

NameFLOW-Paradise

Quarterly Service Report October-December 1995

TABLE OF CONTENTS

Introduction

Operations

1. Operations/helpdesk
2. Outages
3. Issues
4. Statistics
5. NP 93 Pilot Update

Information

1. Information servers
2. Reports
3. Papers

Liaison

1. EEMA
2. EDF
3. IETF

Appendices

1. APPENDIX 1 Helpdesk summary
 2. APPENDIX 2 World Root DSA and LDAP summary statistics
 3. APPENDIX 3 Public DUA summary statistics
 4. APPENDIX 4 WWW/FTP/Gopher/mail summary statistics
 5. APPENDIX 5 IETF ASID WG meeting minutes
 6. APPENDIX 6 ETF IDS WG meeting minutes
 7. APPENDIX 7 NOMENCLATOR Internet Pilot
 8. APPENDIX 8 IETF FIND WG meeting minutes
-

Introduction

This Quarterly Report reflects the NameFLOW activities and operations for the months October, November

and December 1995. The report is intended for people interested in the NameFLOW service and in particular those working for the national networks responsible for the National Directory Services. The report deals respectively with the operational aspects, the information aspects and liaison activities.

The Quarterly reports will be available in paper format to DANTE's customers. An electronic copy will be made publicly available via the web*, without customer sensitive information where appropriate.

For questions about this report or the NameFLOW service, please contact:

Vincent Berkhout
DANTE
Lockton House
Clarendon Road
Cambridge
CB2 2BH
UK

Tel: +44 1223 302 992

Fax: +44 1223 303 005

E-mail: V.Berkhout@dante.org.uk

* URL: <http://www.dante.net/nameflow.html>

Operations

Introduction

This is an overview of operational activities as carried out by ULCC on behalf of DANTE during the months October, November and December 1995.

1. Operations/Helpdesk

There were several changes to the root of the DIT during the quarter. Argentina (c=AR) joined the Directory with its masterDSA cn=Gorrion. An experimental FLDSA, cn=Giant Panda, was registered. This is managed by a researcher at Northeastern University in China, who is looking into the possibilities of running X.500 services in CERNET, the China Education and Research Network. However, c=CN is not yet ready to appear in the Directory, and will be connected on a "read-only" basis.

The FLDSA cn=Arowana was registered; this is a slaveDSA for Singapore (c=SG).

The masterDSA of the top-level entry o=North Atlantic Treaty Organisation (NATO) changed from cn=Violaceous Trogon to cn=Manfred Woerner; the entry for Violaceous Trogon was removed from the root and placed under the organisation.

Problems of accessibility to two FLDSAs (cn=Guinea Pig/c=HU and cn=Dorcan Gazelle/c=IL) were resolved, but problems remain with several other FLDSAs, which have been uncontactable for extended periods.

In cooperation with DANTE and FLDSA managers, the Directory entry for each country was modified to include a labeledURI attribute pointing at a suitable web page for that country. Storing URLs/URIs in the Directory is the subject of a current Internet Draft.

A discrepancy in Quipu's use of the dspchaining tailor parameter was noted, between the default behaviour described in the reference manual and that exhibited by the DSA. This difference was pointed out to FLDSA managers, since it might significantly affect the performance and availability of FLDSAs.

2. Outages

Significant outages of service totalled approximately 81 hours in October (including 2 hours scheduled "at risk"); 64 hours in November (including 2.5 hours scheduled "at risk"); and 149 hours in December. An X.25 problem, affecting EuropaNET X.25 and public X.25 access to the public-access DUA and to the DSA Ocellated Turkey, went unnoticed for a large proportion of the quarter (approximately 66 days out of 92), and on into January. Details of outages are available in the monthly reports.

3. Issues

Several aspects of DIT structure came up during the quarter. Singapore (c=SG) is using the non-standard attribute type sectorName in RDNs immediately below c=SG, for example c=SG@sectorName=Electronics. DSAs/DUAs that do not have the Singaporean OID definitions will see sectorName as the numeric OID 1.2.702.1.1001.2.1.

The Directory manager for NATO has reorganised its DIT, such that the subordinate entries of o=North Atlantic Treaty Organization are now in themselves organizations, rather than organizationalUnits.

In both cases the resulting DIT structure is unusual, and does not comply with the example in X.521 annex B. While this is believed not to contravene X.500, there may be "visibility" problems if DUAs expect the DIT to be structured like X.521 annex B.

It was decided at the NameFLOW-Paradise meeting in September to discontinue the central public-access DUA/DSA/LDAP server in January 1996. This will require some reconfiguration of the central systems, which must be performed alongside the IP readdressing mentioned in the previous quarterly report, yet to be done. This will be an opportunity to assess whether other of the central services should also be discontinued, for example the mail info-server.

4. Statistics

Summaries of the service statistics for the quarter are attached in the Appendices. Full statistics and world-root DSA hourly operations figures are available on the NameFLOW-Paradise FTP/Gopher info-server, under:

```
gopher://gopher.nameflow.dante.net/11/statistics/
```

```
ftp://ftp.nameflow.dante.net/statistics/
```

Two things are noticeable from the statistics: the decline in the number of accesses to the central DUA

service, and the reduced demand for PARADISE reports from the helpdesk (not visible in the summary). Although there is usually a drop in the use of the service over the last quarter of the year, it seems sharper this year than previously. In fact the October figure was less than that for September, which has not happened before.

Changes to the logout message on the DUA service made at the start of October are thought to be the reason for this. The message now announces the imminent closure of the service, gives pointers to alternative access points, and refers interested users to the online copy of the 1994/1995 annual report. The statistics for the use of the UK academic public-access DUA, also based at ULCC and referred to in the new logout message, show that demand for that service increased by many times what would have been expected for the quarter.

5. Update on the NameFLOW 93 pilot

[Distributed on 19 January 1995]

The NameFLOW-Paradise Directory System is mainly based on Quipu, a specific flavour of X.500 with several enhancements to the ISO standard and therefore not fully compliant. The initial edition of the standard, published in 1988, had a few shortcomings, such as the lack of an replication mechanism. The shortcomings were acknowledged and an improved edition of the X.500 standard was published in 1993. This test plan describes some of the advantage of the 1993 standard, explains why NameFLOW-Paradise wants to migrate to 1993 from the current Quipu model.

Before a migration to 1993 can take place a test is planned in the beginning of 1996. The test plan describes a one-week test, methods and participating organisations. If possible the test structure should remain in place to provide the basis for the future operational service. This paper is not an introduction to X.500(93) and the reader is supposed to have basic knowledge of X.500 Directories.

