

NameFLOW-Paradise

Quarterly Service Report April - June 1996

TABLE OF CONTENTS

Introduction

Operations

1. Operations/helpdesk
2. Outages
3. Issues
4. Statistics
5. NameFLOW-Paradise Managers Meeting

Liaison

EEMA
IETF
INET'96
Finnish Directory Forum
Unisource
EuroSInet

Information

Information servers
Reports
Publicity/promotion

Appendices

APPENDIX 1 Helpdesk summary
APPENDIX 2 World Root DSA and LDAP summary statistics
APPENDIX 3 Public DUA summary statistics
APPENDIX 4 WWW/FTP/Gopher summary statistics
APPENDIX 5 Draft Minutes EEMA Directory Committee
APPENDIX 6 Draft Minutes IETF IDS
APPENDIX 7 Draft Minutes IETF ASID
APPENDIX 8 Draft Minutes IETF FIND

Introduction

This Quarterly Report reflects the NameFLOW activities and operations for the months April, May and June 1996. 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: nameflow@dante.org.uk

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

Operations

Introduction

This chapter summarises operations and related activities of the NameFLOW- Paradise Directory service for the three months April, May and June 1996.

1. Operations/Helpdesk

A new country, Malaysia (c=MY), together with its FLDSA Rimau (Malay for tiger), is in the process of joining the Directory. The entry for cn=Rimau has been added to the top level of the DIT, but c=MY has not yet been added, as the connection to Rimau is affected by a low-level network problem. An IP router in Malaysia appears to be behaving strangely, and this is under investigation in MY.

A problem with the FLDSA cn=Margay was pointed out to the managers for c=DE. The DSA's local copy of its own entry had lost the mandatory quipuVersion attribute, which prevented Giant Tortoise from shadowing the entry.

In May, helpdesk staff assisted DANTE in the preparation of the latest NameFLOW- Paradise publicity brochure, for release at the EEMA exhibition in Brussels.

In June, a count was prepared of organisations and DSAs in the Directory Information Tree (DIT).

It had been intended for quite some time that the FTP server would hold mirror copies of software and documentation related to directories, both X.500 and others, but plans were postponed due to a shortage of disk space. A new 4Gb disk was installed in May, and some mirroring agreements have been established.

Draft versions of future ISO Directory documents are mirrored from a Groupe Bull server in the USA. A request to establish a mirror of whois++ software from Bunyip was refused, due to the imminent release of a new version of Digger; Bunyip wish to monitor downloads. Various X.500-related items are mirrored from a server at SURFnet in the Netherlands, which itself mirrors from others. This is an interim measure prior to arranging mirror agreements directly with the relevant master servers.

2. Outages

Significant outages of service totalled approximately 12 hours in April; 16 hours in May; and 11 hours in June. Details of outages are available in the monthly reports.

3. Issues

An office of the Ministry of Research and Technology in Indonesia (c=ID) has expressed an interest in joining the NameFLOW-Paradise Directory, and running the FLDSA for c=ID. However they are not as yet running a DSA, so there is no short- term operational involvement.

The University of Michigan released version 3.3 of their LDAP implementation at the end of April. Currently the central LDAPD is version 3.2. Possibly it is time to upgrade to the latest version.

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 info-server, under:

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

The major change in statistics this quarter is that the figures for the DSA Ocellated Turkey, and for the mail info-server, are no longer shown. Both servers were closed down in February.

Use of the LDAP server continues to increase dramatically. All use attributed to the .uk domain is from a machine at DANTE, and is presumably due to the WWW/X.500 gateway. It is expected that use will drop over the summer months due to the long academic vacation. However, the figures for June are much lower than anticipated, due to a change in the configuration of the gateway. It was accidentally modified by DANTE to use another LDAP server, for most of June.

The figure for local operations for Giant Tortoise in June is affected by the "count of the DIT" performed on 9th/10th June, causing an unusually large number of search operations.

Some of the increase in the use of the FTP server in June is an effect of secondary mirroring. As this server mirrors from others, so other servers in turn mirror from this one; in particular, a server at Macquarie University in Australia.

Although attempted use of the withdrawn public-access DUA falls steadily, there is a sharp decrease from the previous report. It is assumed that this results from the change to the "home" Gopher menu at the University of Minnesota, mentioned in the monthly report for March.

5. NameFLOW-Paradise Managers meeting

The Managers meeting was co-located with the EEMA annual conference, and although it was a short meeting it was rather well attended (16 persons). There were only a few people who attended both the EEMA conference/exhibition and the managers meetings. It is questionable whether future managers meeting will be co-located with EEMA events as most people reacted quite indifferently to this approach.

To give you a rough impression what has been discussed, the minutes from the meeting are included here. This information can also be found on the web

<http://www.dante.net/np/meeting.html>

Minutes 11/06/96 (vb(96)028)

Country File

Performance probe from SURFnet, will be made available when finished.

Dan Net

Public directory service in Denmark.

Co-operation between e-mail service providers and government. One of the agreements was to establish an X.500 directory. First a paper directory was established (72000 entries in 95).

Digital DSA (X.500 93), new user interfaces for different types of user systems.

- query by email (response time 10 seconds)
- WWW (own gateway) & gopher service, finger,
- whois,CSO/Ph

Information represented: organisation, ou (one level),organisational persons, email addresses, EDI information, Web addresses.

Extensive use, 8000 web transactions/week (90% of total)

Problems with difference of input sources, maintenance and update problems.

- WWW update service. Security based on email address, message is sent to owner of the entry that is requested to change. You can change single entry or whole organisation.
- Email update service.

