in X.521 as well as Domain Component directory naming defined in RFC 2247. NameFLOW will not act and has never acted as naming registration authority but may assist in resolving naming conflicts. The proposal defines requirements for LDAP client software as well as for the backbone and country level LDAP servers.

The lively discussion focused mainly on three fields: General reservations, indexing, and naming schema.

General reservations brought forward were that LDAP technology is not mature yet for deployment in a production service (CR). JH made clear that the proposed project is a pilot, i.e., research work and only if successful will it be deployed. There should therefore be an alternative technology in the background in case LDAP fails to provide a dependable distributed service before the QUIPU Y2K bug makes the public domain X.500 technology unusable. A hybrid service was proposed combining an X.500 based backbone of single vendor root and country level DSAs with organisational LDAP

servers. The present representatives of the DANTE customers voted with a majority in favour of this hybrid service.

AF expressed general doubts on the usefulness of indices, arguing that the result would be similar to web indexing a la AltaVista. A search for "John Smith" in the whole world could result in up to thousands of referrals. Also general problems of scalability have to be handled. The current pilot proposal includes only an organisation index, which was acknowledged to be a good step forward. The group also agreed to design the index user interfaces in such a way that default searches are not world wide.

The support for both, the X.521 and the Domain Component naming schema, include mapping problems and the decision where to locate the actual DIT. In any case such a mapping would be a valuable contribution for the whole Internet.

The group generally agreed on the following points:

- To do indexing in DESIRE II project
- To investigate a hybrid X.500-LDAP service in place of the current QUIPU infrastructure
- To move the LDAP Pilot forward

6. Actions

To be performed by DANTE:

- Formation of two Pilot core groups: one for indexing as part of the DESIRE II work and one for the LDAP NameFLOW Pilot.
- Set up mailinglists for the two core groups
- Investigate a hybrid X.500 - LDAP NameFLOW service until June 1999 as fallback possibility.
- Re-issue the pilot plan

7. Date of next Meeting

The next NameFLOW meeting will be on 16 November in Brussels or alternatively in Amsterdam.

Move to [[DANTE](#)][FLOW Services](#)[NameFLOW](#)]

Contact: Konstantin.Chuguev@dante.org.uk