The Objectives

The NP-93 migration has two objectives:

1 Introduce a Directory system based on X.500(93)

2 Phasing out the Quipu system.

While deploying the X.500(93) system the Quipu system will be gradually phased out. Once a X.500(93) service is extensively tested and operational, Directory services based on Quipu systems can be discontinued. The long term goals is to have an X.500(93) system without the need for "backwards compatibility" (or gatewaying from X.500(93) to Quipu or X.500(88)).

Why a 1993 pilot?

The pilot is needed to gain experience, find failures in the software or detect defects in the 1993 edition of the X.500 standard. The pilot will pave the way for an operational Directory infrastructure primarily based on X.500(93). NP sees the opportunity to perform the first field test as there is limited practical experience deploying X.500(93) on a large scale. The goal of NameFLOW is to use X.500(93) to provide a multi system/multi vendor Directory service allowing other Directories (also non-R&D!) to connect and share

information.

The X.500(93) edition of the standard was improved in the following areas:

- * access control
- * replication (sub tree and incremental)
- * schema knowledge

The upgraded X.500(93) system will be improved with:

- * new management tools
- * better interfaces to the Directory (e.g. using LDAP)
- * more reliable with complete data
- * shorter response times

Why not a 1988 test?

There are several reasons for testing the most recent version X.500(93) and not X.500(88). The most important reason is that X.500(88) does not provide the required operational functionality, such as replication. Using X.500(93) provides at least similar functionality to Quipu model and has additional enhancements e.g. in the security area. Another reason is that software development is currently focused on X.500(93), and very little attention is paid to further development X.500(88) or Quipu. Only a limited number of X.500(88) implementations exist and it is expected that this number will not increase. Most X.500(88) implementations can now be upgraded to X.500(93) without major problems. The final reason not to test X.500(88) is that it already successfully tested and operational. (See [1] for further information.)

The Transition Scenarios

There are two possible scenarios for the transition:

One-step transition:

Quipu -> X.500(93) (direct)

Two-step transition:

Phase one: Quipu -> X.500(88)

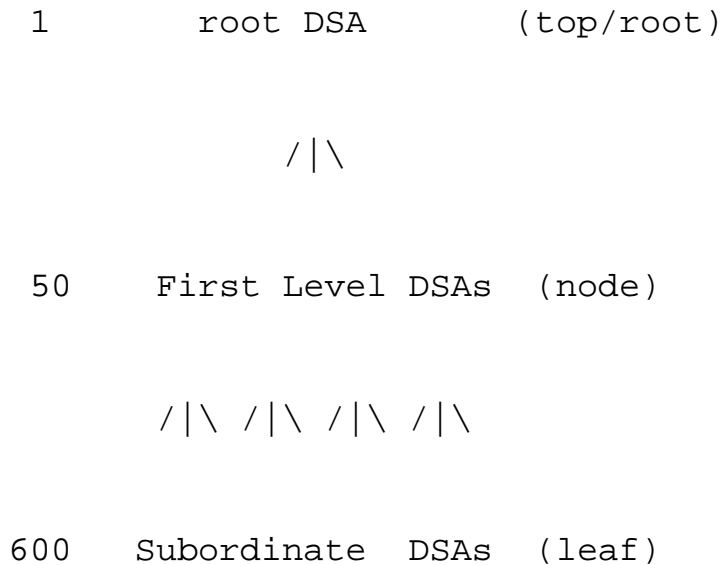
Phase two: X.500(88) -> X.500(93)

A one-step transition (Quipu -> 93) is preferred over a two-step-transition (Quipu -> 88 -> 93), even though it will be more challenging as not all implications can yet be foreseen. The one-step transition is expected to require less effort and time. Currently there are only a few pure X.500(88) implementations connected to NameFLOW-Paradise and introducing this third X.500(88) flavour on a large scale would be undesirable. A major benefit of the one-step transition is that once the X.500(93) model is supported, it should allow both Quipu and X.500(88) implementations to connect (Quipu -> 93 <- 88). The reason for this is that the 1993 version of standard supports all 1988 features.

The major change from the current Quipu model and the an X.500(93) model is that the root DSA can be obsolete as each First Level DSA can have a root entry by itself. In the current Quipu model the information is managed one level up, e.g. the Root DSA manages all country entries, a country DSA (FLDSA) manages all organisations. The principle difference is that in the X.500(93) model the country (FL)DSA manages the country entry itself!

Bottom-up vs. Top-down

The current functional model (depicted below) is based on one root DSA at the top level, the next level are called First Level DSAs (50) and finally other DSAs (600) at the lower levels.



The bottom-up approach requires that all subordinate DSAs must to be replaced before a superior DSA can be replaced. Bottom-up will mean to replace the greatest number (600) of "leaf" DSAs first. In practice this will be unfeasible as a number of leaf DSAs are poorly maintained. In addition it would need a great effort to co-ordinate and some organisations might be reluctant to migrate. To minimise the impact on the current operational service a separate NP-93 tree ("DIT") will be used, whereby the test tree a partial shadow copy of the real tree is. This will avoid some of the migration problems but requires synchronisation (conversion tools) between the two trees.

So why the top-down approach? As explained above the root DSA functionality could be reduced, in theory the root DSA could even be considered obsolete. In practice this means that every FLDSA must have it's own root entry. All that is necessary for a FLDSA are knowledge references to the other FLDSAs. In short: the root DSA as known in the Quipu model can be replaced with a simple file containing the references making the Root DSA obsolete and the top level is replaced. This step could be repeated for each node (level) of the tree. The implication is that each branch of the tree can be replaced at a time convenient to the National Service provider managing their branch of the DIT, and repeat this for the subordinate level. Using this strategy does not require a "switch day" where one complete level or even the complete tree will switch to a X.500(93) based infrastructure, but will allow a node to run two DSA in parallel until the subordinate level fully support X.500(93).

Although a simple reference file would suffice in the beginning some root DSA functionality should be maintained as the first experiences with X.500(93) learned that some functionality of the current Quipu model/root DSA must be maintained to make a global system manageable. The two "root DSA functions" in the 93 model are:

1. distribute all knowledge references to First Level DSAs,
2. some form of replication. (See [2] for full details).

When?

The time frame strongly depend on the availability of X.500(93) software. Although software vendors state that their software is fully compliant to the 1993 edition of standard, it still has to be seen what is implemented it is proposed to have dedicated weeks or consecutive days to perform the tests. The Experience Test (Phase 0: done in August/September 1995) will be followed by tests done in three phases:

1. Root Context Test (Pure X.500(93) replication)
2. Inter working & Scaling Test (using X.500(88) and Quipu)
3. Operational Transition

The Initial Experience Test, referred to as phase 0 was a small test done by only a limited number (three) of countries and several problems were encountered: unavailable '93 software, Root Context replication (DISP agreements, replication errors) and unsupported EDB conversion tools.