Central exchange station for addresses

- data manager function: for old entries it sends out messages to owners for updates.

Now only one DSA, later more distributed. Naming problems are encountered.

- ~ 50000 entries, ~4500 orgs
- software (all based on ldap, email security based on XAPI)
 - WWW gateway
 - mail gw
 - administrative tools

Root context

Replace quipu model with standard X.500 (93) failed.

Following should be done:

- add access control information in DISP
- enhance DISP to allow one level searches in shadow DSAs
- June 18 DANTE meeting with manufacturers to implement DISP enhancements (DISP version 2)
- add text to describe how HOBs carry context prefix
- information

EWOS produces shadowing profiles now. Replication of single entry is missing in the profile. May already be provided by chopping.

SOLO

SOLO is used now in a corporate (Intranet!) directory , a product of Telis Systems & Communications
Performs lookups

1. search + additional info: pointers to managers, secretaries
2. browse

SOLO is used since it is lightweight and has good performance
Integration with other software (email software, WWW, phone)

SOLO

- TCP/IP & text based
- X.500 datamodel and naming & centroids from whois++
- good integration with x500/ldap/www
- potential for whois++/CIP/centroids
- no data management protocol, just simple lookups
- well suited for Intranets
- Internet directories: ???

NP-93

Presentation of the result of the first test phase. Full details available via <http://www.dante.net/np/93pilot.html>

T.61 in DNs

Polish approach to internationalisation of the directory, i.e. Polish characters for Polish users in the

Directory.

Multiple attribute values don't work because you can't decide which one is meant for Polish users

Solution:

- additional object classes and attributes with T.61 and a polishRDN attribute.
- modified WWW interface
- mapping of DNs to polishRDN, these are cached in the WW interface
- different charsets are provided

Idea for general approach by using local entries in each country

"localizedCommonName" and "localizedRDN"), and localized

Overlap with MAITS project of the EC. MAITS tries to solve charsets problems in X.500 and other applications (like WWW). MAITS tries to deliver a set of API for mappings and a mechanism for presentation in ASCII. Further translation is supported. Nexor and David Chadwick are involved in MAITS. Merge of the two approaches possible?

Indexing

University of Brussels

- Whois++ centroids in X.500 (not CIP based)
- Three new object classes
 - person centroid
 - organizationalunit centroid
 - organization centroid
- collects
 - person -> surname
 - organizational unit -> ou
 - organization -> o
 - can be adjusted
- special WWW gw to ldap for searching
- centroid building software (ldap based)
- URLs: <http://sun7.iihe.ac.be/centroid/index.html>
<http://sun7.iihe.ac.be/centroid/build.html>

David's paper -> few remarks

It was decided to form an interest group to discuss integration of three approaches (the ones from Nils, Roland & David).

Next meeting

26 & 27 November 1996

Liaison

EEMA

The biggest event of the EEMA is the Annual Conference and Exhibition, this year hosted in Brussels, Belgium. To promote "academic applications" and Directories in particular DANTE had a stand with various sorts of brochures and other publicity material, such as DANTE in PRINT#19 IndeX.500 by David Chadwick. There was quite some interest in what we (academic and R&D community) do and some people enquired what was needed to get connected to NameFLOW-Paradise, as a user or as an information provider.

This year one of the NameFLOW-Paradise meetings was co-located with the EEMA annual conference. The other meeting of Interest is the EEMA Directory Committee.

EEMA-Demonstration

DANTE was well represented with a demonstration booth, where we were able as one of the few exhibitors to show a working public directory (NameFLOW-Paradise). Besides the usual web browsers and other interfaces, the most recent Indexed servers of Roland Hedberg (UMEA University, Sweden) and Nils Meulemans (Free University of Brussels, Belgium) were demonstrated with quite some success. On average people could not believe that the Directory was so fast (and especially over the Internet). A quick overview of what was demonstrated at the conference is still available on the web <http://www.dante.net/np/brussels.html>

EEMA-Committee (New project manager) So far the cooperation between DANTE and the EEMA has been fairly limited, due to the way the EEMA is organised. The way committees are currently organised is significantly changing from what could be seen as a "mainly consultants" organisation to a more open forum where work is contributed by the participants, similar to the IETF approach.

To improve the relation between the EEMA and DANTE (and R&D environment) the position of project manager was offered to Vincent Berkhout , which he has accepted (still subject to contract). This way we hope to increase our influence in the more commercial side of Directories and keep a close eye on developments. The input from the NameFLOW-Paradise community could be a valuable asset for the EEMA as many of the relevant Directory experts are closely working with, or even within, the NameFLOW-Paradise service. As a first objective as project manager for the Directory Committee Vincent will try to improve communication between its members.

IETF

As a general observation it seems that the battle between Whois++ and LDAP seems to have been fought. With support from major vendors it looks like LDAP is going to win, but still leaving many questions to be answered. As LDAP is only an access protocol it does not answer the question of how a truly distributed Directory is to be implemented. As a result of these recent changes the Whois++ community is now focussing on the FIND working group where the Whois++ concept can very well be used. Below you'll find a bit more about the working groups ...

IETF-IDS (see Appendix 6)

An average meeting, but with some intensive discussions. The discussions seem to focus on the trivial things and the group does not really discuss the protocol issues. On the whole it seems to become a "liaisons report" meeting. However, developments with Ph are going strong and it is expected to go up for proposed standard. The work from Martin Hamilton on DNS aliasing and "finding stuff" is useful, but is a fix where the use of SRV records in DNS is perceived as the solution.