Phase One: Root Context Test

The Root Context Test will be done on 12 to 16 February 1996 and done by a small group (ten) of participants. The actual test will be defined at the EuroSInet test writing workshop and the target is to test the top level/Root Context, in particular DISP. If possible the test can be extended to DAP and DSP. It is advised to install needed software in the week prior to the test. A week is reserved to complete the test and this will be more than sufficient time.

- 12-02: installation and conversion of test entries
- 13-02: test DISP
- 14-02: Test DSP and DAP and connecting to external DSAs
- 15-02: full interconnection test
- 16-02: reporting/discussions

It should be noted that the time frame above will be more than sufficient to perform all the tests and more activities can be done in one day. A meeting is planned on 22 February 1996 at Schiphol Airport to discuss the results and prepare Phase Two.

Phase Two: The Small Transition (Or Inter Working & Scaling Test)

On the precondition that the Root Context works properly the next phase will be to test inter working (including Quipu and X.500(88) coexistence) and scaling. The test is planned for May/June and will allow involvement from organisations with an interest in X.500(93) and the new infrastructure. A meeting to discuss the result is planned for the June NameFLOW-Paradise Managers meeting (11-14 June, Brussels)

This will be a large scale test, focusing on ACI, coexistence between Quipu/X.500(88) and LDAP server and clients tests. In X.500(93) ACI is differently modelled (outside an entry) and this could have an impact on performance. This large scale test is intended for big organisations with large data sets. During this phase the first non-Quipu (derivative) DSAs can be introduced. A major part of access to the Directory is via LDAP servers and clients and should therefore be tested.

Phase Three: The Big Migration (Or Operational Transition)

All will depend on the outcome of the two previous test. If both are successful The Big Migration to an X.500(93) infrastructure can start. It is hoped that this can start in September 1996 after the Holidays. It will involve all DSAs out there. The Big Migration could be postponed if documentation for subordinate level is not yet available.

The Tests

The tests to be used will be defined in co-operation with EuroSInet on 24, 25 January 1996 in Cambridge where four NameFLOW participants will participate in an X.500(93) test suite writing work shop. It is not expected to perform a full X.500(93) test in Phase One, as the target will be to test the Root Context using DISP. If possible the test the test can be extended with DAP, DSP and the modified ACI approach. It is acknowledged that the test suite will not cover everything so there will be a need for extensions, such as the needed enhancement for the Root Context, LDAP service and so on.

EuroSInet Test Writing Workshop

NameFLOW-Paradise will have four representatives at the EuroSInet workshop.

The basis for the EuroSInet the X.500(93) work shop is the X.500(88) test suite. The test suite has an emphasis on inter working rather than conformance testing and the current suite already includes DAP and DSP, but exclude DISP and ACI test. The work shop will be split in two subgroups, one group defining DISP tests, the other defining ACI test:

DISP: Ronan Flood and Vincent Berkhout

ACI: Tomasz Wolniewicz and Damanjit Mahl

Although the schedule is tight (work shop 24 January 1996 and test Phase one 12 February 1996) the tests should be ready and distributed before phase one commences. The tests of EuroSInet have an inter working (rather than "conformance") character reflecting real usage.

Documents

The outcome of the test needs to be documented.

- * Test Experiences (This plan/Tests/Problems/Evaluation report)
- * Discussion paper for continuation (Suggestions for next phase).
- * Possibly a first draft of technical manual "How to upgrade a Quipu (or 88) DSA to an X.500(93) DSA"

Challenges Ahead

Problems may occur in the following areas:

- * Conversion of the Quipu EDB format to X.500(93) format and vice versa.
- * Effective list and search operations using the 1993 model
- * Access Control Information is differently modelled, conversion could be problematic.
- * Implementing the Root Context, how to manage knowledge references and what level of replication will needed.
- * How to support Quipu (ISODE version 8.0) and pure X.500(88) DSAs in the NP-93 infrastructure.

* Will current LDAP clients work, in particular those using University of Michigan LDAP servers 3.1 and 3.2.

Mailing List

A dedicated NP-93 mailing list is already operational.

Name list: NP-93@dante.org.uk

Subscription: NP-93-request@dante.org.uk

Archive: not available

Acknowledgements

Many thanks to the following organisations and people who participated and made valuable contributions to the tests and documents.

Brunel University (UK)

Damanjit Mahl damanjit.mahl@brunel.ac.uk

DANTE (UK)

Vincent Berkhout vincent.berkhout@DANTE.org.uk

Nicholas Copernicus University (PL)

Tomasz M. Wolniewicz twoln@mat.uni.torun.pl

SURFnet (NL)

Peter Jurg Peter.Jurg@SURFnet.nl

SWITCH (CH)

Thomas Lenggenhager lenggenhager@switch.ch

SUNET (SE)

Roland Hedberg Roland.hedberg@umdac.umu.se

UKERNA (UK)

Rodney Tillotson Rodney.Tillotson@ukerna.ac.uk

ULCC (UK)

Ronan Flood Ronan.Flood@noc.ulcc.ac.uk

Delft University of Technology (NL)

Hans Klunder Hans.Klunder@rc.tudelft.nl

NEXOR (UK)

Colin Robbins c.robbins@nexor.co.uk

ISODE Consortium (UK)

IC-info@isode.com

EuroSInet (UK)

Chris Ridd chris@imc.exec.nhs.uk

John Horton J.Horton@imc.exec.nhs.uk

and many others.

References:

[1] OIFP reports <ftp://ftp.nameflow.dante.net/paradise/oifpfinal.txt>

[2] "Managing the Root Context", David Chadwick, "DANTE in Print #18",

<<http://www.dante.net/pubs/dip/18/root.txt>>

Information

Information servers

As part of the information service of NameFLOW-Paradise DANTE operates several servers. There are 'historical' PARADISE information servers, such as ftp, gopher and e-mail operated by ULCC. In addition a web server is now fully operational as part of the DANTE World Wide Web service. Usage statistics for each server are included in Appendix 4.

Reports

Quarterly and individual monthly reports are available on-line from DANTE's WWW server:

3rd Quarter 1995: <http://www.dante.net/np/report/qr/95Q3.html>

October: <http://www.dante.net/np/report/mr/mr9510.html>

November: <http://www.dante.net/np/report/mr/mr9511.html>

December: <http://www.dante.net/np/report/mr/mr9512.html>

Papers

During the fourth quarter of 1995 two papers have been produced relating to NP developments issues. Both are available from the NP www server.