IETF-ASID (See Appendix 7)

The main thing that is happening within the ASID group is the development of the LDAP Version 3. This version of LDAP is coming very close to full DAP and it was suggested to call the third version Internet DAP. There was some discussion that LDAP could derive from the X.500 standard where needed, but people thought this was a bad idea. After the working group meeting a group of three (Tim Howes, Mark Wahl and Steve Kille) discussed the issues raised during the meeting and have published a set of resolutions and rationales on the ASID mailing list.

In addition the group works on Whois++, labeledURI, PGP keys in X.500, directory MIME types (a framework document and the versit profile) and on the exact URL format to distinguish whois, rwhois and whois++ queries.

IETF-FIND (See Appendix 8)

The FIND group has changed its approach significantly and the bias to Whois++ is now much more balanced opening the way for a true Common Indexing Protocol. It is very good to see that the group is now seriously looking at integrating Whois++, LDAP and CCSO. The current proposed architecture looks promising, but there is still quite some confusion in the group itself. The presentation on "Scaling of the CIP" was most interesting and showed the characteristics of centroids and how they scale.

INET'96 - Session on Directory Services

The IETF was co-located with INET'96. There was one interesting "Panel Discussion" hosted by Erik Huizer with David Chadwick, Tim Howes, Tony Genovese and Patrik Falstrom. After each panel member gave a presentation on their personal view on Internet Directories the audience was allowed to ask questions. A good session and a good indicator where the Directory problems lie although the session had a commercial flavour.

Finnish Directory Forum (Trip note - 10 April, 1996)

To promote European Directories and to show what NameFLOW-Paradise does for individual countries, two presentations were given during a meeting where the Finnish Directory Forum was launched. The Finnish Directory Forum is the cooperation between the two national telephone operators using X.500, EUnet using Whois++ and FUNET, the Finnish national research network, currently supporting both. The set up of the FUNET is a moderate but effective one: the access is mainly via the Web using LDAP gateway(s) and operating a single DSA. The data are centrally stored and there will be a given set of interfaces.

During the meeting the greatest concern of the audience was security and access control on private information. As a general remark the actual synchronisation of the information will need some further attention. FUNET is definitely leading the way with their cooperative directory by multiple service providers.

Unisource (Trip note - 17 June, 1996)

On 17 June an informal meeting took place with Unisource. During this meeting Unisource expressed interest in participating in the further development of NameFLOW-Paradise. As yet no concrete steps have been taken as a result of the meeting.

EuroSInet

(Press release, 23 may 1996)

EuroSInet Interoperability Testing Workshop News - Yet another successful event

EuroSInet, the European Interoperability Testing Association, held its 7th workshop event from May 13th until May 17th at NCR in Copenhagen, Denmark.

The nine members listed below were present at the event verifying the interoperability of their X.400 electronic mail and X.500 directory products.

Vendor	Testing
Data Connection	X.500'93
Digital Equipment Corporation	X.500'93
ICL	X.400'88 & X.500'93
ISOCOR	X.500'93
ISODE Consortium Limited	X.500'93
NCR	X.400'88 & X.500'93
Net-Tel	X.400'88
NEXOR	X.400'88 & X.500'93
Siemens Nixdorf	X.400'88

A notable feature of the event was the number of X.400 suppliers who are now able to offer Message Stores and P7 Remote User Agents.

Participants at the event were able to use EuroSInet's newly developed X.500'93 test suite. In addition to basic directory functionality this suite includes tests to verify Access Control and Replication (Shadowing), both of which are of vital importance to users.

The importance of the X.500 Directory standard is further emphasised by recent press announcements by over 40 major vendors including, Netscape, Microsoft, AT&T and others pledging support for the Lightweight Directory Access Protocol (LDAP). LDAP complements the fully featured Directory Access Protocol (DAP) of the X.500 standard.

EuroSInet members are currently developing an LDAP interoperability test suite which will be available by the end of June 1996.

Building on the success of the workshops to date, EuroSInet plans further workshops in the Autumn of 1996 and Spring 1997.

Clive Betteridge
Chairman, EuroSInet
21st May 1996

[Issued by EuroSInet Secretariat, c/o Apertus, 10 The Cedars, Tilehurst, Reading, UK]

Information

Information servers

As part of the information service of NameFLOW-Paradise DANTE operates several servers. There are the 'historical' PARADISE information servers, such as ftp and gopher, 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:

1st Quarter 1996	http://www.dante.net/np/report/qr/96Q1.html
2nd Quarter 1996	http://www.dante.net/np/report/qr/96Q2.html
April 1996	http://www.dante.net/np/report/mr/mr9604.html
May 1996	http://www.dante.net/np/report/mr/mr9605.html
June 1996	http://www.dante.net/np/report/mr/mr9606.html

Publicity/promotion

DANTE was present with a stand at the EEMA Annual Conference in Brussels, 10- 14 June 1996. Its purpose was two fold: firstly to increase the visibility of DANTE's activities in the directory area and secondly to increase awareness of DANTE as a player in a pan-European networking context.

For the occasion a special set of web pages reflecting the latest developments in the NameFLOW-Paradise service was created and the new brochure was distributed.

APPENDIX 1 Helpdesk summary for Apr/May/June 1996

Country	Number of queries				
Full Name	ISO Code	April	May	June	Quarter
Australia	AU	1	1	-	2
Canada	CA	-	-	1	1
Germany	DE	1	-	-	1
Finland	FI	2	-	-	2

United Kingdom	GB	3	1	1	5
(Indonesia)	ID*	1	1	1	3
India	IN	-	1	-	1
(Malaysia)	MY*	-	-	2	2
Netherlands	NL	1	-	1	2
Portugal	PT	-	1	-	1
Russia	RU	1	-	-	1
Sweden	SE	1	-	-	1
United States	US	5	3	-	8
<hr/>					
Total Requests		16	8	6	30

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

APPENDIX 2 World Root DSA and LDAP summary statistics for Apr/May/Jun 1996

Summary of calls to DSA Giant Tortoise

From 0:05:12 on 31 March to 0:04:43 on 30 June

No. of binds	April	May	June	Quarter
Local	197	247	266	710
Remote	6042	5865	5253	17160
<hr/>				
Total	6239	6112	5519	17870

No. of operations	April	May	June	Quarter
Local	5	27	7275*	7307
Remote	46318	47877	47538	141733
<hr/>				
Total	46323	47904	54813	149040

(* higher than normal due to a count of the DIT being performed)

System usage (calls received)	April	May	June	Quarter
-------------------------------	-------	-----	------	---------

Binds by Directory technicians	1837	1646	1430	4913
Reads of DSA entries	24	34	26	84
Other ops on DSA entries	0	5	23	28
Getedb operations (inc slices)	46151	47517	46997	140665
Spot shadows	10	13	19	42

Total	48022	49215	48495	145732
-------	-------	-------	-------	--------

LDAP usage

LDAP usage from Apr 1 1996 to Jun 30 1996

	April	May	June	Quarter
Connections	5742	16763	8152*	30657
Total connect time (seconds)	119961	477470	100009*	697440

(697440 seconds is 193 hrs 44 mins 0 secs)

(* lower than expected due to configuration change in WWW/X.500 gateway)

APPENDIX 3 Public DUA summary statistics for Apr/May/Jun 1996

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

Note: DUA access was withdrawn during February 1996, so these figures - reflect attempted rather than actual use.

Network	April	May	June	Quarter
Internet	3053	1873	1171	6097
UK academic X.25 (JANET)	0	14	4	18
EuropaNET X.25	0	0	1	1
Public X.25	9	11	5	25

Total	3062	1898	1181	6141
-------	------	------	------	------

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

Domain	April	May	June	Quarter
edu	1390	523	320	2233
unresolved	437	310	203	950
uk	302	326	177	805
com	183	132	79	394
net	139	119	59	317
nl	65	41	39	145
us	63	38	30	131
dz	41	51	31	123
org	42	30*	40	112
ca	52	29*	23	104
Total	2714	1599	1001	5314

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

APPENDIX 4 WWW/FTP/Gopher summary statistics for Apr/May/June 1996

WWW server

	April	May	June	Quarter
Unique hosts	780	1219	1252	3043
Number of HTML requests	1885	2201	2207	6293
Number of non-HTML requests	503	296	247	1046
Number of malformed requests	57	37	57	151
Total number of all requests/errors	2448	2534	2511	7493
Total number of Kbytes requested	76757	49384	38759	165900
Average requests/day	81.6	81.9	82.7	82
Kbytes/day	2559	1596	1277	1810

FTP server

TOTALS FOR SUMMARY PERIOD Mon Apr 1 1996 TO Sun Jun 30 1996

Note: the June figures are unusually large, partly due to secondary mirroring.

	April	May	June	Quarter
Files Transmitted	96	104	367	567
Bytes Transmitted	30527360	21997924	102171133	154696417
Average Files Daily	4	5	13	7
Average Kbytes Daily	1387	1157	3648	2064

Gopher server

Gopher usage from Mon Apr 1 1996 to Sun Jun 30 1996

	April	May	June	Quarter
Total connections	17	24	17	58
Total files retrieved	2	35	27	64

APPENDIX 5

EEMA Directory Committee Draft Minutes

Permeke Room, Sheraton Hotel & Towers

Place Rogier2, B-1210 Brussels

Thursday, 13 June 1996 - 12.30-16:15

The meeting started with a lunch offered by EEMA.

David Goodman (Committee Chairman) welcomed the delegates .

1 Adoption of the Agenda

The Agenda was issued to delegates who introduced themselves and gave a resume of there interest in the Directory Committee and the topics they would be interested in discussing as follows:(Note that this is not an exhaustive list)

security and security needs directories; organisational use of X.509; NADF liaison and promotion of the challenge -97 security event; the most recent developments, especially in LDAP; NameFLOW-Paradise; Naming Structure for Directory Information Tree; TOPOL (Top Level Naming in Europe; a cooperative environment; recommendations for global connectivity (root

DSA); Directory Service Operations; ADDMD setting up PRDMD; business cases to help start a service; promotion of directory deployment; networking; cooperation regarding a DG3 directory project called Euroview; where Europe fits in to American perspective; after this round a short introduction by Vincent Berkhout explaining that he would like to improve communication for the DC and wanted to push projects and their champions to continue the valuable output so far.

Additional agenda item: WEMA and the challenge-97 proposal

The objective of the committee should be to drive the market and have a positive effect on electronic commerce. It should provide proof of concepts, interworking demonstrations and be involved in top level naming.

2 Minutes of the Last Meeting

Minutes from previous meeting were distributed.

3 Liaisons/Education

Agenda item was skipped as the available time for this meeting was limited.