Whois X.500++ ?

David Chadwick

<http://www.dante.net/np/ds/whoisx500.html>

Managing the X.500 Root Naming Context

(submitted as Internet Draft)

David Chadwick, January 1996

<http://www.dante.net/pubs/dip/18/root.txt>

Liaison

EEMA Directory Committee *15 December 1995, Lotus Headquarters, Staines, UK*

Vincent Berkhout

Only a few "core" group members of the EEMA Directory Committee were present. During the introduction (tour de table) two interesting topics were mentioned, namely the start of two national initiatives: a Danish Directory Forum and a Finnish Directory Forum.

The EEMA black book "X.500 Directory Synchronisation Product Guide" (by Roger Molesworth) was discussed and will be finalised before the next meeting.

The EEMA Web server is now operational and can be reached via <http://www.eema.org/eemahq>. The web site gives access to an internal directory containing all its members and distribution lists.

The document "Guidelines for Corporate Directory Deployment" by Amanda Edwards and John McKinnon is basically complete and needs "minor" editorial changes. The document has the popular subtitle "Hitchhikers Guide to the X.500 Directory". The project "Top Level Naming for Europe" seems to be one of the hottest items at the EEMA and consumed a major part of the meeting. It seems that expectations about the contents vary and it will need a major revision. The document is being rewritten chapter by chapter and it will take another half year before it will be finalised.

With respect to the organisation of the meeting VB put the suggestion forward to distribute the necessary documents well before the meeting. (i.e. minutes of previous meetings, reports, agenda). The reply was that this was a voluntary committee and one could not demand full commitment from its members.

A financial budget for Directories was put forward for the years 1995 and 1996. The following are the current, proposed new and back burner projects:

Current Projects (1995)

Budget*

Guidelines for Corporate Directory Deployment

24.9

TOPOL (Phase 1 & 2)

13.8

X.500/WWW services

5.0

Product Guide

11.6

Global

Dictionary

?.?

Total:

55.3

Proposed New Projects (1996)

Budget*

TOPOL (Phase 3)

40

Brussels Demonstration: The Network

2

Brussels Demonstration: The Applications

3

Directory Technologies
10
Market Forecast
10
Regional Support for X.500
10
Guidelines for Common Values
20
Case Studies
3

100

Total:

Back burners (1996)
Budget*
Coexistence among Directories
--
Support for Euroview
--
Business Case for Enterprise Directories
--
--

Total:

(Note on budget: all amount in Swish Francs)

There are two interesting projects with a clear overlap:

- 1) Directory Technologies; which technologies can be used to build a corporate directory, and who will use which technology to provide a global Directory.
- 2) Coexistence amongst directories. A two phase project, the first phase to analyse current directory technologies and address the problem of deploying multiple directory services (and multiple providers), where one possible solution could be "Common Indexing". The second phase to decide on the technology and perform a test case, e.g. connecting two different directories.

The first project was accepted and also the second phase of the second project (but not funded!).

European Directory Forum

17 January 1996, Stockholm, Sweden

Vincent Berkhout

The new scope (~ mission) statement explains what the EDF is all about:

"The members of the EDF support the goal of providing every European customer with access to Europe-

wide directory services, at affordable prices.

The EDF will assist in the elaboration of solutions for a "universal" virtual Europe-wide directory."

The goal of the group is not to solve today's problems, but to provide goals and guidelines for a universal Directory with a level playing field for all European directory service providers. The EDF group can do valuable work as they can influence European regulations.

The EDF is still in the founding phase but expected to become operational starting from 1 January 1998 (although this was already questioned!). Before actually founding an EDF two questions have to be answered:

1) Is a European wide "universal" directory technically feasible?

2) Is the EDF politically desirable?

The first question addresses whether a 150 million people directory by multiple providers is technically possible and workable. (With the option of accessing at least another 150 million non-European person entries). The second question involves the issue whether the EDF has sufficient mandate to survive the "great forces that be". An EDF can officially be founded as the answer to both questions would be 'yes'. In the structure of all European organisations the EDF will be parallel to the current European Numbering Forum (ENF).

The Players

There are quite important players around the table: CEC, ENF(European Numbering Forum), ECTRA (European Conference of Postal and Telecommunications Administrations), ETNO (European Public Telecommunications Network Operators Association) and ETSI (European Telecommunications Standards Institute). Most people (and their organisation) have a "telephone operator" (TO) background. In addition to these TO there is room for small/coming pan-European service providers to participate, e.g. Telia TeleRespons (SE) and IBM (DE) were represented. The EDF seems to be a sort of consultation/input group for the CEC. What the CEC is looking for are non-traditional service providers to increase pressure on TOs, if possible before the 1998 liberalisation. They are looking for non-TOs to set up a European service and if possible via an organisation such as the EDF.

The meeting

Started with a word of welcome, new agenda and presentations of all the people around the table. I gave a short talk (+/-25 min, no sheets) about NameFLOW. Promised to send DANTE and NameFLOW Information package to all the people who gave me their business cards (done). Then a two hour brainstorm session on the EDF its context (changing environment), requirements, problems and goals. The afternoon session was used to determine the terms of reference (ToR) The ToR will only allow pan-European organisations to participate (and **no** consultants). The ToR were heavily revised to allow non-traditional, but European organisations to become members; the EDF will not be exclusive to founding members and pan-European organisations.

The Flight Back (with Leo Koolen of CEC)

I was on the same flight as Leo Koolen (CEC DG XIII) back to Brussels. He could be our "ticket in" to the EC to promote Directories and we had a talk about the current pan-European directory service providers. One requirement is that such a service needs to be based on a trans European Network represented in most European countries. I explained about the TEN-34 project where DANTE acts as co-ordinating partner and proposed a similar strategy for directories. He was interested in what DANTE and its customers could do to "pioneer" a more commercial service. We agreed upon a meeting in Brussels to discuss future co-operation with the CEC and a second meeting with commercial directory providers who already made an offer to provide a Directory service to the CEC itself (?).

IETF

Minutes of the IETF ASID, IDS and FIND Working Groups, as well as an overview of the NOMENCLATOR Project, can be found in Appendices 5-8.