4 WEMA Challenge-97

The idea behind Challenge-97 is to have a security demonstration under a global banner, where the challenge is to "send a secure message from continent to continent without violating any laws". The event is initiated by the Japanese EMA (JEMA) who are the Security champions but as security needs Directories they requested "HELP" from the EEMA-DC(EEMA is the champion of Directories). Alexis Bor, the NADF chair and sub- chair from EMA-DC, explained the goal. The first kick-off meeting to discuss the organisation of the Directory Challenge-97 event will be held in Tampa, US on 15- 16 July 1996

5 NADF (used to be EDF Naming)

The use of the abbreviation EDF for European Directory Forum is rather loaded within Europe as it is also the name of a sub committee of the European Numbering Forum (ENF). The suggestion was made to change the use of the abbreviation NADF where the word 'American' would be replaced with 'Atlantic'. The name of the connected European and American forum would then be "North Atlantic Directory Forum". The committee indicated that they were interested in participation with the NADF and will set up a subcommittee of the committee to cooperate/liaise with the American part of the NADF. An inconclusive discussion who would be representing the committee followed, two volunteers (David Goodman and Nick Laszlo) were appointed as they will be attending the next meeting anyway. They will be a focal point between committee and the NADF- west. Others who showed interest in Tampa meetings were Colin Robbins, Roger Mizumori, Bernard Tardieu and Keith Richardson.

6 Challenge-97

The challenge-97 will be driven by a group of four. The challenge-88 was mainly driven by Joanne Ghahremani from BTNA but she has been allocated to a different project within the company and will not

be able to spend a same amount of time on challenge-97 and was looking for some financial compensation to supplement her voluntary effort. Back to the group of four, Colin Robbins offered to champion the European effort and set up smaller team per continent. US customs were interested in participation as they see the challenge-97 as a proof of a concept. to get more money in 1998. The US Customs knew that the Australian and French customs would also be interested in participation and possibly the European Commission.

[Need for security meeting to look at the legal side of things?]

[What is driving directories? electronic commerce and Vendor support]

[Two drivers of directories will be world wide banking and crime prevention]

[and electronic business environment, what is important for the Directory Committee and what for SPLC?]

Concern was expressed about the overlap of committees ie Directories, SPLC and the User Committee and what they would do in the challenge-97 and which "committee" would get the credit. The goal of the challenge is point-to-point security and it was decided that what is in-between is relevant. The disadvantage of Directories is that they are not on the foreground and are not really seen by users. For the challenge: setting up a continent to continent messaging connection there is a need for a (alternative) Certification Authority. A member of the SPLC said they have a project how to set up a Certification Authority and requested that the project be validated by members of the Directory Committee.

7 Project Reports

7.1 Top Level Naming in Europe

First phase is almost ready and the document is in final draft. There was a workshop on Monday to kick off phase two where the actual recommendations are defined. The workshop was attended by 50 people. The goal of the document seems to be shifting, and the contract may need to be revised in agreement with the authors. The current development of LDAP should be included to take the Internet naming into account. This will mean the analysis has to expand. Colin Robbins wants to add his "conspiracy theory" where SLAPD (Stand Alone LDAP servers) take over and be the most important Directory and all work will be in vain.

7.2 Guidelines for Corporate Directory Deployment

The project is finished and the document is ready to be printed as an EEMA black book. Unfortunately the document was not in time to have it printed and presented at this meeting. The MAFF people and project champions have decided that the fourth business case planned to be appended could be dropped and was accepted by the committee. It was proposed to put electronic copies on the EEMA web server. EEMA should decide if to sell the document or use the document as a demonstration/publicity tool. (Note: EEMA policy of storing document on their web server and possible access control is not clear to the committee.)

7.3 X.500 Product Guide

The second version of the X.500 product guide describes the directory (X.500) products of forty vendors. It is significantly more elaborate than the first version with twenty vendors. The objective of the document is to counter misinformation and show that directory products exist. The document is different from the more general WEMA product guide as the information is more objective and should give a fair overview to compare products and vendors.

[The issue was raised that the document should be made publicly available if the EEMA believes the number of sold copies would be limited, see above. Discussion of all free catalogues available and how many copies would be sold. An alternative, storing web pointers to the vendors home pages and products.

Vee Baker suggestion was that the document is member information anyway]

Action -David: discuss web/black book policy with EEMA.

Action - Vincent: for EEMA members an electronic copy in PDF format on the web server.

7.4 Powered by X.500 (Theme of Brussels exhibition)

The demonstration of X.500 products was not as good as last year. The EuroSInet workshop in Copenhagen was considered more useful. Most people did not get the message "powered by X.500". The intention was to show operational directories using attendees of the committees and visitors of exhibition. Access should be provided via a public network, not a special one as last year. The question was asked which exhibitors could demonstrate a public directory only two (three) exhibitors provided a public directory. In a previous meeting it was decided not to have the interworking demonstration but make this exhibition a pure marketing event.

Exhibitors would showed applications using X.500 would say "powered by X.500" similar to the "Intel Inside" campaign. The volunteers who should have been organising the marketing event were not attending this meeting.

8 New Projects

8.1 Next Years Exhibition

The 97 Annual Conference will focus on security (challenge-97). For next year's security event the EEMA should look for a sponsor of smartcards for all the visitors to the exhibition. Security is "showable" but underlying directory is taken for granted. There should be a clear focus and a single objective for next year.

8.2 Introducing LDAP

A paper introducing and explaining LDAP was proposed by Colin Robbins. It will be a document similar to a FAQ, an overview of Frequently Asked Questions. Please send questions and if possible answers to cjr@nexor.co.uk and Colin will write a document.

8.3 Business Cases

There was an interest in getting more business cases. Larry Lee has a collection and could be used as a base. First look at his work. The difference between "case study" and "business case" was emphasised. There was no volunteer to do the work - Chris Ray maybe. David Goodman will champion the project as he proposed it and was interested in the work.

9 Any Other Business

Some discussion on OIDs, as a continuation of the TOPOL workshop on Monday. A short introduction to MAITS was given by Colin Robbins. The project is now in the phase were they will work on the user requirements document where the Directory requirements will be incorporated.

10 Date and Time of Next Meeting

The 26-27 September in Moscow is the first option but a very limited number of committee members will be there. Therefore the next meeting will between 15-17 October in London collocated with Electronic Commerce Ô96

APPENDIX 6

Draft IDS Minutes

Montreal IETF

9am to 11:30am June 26th 1996

1. Previous meeting's Minutes and current Agenda were agreed.

2. Whois++ Projects Presentation

The Internet Whois++ Project NSF Testbed is currently running Digger 2.0 and moving in to production. There are currently 4 participating campuses with a total of 100,000 records with 8 attribute/value pairs per record. The anticipated size is 250,000 to 300,000 records. There is one centroid at present which receives 400 hits/day. The final report of this project is due in July/August 1996.

The Swedish CIP Project is scheduled to run from August 1996 until January 1997 and is required to span ISPs, companies with their own catalogues and companies which run the service. Access will be via Whois++ with LDAP following later. Initially this project is for email addresses.

Hudson is a White Pages for the Government in Canada which is using Whois++ and CIP. This has been moved from test into production and can be accessed via <http://whitestar.ic.gc.ca/hudson>

3. Status of on-line Catalogs

The X.500 I-D has been sent to the list and needs to be reviewed. Sri and Linda offered to do this. The next stage will be for the implementors to be given two weeks to check their entries before moving to last call.

The Whois++ Catalog currently has three entries, Digger server from Bunyip (just released), IMAP client from ICL and Wombat which is a stand-alone X-Windows client from Bunyip.

4. Ph Drafts Status

"The CCSO Nameserver (Ph) Architecture" Draft needs to be reviewed by the Ph people. A two week deadline has been imposed to make a decision on this being put up for proposed standard.

The "Preferred Practices for Ph Directory Services" Draft has not been circulated. Joann will send a Draft to the list by October with the target status being informational.

5. DNS Aliases

Martin Hamilton gave an overview of the "Use of DNS Aliases for Network Services" Draft. A lengthy discussion ensued about this being

a short term solution until SRV available and that there seemed to be no single commonly used alias for CCSO. It was decided to move the discussion about the CCSO alias to the mailing list and the group felt that this Draft should be progressed as a BCP.

6. Discovery

The "Finding Stuff (Providing information to support service discovery)" Draft was discussed and the authors agreed to look at the Service Location Draft and if more work needed to be done the information should be fed back to that WG.

7. BCP on Directory

Concern was expressed that even though this Draft had be sent to the list there had been no comments. It was suggested that specific pointers to existing services would be useful.

8. Managing the X.500 Root Naming Context

Dante wants to migrate from Quipu to an X.500 1993 service but has run in to problems with the 1993 way of doing one level searches. Defect reports on the Standard are currently being progressed with the aim of getting a new profile to shadow a single entry. The aim is to re-issue this Draft as Experimental by September.

9. White Pages Schema Draft

Various comments on syntax and names have been received and they will be circulated to the mailing list for further discussion. The section on owner/creator needs to be made consistent and some clarifying text needs to be added about issues such as attributes should be supported by the server but don't have to necessarily be populated. Due to lack of time further discussion on multi-valued attributes will be held on the list. Tony will circulate a revised Draft by August 1st 1996.

10. Normandy

Microsoft gave a presentation on Normandy the directory service they are building using the IDS work.

11. Nomenclator

Nomenclator integrates CCSO servers and is currently an experimental tool. The Query Resolver and Distributed Catalog Service will be run at Bell Labs initially in order to gain experience before making the software available.

More information is available from :

<http://cm.bell-labs.com/what/nomenclator>

<http://cm.bell-labs.com/who/joann>

APPENDIX 7

Access and Searching of Internet Directories WG Meeting

Meeting Minutes

Wednesday, June 26, 1530-1730

- Agenda review/changes

The proposed agenda was accepted with minor reordering of some items.

- application/directory MIME type drafts
- application/directory framework

Tim reported that the application/directory framework draft had received several minor comments which would be incorporated into the next version. The group agreed that after these changes are incorporated, the draft should go forward as a proposed standard.

ACTION: Tim to produce a new version of the draft and submit it for proposed standard.

- versit/pdi vcard profile

Patrik Falstrom led the discussion of a number of problems/changes he has encountered with the vcard profile and application/directory framework. The issues included:

- Character set: The default is currently 7-bit ascii. The group agreed that this default could safely be changed to UTF-8, since it is a superset of 7-bit ASCII, and since there were alternate methods of selecting character set within the specification already.
- Encoding mechanism: The default is currently 7-bit clean encoding. The group agreed to change the default to 8-bit encoding, with the appropriate MIME or application/directory field-specific encoding to be used if a 7-bit encoding was required.
- Linebreaks: The current versit vcard draft states that CR, LF, and CRLF sequences should all work as line break sequences. Frank Dawson confirmed that the intention here was to do the same thing RFC 822 does by allowing whatever end-system representation was natural, but using a canonical representation on the network. Frank, Patrik, John Myers, and Dave Crocker agreed to come up with clarifying wording for the draft.