APPENDIX 1 - Helpdesk summary for Oct/Nov/Dec 1995

Country		Number of queries			
-----		-----	-----	-----	
Full Name	ISO Code	October	November	December	Quarter
-----	-----	-----	-----	-----	-----
Argentina	AR	3	-	-	3
Australia	AU	1	-	-	1
Brazil	BR	2	-	-	2
Canada	CA	1	1	-	2
(China)	CN*	1	14	4	19
Germany	DE	1	-	-	1
(Algeria)	DZ*	-	-	1	1
United Kingdom	GB	13	13	5	31
Ireland	IE	1	-	-	1
India	IN	1	4	3	8
Italy	IT	1	-	1	2
Korea	KR	1	-	-	1
(Macau)	MO*	1	-	-	1
Netherlands	NL	-	1	3	4
(Romania)	RO*	-	-	2	2
Singapore	SG	-	1	1	2
Slovakia	SK	1	-	1	2
Taiwan	TW	1	-	-	1
United States	US	13	4	2	19
(South Africa)	ZA*	-	1	-	1
-----	-----	-----	-----	-----	-----
Total Requests		42	39	23	104

(A * by the country code shows that this country has no Directory entry)

APPENDIX 2 - World Root DSA and LDAP summary statistics for
Oct/Nov/Dec 1995

Summary of calls to DSA Giant Tortoise

From 5:35:38 on 30 September to 5:29:26 on 31 December

No. of binds	October	November	December	Quarter
Local	3520	3740	3363	10623
Remote	11470	10999	9702	32171
Total	14990	14739	13065	42794

No. of operations	October	November	December	Quarter
Local	659	583	730	1972
Remote	160862	124696	105409	390967
Total	161521	125279	106139	392939

System usage (calls received)	October	November	December	Quarter
Binds by Directory technicians	10201	10049	9042	29292
Reads of DSA entries	284	213	269	766
Other ops on DSA entries	16	15	45	76
Getedb operations (inc slices)	56624	54027	56472	167123
Spot shadows	24	27	16	67
Total	67149	64331	65844	197324

LDAP usage

LDAP usage from Oct 6 1995 to Dec 28 1995

	October	November	December	Quarter
Connections	36	6	5	47
Total connect time (seconds)	16263	54	694	17011

(17011 seconds is 4 hrs 43 mins 31 secs)

Summary of calls to DSA Ocellated Turkey

From 0:05:27 on 30 September to 0:04:32 on 31 December

No. of binds	October	November	December	Quarter
Local	19243	12890	8540	40673
Remote	10551	9844	10171	30566
Total	29794	22734	18711	71239

No. of operations	October	November	December	Quarter
Local	617946	409640	291135	1318721
Remote	51328	41174	38633	131135
Total	669274	450814	329768	1449856

System usage (calls received)	October	November	December	Quarter
Binds by Directory technicians	23501	17279	13122	53902
Reads of DSA entries	5667	3688	3120	12475
Other ops on DSA entries	31	19	13	63
Getedb operations (inc slices)	639	591	574	1804
Spot shadows	2229	2011	2023	6263
Total	32067	23588	18852	74507

APPENDIX 3 - Public DUA summary statistics for Oct/Nov/Dec 1995

DUA usage (logins to Directory Enquiry service at nameflow.dante.net)

Network	October	November	December	Quarter
Internet	22147	14361	9509	46017
UK academic X.25 (JANET)	200	296	139	635
EuropaNET X.25 10*	4	4	2	
Public X.25 35*	6	11	18	
ULCC dialup	47	11	5	63
Total	22404	14683	9673	46760

(* figures affected by the X.25 problem mentioned in Outages)

Top ten Telnet DUA logins by domain, selected and ordered by quarterly total

Domain	October	November	December	Quarter
-----	-----	-----	-----	-----
edu	12236	7218	4629	24083
uk	3486	2287	985	6758
unresolved	1887	1261	834	3982
com	908	685	679	2272
ca	822	579	357	1758
net	476	460	505	1441
us	384	244	228	856
org	256	173	163	592
nl	219	187	142	548
mil	142*	113	105	360
-----	-----	-----	-----	-----
Total	20816	13207	8627	42650

(* indicates that the domain was not in the top ten for that month)

APPENDIX 4 - Web/FTP/Gopher/mail summary statistics for Oct/Nov/Dec 1995

Web server

TOTALS FOR SUMMARY PERIOD Sun Oct 1 1995 TO Sun Dec 31 1995

	October	November	December	Quarter
-----	-----	-----	-----	-----
Unique hosts	623	514	387	1392
Number of HTML	1890	1769	1208	4867
Number of non-HTML requests	143	182	116	441
Number of malformed requests	93	177	27	297
Total number of all requests/errors	72	71	44	63
Total number of Kbytes requested	19259	18621	22233	60114
Kbytes/day	651810	621602	718937	671468

FTP server

TOTALS FOR SUMMARY PERIOD Sun Oct 1 1995 TO Sun Dec 31 1995

	October	November	December	Quarter
Files Transmitted	1149	1190	1004	3343
Kbytes Transmitted	125021	106738	98382	330142
Average Files Daily	37	40	7	38
Average Kbytes Daily	4032	3557	3643	3744

Gopher server

Gopher usage from Sun Oct 1 1995 to Thu Dec 28 1995

	October	November	December	Quarter
Total connections	48	36	18	102
Total files retrieved	80	41	29	150

Mail server

Mail-server usage from Oct 11 1995 to Dec 21 1995

	October	November	December	Quarter
Copies of help sent	3	1	0	4
Total files retrieved	2	7	3	12
Total requests	5	8	3	16

APPENDIX 5

Access and Searching of Internet Directories

Meeting Minutes

What: Access and Searching of Internet Directories

When: Wednesday, December 6, 1530-1730

- Agenda review/changes

The chair apologized for getting the agenda out so late, and for not producing a proper document archive.

The proposed agenda was reviewed and accepted without changes.

- WHOIS++ status

Patrik Falstrom gave a brief status report on the WHOIS++ query protocol documents. They have been submitted to the ADs for proposed standard status, and should be reviewed at the next IESG meeting.

ACTION: Harald to submit WHOIS++ documents for proposed at the next IESG meeting.

- CIP status

A new working group (FIND) is forming to define the Common Indexing Protocol, and the BOF met just before the ASID meeting. ASID will drop this item from its charter.

- LDAP status

LDAP has been at draft standard status since last March, and the group discussed whether LDAP is ready to progress to full standard. There are multiple independent interoperable implementations. There was a question raised as to whether the Kerberos BIND credentials were supported by any implementation other than the one from University of Michigan. The group agreed that this question should be resolved before LDAP goes forward, and Tim agreed to try and find out. There was another question raised as to whether LDAP had seen enough operational experience. The consensus of the group is that it has.