ACTION: Frank, Patrik, John, and Dave to clarify wording in the vcard draft.

In the course of the discussion, Dave Crocker volunteered to write up a separate Internet Draft describing how to handle the linebreak problem in the general case, since it appears to be something many people have run into before.

ACTION: Dave to produce "how to handle linebreaks" draft.

- Timezone representation: The current draft allows a number of timezone representations. The group agreed to standardize on always using the single format defined in RFC 822.

ACTION: Frank and Tim to produce a new version of the vcard draft with above revisions.

- Brief status of standards track documents

- WHOIS++ documents

Patrik Falstrom reported that revisions to the WHOIS++ documents are on hold pending the outcome of the FIND group, to determine how the indexing portion of WHOIS++ will change when using the common indexing protocol.

- LDAPv2 documents

Revised drafts have been submitted. One issue was raised regarding the language describing how T.61 strings are encoded in the attribute draft. Harald was volunteered to come up with some better language. It was also noted that the protocol version number (2) was not mentioned in the attribute draft. Tim volunteered to fix this.

ACTION: Harald to produce revised T.61 language.

ACTION: Tim to revise drafts and resubmit.

- labeledURI objectclass/attribute

The IESG objected to the definition of two attributes in the draft. The group agreed to revise the draft as the IESG requested, standardizing on only one attribute (labeledURI), mentioning the other as historical. After this revision, the draft will be put forward for proposed standard.

ACTION: Mark Smith to revise the labeledURI draft and resubmit.

ACTION: Harald to submit revised labeledURI draft for proposed standard.

- PGP objectclass/attribute

Roland reported that he had received no comments on the latest version this draft, except from David Chadwick, who pointed out that the draft does not allow storage of multiple certificates. Roland and David volunteered to work this out offline and submit a new version of the draft. The group agreed that once these changes are made, the draft should go forward for proposed standard.

ACTION: Roland and David to revise the pgp draft to handle multiple values and resubmit for proposed standard.

- WHOIS++/WHOIS/RWHOIS URL formats

The group agreed that each scheme deserved its own URL prefix, and decided on the following three prefixes: whois:, rwhois:, whois++:. There was some discussion of whether the "++" in the whois++ prefix would break any existing implementations. The characters are URL-legal, and the only problem people could think of was with some C++ preprocessors, which the group doubted were in wide use as HTML-parsing or generating tools.

ACTION: Martin to submit a revised whois++ URL draft.

- LDAPv3

- Summary of changes from version 00 to version 01

Mark Wahl presented a summary of changes between version 00 and version 01 of the drafts. Changes include:

- subschema subentries: Addition to allow schema elements to be stored in the tree itself, accessible as a group or from individual entries.
- manageDsaIT service control: Allows knowledge references within a server to be manipulated, rather than the entries they reference.
- preferredLanguage service control: Allows the client to select a preferred language. There was some discussion of the scope of this service control, and the group agreed it should apply to attribute values as well as error messages. An issue was raised on the list that language was not enough to solve this problem. One also needs to know the representation of the language (e.g., Kanji or Roomaji for Japanese).
- multi-stage binds: Allows a general challenge-response bind mechanism. There was discussion as to why LDAPv3 did not use SASL (Simple Authentication and Session Layer) for authentication.

Discussion focused on two issues:

1) Does the LDAPv3 bind mechanism support the same functionality as SASL? John Myers thought not, but Mark Wahl thought yes, with the use of other credentials in the LDAPv3 bind.

2) Is SASL appropriate for inclusion in the LDAP protocol? There was much discussion, and the group agreed to continue discussion on the list.

ACTION: John Myers to send a summary of his concerns to the list to stimulate discussion.

- binary attributes requested with ";binary": Allows a client to request that any attribute be sent back in binary (BER-encoded) mode. Servers are only required to support some attributes in this format (certificates and the like). If a server does not support ;binary for an attribute, it should act as if the attribute does not exist.

- paged results from search: Allows a client to request search results a page at a time, and to move around arbitrarily within the result set.

- modifyRights: This is currently a protocol addition. The suggestion was made on the list that this be an operational attribute rather than protocol extension. There was much debate over this, but in the end the group agreed that this was marginal functionality and should just be removed.

- add entry target system: From X.500, an extra parameter to the add request that allows an entry and its knowledge reference to be added to separate DSAs in a single operation. There was much discussion over whether this was needed. There were two objections: 1) The current form was X.500-specific, and included a presentation address; 2) The same functionality could be provided by multiple operations using the manageDsaIT service control. On the pro side, it was pointed out that X.500 already uses this feature to solve a problem and why can't we use the same thing? The group agreed to continue discussion on the list.

ACTION: Tim to revive discussion on the list.

- mapping onto SSL: A short section of the document was added describing how to map LDAP onto an SSL connection. There was some discussion as to whether SSL or TLS (Transport Layer Security) was the appropriate thing to be mentioning here. The group agreed that TLS should be added as soon as a stable document was produced that we could reference.

- Discussion of dynamic directory support

There was much discussion on this topic, which involves whether LDAP should be extended to support dynamic directories. A dynamic directory is one whose data changes frequently, and must be refreshed often. The prototypical application for this is a directory containing the current IP address of users who dial in or are otherwise mobile. There were several issues raised during the discussion.

First, does support for this belong in LDAP? Several people at the meeting had implemented separate dynamic directory schemes only to find that they ended up duplicating much of the functionality already found in LDAP. Based on this experience, they felt that the small extensions necessary to LDAP to support the required functionality were a good idea and well worth the price. Arguments against include the fact that LDAP is based on the X.500 models, which define a static directory, and the added complexity that dynamic directory support would introduce.

Second, how should this functionality be included? There were three proposals, with variants, put forward. The first was to extend the protocol with a new "refresh" operation allowing a client to continually tell a server that it is still alive. Responses would also be extended so the server could tell a client how often to refresh. Information from a client that does not refresh would be deleted from the directory. The second was to use the existing modify operation to maintain dynamic directory information. In this scheme, the "refresh" operation is accomplished by a normal LDAP modify operation. The server could communicate the refresh time to the client in an operational attribute the client reads using the search operation. The third was to define extended operations using LDAP's extensible operation capability to handle this situation.

Objections to the first scheme were primarily related to the added protocol complexity it introduces. Objections to the second scheme centered around the abuse of the LDAP modify operation.

Third, how often must refresh occur? The answer to this question depends in part on the application, how many clients are using the directory, etc. The wide variety of answers to this question led to the requirement that the server be able to tell clients how often it would tolerate refreshes.

Much discussion ensued, which had to be cut short for lack of time. The group agreed to continue the discussion on the list.

ACTION: Tim to revive dynamic directory discussion the list.

- Discussion of typeless DN support

This discussion item had to be cut due to lack of time. The group agreed to discuss it on the list.

ACTION: Tim to revive typeless DN discussion on the list.

ACTION: Mark to revise the drafts to reflect consensus reached at the meeting and during subsequent mailing list discussion.

- Any Other Business

The meeting concluded, slightly late, with an agreement to meet again in San Jose.

APPENDIX 8

Minutes for the FIND meeting at the 36th IETF in Montreal

Time: WEDNESDAY, June 26, 1996, 1930-2200

Chair(s): Patrik Faltstrom <paf@swip.net>
Roland Hedberg <Roland.Hedberg@umdac.umu.se>

Author of the minutes:
Roland Hedberg <Roland.Hedberg@umdac.umu.se>

Helping us with spelling corrections etc:
Ken Weiss <krweiss@ucdavis.edu>

0. Agenda review/changes

The proposed agenda was accepted without changes.

1. Why two parallell CIP drafts ?

Patrick explained that he and Roland shared the view that Chris Weider's draft didn't reflect the consensus of the group reached at the LA meeting and also had too much whois++ stuff in there. Therefore a second draft was produced by by Jeff Allen and Patrik Faltstrom. The intended outcome of this is that these two drafts will be merged into one.

2. Charter of the find group

There where some discussion about which papers were going to be produced. The consensus was that there should be one document specifying the CIP, another one specifying how to use centroids as one special case of indexes within the CIP and further for each

client - server protocol that is going to use the CIP one paper describing the mapping between the data representations and one describing the access method.

3. LDAP/CIP work at Umea University

Roland Hedberg presented the work he has been doing to enable a X.500 DSA to work as an index server and he also presented a WWW-interface that can use this index server.

The WWW-interface can be reached at

<http://macavity.umdc.umu.se/~roland/query2.en.html> and the indexserver it accesses contains all the information presently accessible in the Swedish branch of the X.500 DIT (~50.000 entries). For the time being the index only contains names of people. Roland will produce a draft describing the objectclass and attributes needed to accomplish this .

4. The new CIP draft

Jeff Allen presented the gist of the new draft. The discussion following the presentation led up to a couple of unresolved items:

The use of MIME - should/can INDEX-CHANGED be structured as a MIME message

Aggregation ala CIDR - facilitate query routing.

Incremental updates - per application domain or general.

Security - both regarding exporting indexes and data protection.

Centroid scaling issues - certain datasets only contain unique items which means that the resulting index is no smaller than the original dataset.

Frontends to indexservers might only speak one access protocol - clients speaking another access protocol can not pass this server, while climbing the tree upwards or downwards, which means that parts of the mesh might be unaccessible to the client.

5. Workshop of Distributed Indexing and Searching

Erik Selberg presented some ideas on using query routing within the Web indexing sphere which came out of the workshop .

It was felt that introducing query routing and distributed index servers is a necessary step in the development of the Web indexes since the current centric approach doesn't scale.

More info on the workshop can be found at

<http://www.w3.org/pub/WWW/Search/9605-Indexing-Workshop/>

It was agreed that followup work undertaken by the query routing contingent from the Distributed Indexing/Searching Workshop would be folded into the FIND working group.

6. The CIP and CCSO

Martin Hamilton presented his work on integrating CCSO nameservers with the CIP. His conclusion was that it was viable but that there remained some items that have to be resolved. There is no standard URL format for a CIP referral to a CCSO nameserver. For the time being Martin proposed that one could use the gopher one (`gopher://ccso.server.domain.name:105/2`).

Another question is whether the CCSO should the CCSO attribute names and types be normalized to a common schema.

7. Scaling of the CIP

Patrik presented some graphs showing the relationship between the size of a centroid and the size of the actual datasets both when looking a people informations from the phonebook and large document collections. Phonebook information revealed the not very astonishing fact that phonenumbers are unique which means that the centroid increased almost linearly with the growth of the dataset. Removing phonenumbers from the centroid gave a much slower growth and it also appeared to be asymptotic. When indexing words out of documents the curve didn't seem to level off when the dataset grew (max dataset size ~12.000.000 tokens). When applying a stop list weeding out very frequent words and very unusual words the curve became asymptotic, reaching 60.000 and levelling off to be leveling of at that value.