There was some confusion in the group about the difference between the LDAP protocol and the widely-used University of Michigan implementation of LDAP. The LDAP protocol has had one version change. It went from version 1 to version 2 when the MODRDN operation changed. There have been several versions of the U-M implementation of LDAP released, the most recent of which is version 3.2. Earlier U-M releases had some bugs in the BER encoding that were hampering interoperability with conformant LDAP implementations. These bugs have been fixed, and the implementations now interoperate, though there are still some old implementations out there.

There was also some confusion as to what exactly was being proposed for full standard. Again this appeared to stem from confusion between the LDAP protocol as a front-end to the X.500(88) directory as defined in RFC 1777 and RFC 1778, and the University of Michigan implementation of LDAP, which includes some experimental extensions transparent to existing LDAP clients for doing stand-alone LDAP, passing back referrals, indexing information, etc. It is only the former that is being considered for standardization. The latter are only useful experiments that will hopefully feed into the design of the next version of LDAP.

ACTION: Tim to check on implementations of kerberos LDAP BIND credentials, and put LDAP up for full standard if there are others that interoperate.

- LDAP URL format draft

At the last meeting, the LDAP URL format draft was approved by the group, provided that it be passed by the URI working group for review. This was done, to little comment, and the group now suggests that the draft be progressed as proposed standard, after it is passed by Harald's URI checklist.

ACTION: Tim to submit LDAP URL format draft to ADs for progression as proposed standard.

- labeledURI draft

At the last meeting, changes to Mark Smith's labeledURI draft were discussed. The group consensus was that both labeledURI and labeledURL attributes were useful. Mark changed the draft to incorporate both attributes. The group agreed that with these changes the draft should be progressed as proposed standard.

ACTION: Tim to submit labeledURI draft to ADs for progression as proposed standard

- LDAP/X.500 caching draft

The group agreed that the caching draft was useful, but that the function would be better served by creating an operational attribute, rather than a user attribute to hold the TTL information. Some reservation was expressed about the work, since this is an area the X.500 standard has intentionally avoided. The group agreed that this draft should be revised and progressed as experimental.

ACTION: Tim to revise caching draft and circulate to the list for comment.

- application/directory MIME type draft

Several comments on the application/directory MIME type draft since the last meeting have been incorporated, but a new version of the draft has not yet been submitted. Changes include the addition of a home fax number and change to using multipart/related rather than multipart/mixed.

There was some discussion of potential uses for this draft, from the straightforward carrying of directory information in email from a simple directory query responder, to use as a method of carrying directory synchronization information, to the provision of directory information over HTTP. There was general agreement the draft was useful.

Concern was expressed that the draft defines both a general framework for carrying directory information and the specific content relating to a person. The issue is that the person information implies some schema which should be harmonized across all directory services, if this draft is to be useful as a general mechanism for carrying information. This schema harmonization is already being tackled by the IDS group. The suggestion was made, and the group agreed, that the draft be split into two. One draft would define the MIME type and general framework for carrying directory content of different types. The other draft would define the content for person directory information.

A third draft was proposed to define the necessary content for handling directory synchronization applications.

ACTION: Tim to split the application/directory MIME type draft into two drafts (one framework, one person information).

ACTION: Greg Vaudreil and Ed Reed to write an application/directory MIME type content draft for directory synchronization.

- leaf/nonleaf draft

This draft was withdrawn by the authors (with the blessing of the group), since it had been pointed out on the list that the main function of distinguishing leaf from non-leaf objects could be done by using an already defined X.500(93) operational attribute.

- String encoding of presentation address draft

The string encoding of presentation address draft has been revised by Steve Kille to support the new IPv6 addressing scheme. The group agreed that the draft should go forward, provided that it be circulated to the ASID list for comment. The document is a product of the TOSI group, so not directly in the ASID charter.

ACTION: Steve to circulate the presentation address draft to the list.

- Storing PGP attributes in the directory draft

Roland Hedberg gave a brief presentation on his draft defining an object class and attributes for storing PGP certificates in the LDAP/X.500 directory. The presentation prompted much discussion.

The debate focused on whether it is better to store certificates in the directory directly, or to store a URL pointing to the certificate in a PGP key server. The primary advantage of the latter method is one of easier maintenance. If the user needs to maintain their key(s) in a PGP key server anyway, the added administration and potential for inconsistency introduced by storage in the directory is a bad thing. On the other hand, storing only a pointer in the directory places an extra burden on clients, which will have to implement an additional access protocol to retrieve the key from the PGP key server.

The group was fairly evenly divided between the two approaches, prompting the suggestion that the draft be changed to define attributes appropriate for both solutions. The market could then decide which was better.

ACTION: Roland to revise the PGP draft to incorporate both solutions, and post the revised draft to the list.

- SUM500 draft

Vincent Berkhout gave a brief presentation of his SUM500 draft, which defines a method of mining the Web and other information services for X.500 information. Vincent's idea involves using standard HTML pages that, if present on an organization's web server, could be read and parsed to produce organization and people entries for the X.500 directory.

The group thought the draft useful, and there was discussion of Vincent's proposal to rewrite the draft to use the application/directory MIME type as the standard format.

ACTION: Vincent to revise the draft to reference the MIME type draft, and post the revised draft to the list.

- LDAP API RFC 1823

Tim announced that informational RFC 1823 was available that documented the University of Michigan LDAP API. The information was presented to the group for informational purposes only, though a short

discussion ensued about the appropriateness of doing API work in the IETF.

- LDAPv2 presentations and discussion

Dave Horvath of Chromatix gave a presentation on the US Navy's work to produce a secure version of LDAP. The Navy's approach was to implement MDAP (Minimal DAP - essentially full DAP PDUs over some other transport mechanism) as extensions to LDAP. Their implementation is called SLDAP (Secure LDAP), and it supports strong authentication and end-to-end digital signing of search operations and results.

Dave described how they produced a new Windows LDAP DLL that implemented the protocol extensions and used the Fortezza card for signing. The DLL approach means that existing Windows LDAP clients can be used unmodified with the new DLL and still receive the benefits of strong authentication and signed operations.

Kevin Jordan gave a brief description of the extensions that CDC has made to their implementation of LDAP to support some X.500(93) operations. The extensions include the addition of a ModifyDN operation, an operation to add a context prefix, and the ability to set new operational parameters, such as the dontUseCopy service control.

There were general apologies from the chair and several other working group members because of the general lack of progress made since the last meeting on the LDAPv2 document. More promises were made for a draft by the next meeting.

ACTION: LDAPv2 volunteers to get cracking and produce a draft by the next IETF.

- AOB

No other business was presented, so the group adjourned almost on time, agreeing to meet again at the next IETF in Los Angeles.

APPENDIX 6

[points of interest: 5,7 and 11]

DRAFT Minutes of the IDS Working Group

Dallas, Wednesday December 6th 1995

The previous meeting's minutes were agreed and there were no changes or additions to the Agenda.

1. Liaison Reports

It is preferred that liaison reports be circulated to the mailing list prior to the meetings so that they are accessible and there is a record of them.

AARNET, Nameflow and Long Bud reports will be circulated to the mailing list, Nomenclator has already

been sent.

Patrik Faltstrom gave an overview of the WHOIS++ Pilot. This currently consists of a number of projects around the world with 54 servers and 280 fetches (unique domains). The top node of the world allows for searches on WPS information and on URCs. The next version of the Digger code will be available at the beginning of February and will provide language support, Unicode support, searches for CDR blocks and security in the form of Kerberos 4 between both servers and servers and clients. An NSF project at the University of California is currently looking at management tools and scalability. The expectation is that this project will handle 0.5 million entries. Further information can be obtained from :

<http://www.bunyip.com/products/digger>

2. Papers submitted

The Building a Directory Service in the US Draft will be progressed as an informational RFC. A new ID with a few minor changes to the text about '93 Access Control and corrections will be submitted shortly. After discussion and a show of hands the IDS group decided that the directory paper by Peter Jurg and Erik Huizer should be progressed as a BCP RFC. Harald will take up this matter with the IESG.

3. X.500 Catalogue

The authors are currently waiting for implementor descriptions, a deadline of 31st December was agreed for implementor input with the resultant Draft being out by 15th January 1996. There was some discussion about what to do with the existing but out of date information in the catalogue, the suggestion was that this should be taken out.

4. WHOIS++ Catalogue

This has not been updated since Stockholm but implementors should send their information to Patrik Faltstrom by 31st December then the Draft will be available by 15th January 1996.

5. CCSO Nameserver (PH) Architecture

The aim of this draft by Roland Hedberg, Steve Dorner and Paul Pomes is to get a more rigorous description of Ph and no functionality will be added that would break the currently running software. A BNF description of the protocol is being added. Creating new attributes within Ph is currently open, Joann Ordille found over 900 attributes in her research - many of these were language differences and interpretation differences. A minimal subset of attributes is being put together with stricter descriptions. A first draft of this paper has been sent to the registry and the next version will be circulated within weeks.

6. Ph Directory Server Creation and Support

Roland Hedberg, Joann Ordille and John Noerenberg will submit a Draft by 1st March 1996 which will give preferred practice and will include a preferred schema. The Nomenclator report has information on schema attribute distribution :

<http://netlib.att.com/netlib/att/cd/home/nomen/nomenclator.html>

Alternatively email joann@research.att.com

More information on Ph can be obtained from the following listserv :

info-ph-dev@listserv.cso.uiuc.edu

The final decision about whether the Ph work will be FYI or Standards track will be made at the LA IETF. The current feeling is towards Standards track because of the perceived importance of having well-documented directory services in order to obtain an integrated directory service.

There was some discussion about using the Web as a directory service and the problem of unstructured data.

7. X.500 Root Naming context

David Chadwick gave a brief overview of his paper. He is proposing a phased solution to the problem of managing the Root namespace which uses DISP rather than DOP initially. Nexor will implement these changes next year and they will be piloted by Nameflow.

The decision on how to progress this paper was to discuss it until 31st December then move it to Experimental. Implementors would then be sought and Harald would approach SC21 with the results. If they still were not interested then it would be decided whether to move it on to the IETF Standards track.

8. Internet White Pages Service

There was a lot of debate on what exactly user requirements meant and the resultant agreement was that discussion would be moved to the mailing list. This will be kept focussed by using threads in subject lines. The chairs will send out a Table of Contents within a week to get agreement on a structure. The authors would then rearrange their test in to the new structure then it would be reviewed section by section on the list.

9. Schema Requirements

Tony Genovese went through the document structure and the following items are to be discussed further on the mailing list :

Structured fields

Certificate storage

Tony would like more input from the WG along with some examples. Harald will provide input on character sets.

10. Privacy Requirements

Barbara Jennings will send a paper she has written on this subject to the mailing list and the Group will

decide if they feel this is a relevant topic to be pursued.

11. Quality of Directory Service

This discussion was prompted by a NameFLOW paper on quality of directory services. It was decided that this information should be used as input to the user requirements draft.

12. Charter Revision

Document timescales and revisions will be incorporated into the Charter. It was also felt that the Charter ought to make the directory experience of the Group more visible to the rest of the IETF in the light of current discussions in the DNS and PKIX areas.

13. Schema Task Force

Pressure of work has forced a number of the original group to relinquish their responsibilities to this area. It was agreed that this was an important area that had to happen so Tim Howes will try and reform the Task Force with new people.

APPENDIX 7

NOMENCLATOR

Distributed Search and Retrieval from CCSO Servers

NOMENCLATOR INTERNET PILOT REPORT

The Nomenclator Internet Pilot is currently integrating many of the publicly available CCSO servers around the world. Each CCSO server has a database schema that is tailored to the needs of the organization that owns it. Nomenclator can integrate the different database schema at these servers and provide fast cross-server searches for locating people on the Internet.

More than 40 CCSO administrators have registered their servers as part of the Nomenclator Internet Pilot. We have used the registrations to develop and test a new type of data integration system called "Data Integration By Example." A paper on this technique has been published, and is also available through the Nomenclator home page as at the end of this message. We are currently doing a final testing of the full system, and expect to field a service on the Internet in February. The service provides a Web interface for queries and a query protocol for those who wish to create their own interface. We will also be distributing the query processing software, so you can run the query processor at your site to benefit from the increased speed of local caching and lower local load. An informational RFC on the Nomenclator protocol and the Nomenclator pilot will be available after the system goes into service.

Nomenclator uses distributed indexing techniques to provide fast cross-server searches. It also supports extensive data and meta-data caching to scale the system to Internet size.

One of the best things about the system is that a CCSO site can incorporate its server into Nomenclator without having to change it. All we need is for the administrator of a CCSO server to provide some basic

information about the server. We use that information to make queries faster by constraining searches to those servers where an answer is likely to be found. We also use the information to make queries easier to express and understand by mapping the data at each server to our global schema.

When Nomenclator constrains the search for a query answer, it screens out irrelevant queries from ever reaching a CCSO server. When Nomenclator finds an answer in its cache, it screens out redundant queries from reaching a CCSO server. Each server becomes easier to find and use without experiencing the high loads caused by exhaustive and redundant searches.

Nomenclator is described in detail through its home page at:

<http://www.cs.att.com/csrc/nomen/nomenclator.html> -or-

<http://netlib.att.com/netlib/att/cs/home/nomen/nomenclator.html>

In particular, detailed technical papers about the system are available at

<http://www.cs.att.com/csrc/ordille2.html> -or-

<http://netlib.att.com/netlib/att/cs/home/ordille2.html>

The paper "An Experiment in Integrating Internet Information Sources" describes our Data Integration By Example techniques. The paper "Information Gathering and Distribution in Nomenclator" provides a very short overview of the system. The Ph.D. thesis "Descriptive Name Services for Large Internets" provides the most complete technical description. A shorter, but reasonably complete, technical description is available in "Distributed Active Catalogs and Meta-Data Caching in Descriptive Name Services."

New registrations are always welcome! Registration takes approximately one hour. We will be happy to receive comments about the the registration process or other aspects of the system.

The Nomenclator Internet Pilot is sponsored by InterNIC Directory and Database Services and AT&T Bell Laboratories Computing Science Research Center.

Submitted to the IETF-IDS on December 1, 1995.

Joann J. Ordille	
AT&T Bell Laboratories	Tel 1 (908) 582-7114
Computing Science Research Center	Fax (908) 582-5857
600 Mountain Ave., Rm. 2C-301	E-mail joann@research.att.com
Murray Hill, NJ	
USA 07974-0636	

<http://www.cs.att.com/csrc/ordille.html>

Minutes of the FIND WG, 34th IETF, Dallas

The FIND WG met for the first time at the 34th IETF. Patrik Faltstrom chaired the meeting. He first reviewed the charter. It is pointed out that the email address of the Working Group is "find@bunyip.com", and nothing else.

The Charter gives the WG goal as defining one and only one common indexing protocol which all directory services can use when passing indexing information. Patrik admitted that so far this work has been aimed toward WHOIS++, but that he is depending on the group for help in making it work across directory protocols. Currently there are 2 drafts which came out of the WNILS WG: one on the Common Indexing Protocol by Chris Weider and the other on the WHOIS++ mesh by Patrik. Patrik intends that the second version will include LDAP and PH.

The way directory information is indexed in the CIP is for each leaf node to supply the information to the indexing server(centroid). When the indexing server gets a query it will be able to prune the branches where there will be no information. (Note that the examples are in WHOIS++) The leaf server sends the Data-Changed command:

```
# DATA-CHANGED
```

```
Version-number:
```

```
Time-of-latest-centroid-change:
```

```
Time of message-generation:
```

```
Server-handle:
```

```
Host-Name:
```

```
Host-Port:
```

```
Best-time-to-poll:
```

```
Authentication-type:
```

```
Authentication-data:
```

```
# END
```

The centroid uses the Best-time-to-poll value to send a poll command:

```
# POLL
```

```
Version-number:
```

```
Type-of-poll:
```

Poll-scope:

Start-time:

End-time:

Template:

Field:

Server-handle:

Host-Name:

Host-Port:

Hierarchy:

Description:

Authentication-type:

Authentication-data:

END

The polled machine sends back the Centroid-changed response:

CENTROID-CHANGES

Version-number:

Start-time:

End-time:

Server-handle:

Case-sensitive:

Authentication-type:

Authentication-data:

Compression-type:

Size-of-compressed-data:

Operation:

```
# BEGIN TEMPLATE
```

Template:

Any-field:

```
# BEGIN FIELD
```

Field:

Data:

```
# END FIELD
```

```
# END TEMPLATE
```

```
# END CENTROID-CHANGES
```

Both the template and field are repeatable.

Today the only transfer is on the whole centroid, it is case insensitive, is the 8879-1 character set, and the tokenization algorithm is white space and @.

More information about the CIP is available at:

<http://www.bunyip.com/products/digger>

The question was asked why use this when X500 has replication? The answer is that it is a base for the future. X500 doesn't offer indexing, nor does it provide a common indexing for all protocols. This model is also used for URN to URC resolution at Georgia Tech, and the model may allow for Web indexing. Chris Weider pointed out that it will allow for the 1,000 flowers blooming, a term which refers to the multiplicity of directory protocols becoming available.

Patrik was asked about things not WHOIS++ and he replied that he does not believe there will be any problem handling the indexing information.

Patrik was also asked is the WG should do a survey of indexing schemes and Patrik replied he was looking for volunteers to do so.

There was a small semantic discussion on whether it was an indexing protocol or whether it was exchanging data to create an index. Patrik would like to have a common format for the index if possible. Each directory service would pass its index to the centroid, which would index that index. And in fact, the index would be indexed at each level of the tree. Some of the issues to study will be the trade-offs of the number of levels and the size of the indexes. A lot of factors are involved: data, reduction of indexes, geography.

Tim Howes of U Michigan gave a presentation on an program he's written called centipede. The centipede connects to a directory over LDAP which tells it what information to produce for the centroid index. It is produced, and centipede then connects to the target with the references and uses LDAP to install the index in the entry. It generates distinct values (whole names rather than tokens) which it passes up the tree. The large index allows more precise searches and pruning. Tim gave the following URL for more information:

<http://www.umich.edu/~rsug/ldap>

CIP has the advantage of you knowing who is indexing you, while centipede does not.

Chris W. reminded the group that all of this was experimental and wanted the group to think about what sorts of indexing information would be useful.

The group identified the following issues:

- Character sets
- Tokenization algorithm
- Legal issues
- How to specify for partial centroids
- New server-to-ask records
- Schema translations
- Query result
- Protocol issues
- Security
- Replication
- Dealing with replicated data
- Polling cycle detection

The group focused on what might be simple. These issues might be:

- Common format and schema translation
- Overall model
- Given a name of a company, return a domain name
- Index data stored in WHOIS RWHOIS WHOIS++ and X500.

The group agreed that the plan of attack should be:

- 1) Overall model
- 2) Schema translation

Both Joann Ordille and Roland Hedburg have some experience with schema translations that will be useful to the group.

It was suggested that we would need a registry of schema with descriptions and capabilities.

The group also asked what happens when the search result lives on WHOIS++ but the client only speaks LDAP. Proxies were suggested as one solution. Another solution would be to return URLs which contained queries which could be handed to a server.

The group then elected to defer engineering discussions to the list, and Patrik adjourned the meeting.

Participants of the FIND WG meeting in dallas

...[list deleted by VB] ...