NameFLOW-Paradise

Half Year Report January-June1997

TABLE OF CONTENTS

Introduction

This Report reflects the NameFLOW activities and operations for the first half of 1997. The report is intended for people interested in the NameFLOW service and in particular those working for the national networks responsible for national Directory services. The report deals with the operational aspects, the information aspects and liaison activities respectively.

The 'specials' in this Report are:

- The results of the NP-93 test pilot
- Closely related, a short description of the EEMA Directory Challenge-97, where DANTE participated.
- Trip report of Vincent on the ÒEuropean Directory ForumÓ (EDF) meeting hosted
- EEMA and IETF working group minutes

A paper version of the Quarterly Reports is made available to DANTE's customers. An electronic copy will be made publicly available via the web*, where appropriate without customer sensitive information.

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

Vincent Berkhout
DANTE
Francis House
112 Hills Road
Cambridge CB2 1PQ
UK
Tel: +44 1223 302 992
Fax: +44 1223 303 005
E-mail: nameflow@dante.org.uk

* URL: http://www.dante.net/nameflow

Operations and Development

This quarterly report summarises activities and operations of the NameFLOW-Paradise Directory service, and associated information services, for the six months from January through July 1997. Pointers to the individual reports for each month can be found in the References section.

1. Operations/Helpdesk

The presentationAddress attributes of the FLDSAs Chinchilla (CH/LI), Tarpon, cn=Fruit Bat (c=US) and cn=Elephant Seal (c=IS) (SK) were changed at the request of their respective country managers.

The Internet Draft on storing URLs/URIs in the Directory became RFC 2079 in January. This was announced to FLDSA managers for their information.

Following last quarter's Issues on the results of installing router-filters, the managers of some of the DSAs with out-of-date knowledge references were contacted, with varying results.

Mirroring was set up of the EMA (Electronic Messaging Association) Challenge 97 information, and associated files.

A private FTP area was set up for the upcoming X.500(93) pilot tests.

Helpdesk staff participated in the NP-93 test period, with an emphasis on root-context replication.

2. Outages

In February significant outages of service totalled approximately 30 minutes. In April significant outages of service totalled approximately 11 hours scheduled. There were no significant outages in January, March, May or June, but there was an extended problem with IP routing in May, which primarily affected DANTE's web/X.500 gateway. Details of outages are available in the monthly reports.

3. Issues

EuropaNET/X.25, the trans-European academic X.25 network, was discontinued at the end of March, as agreed between DANTE, DFN, SURFnet, and UKERNA. This should have no impact on X.500 operations, as it appears that there are no DSAs reliant solely on this network for connectivity. DSA managers are advised to remove addresses with DNIC 2043 from their configuration, to prevent attempted connections.

Router-filters installed on a LAN at ULCC in February show attempted accesses to the old IP address for Giant Tortoise, and to the defunct FLDSA Ocellated Turkey. The filter logs will be used to try to identify the DSAs with out-of-date knowledge references, and hence the relevant managers who can update them.

It appears that ISODE/Quipu 8.0 suffers from "Year 2000" incompatibility, which will lead to problems with date comparisons after 31 December 1999. It is questionable whether we should still be offering this code via anonymous FTP.

David Chadwick's document on root-context replication appeared as RFC 2120 in April. It can only be hoped that vendor will follow this to produce X.500(93) implementations that will allow practical management of the upper levels of the DIT.

The manager for c=US has reported that he is unable to link his current FLDSA to standalone LDAP servers using the X.500 Enabler from Critical Angle. It is believed that the difficulty may be with the initial configuration rather than "production" use, and it is hoped that this can be resolved.

In June, a T3 circuit (45 Mbps) was brought into service between the UK academic network, JANET, and North America. This should improve performance of Directory operations for Giant Tortoise and UKERNA DSAs into Canada, the USA, and beyond.

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 information-server, under:

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

In the DSA statistics in appendix 2, the reduction in binds by Directory technicians after January is due to the withdrawal of a top-level probe run by SWITCH in Switzerland.

The LDAP statistics for May in Appendix 2 were affected by the extended IP-routing problem mentioned in Outages.

The numbers of remote binds to Giant Tortoise in Appendix 2 have dropped from the previous quarter. The drop is attributable to a lack of connections from the FLDSA cn=Caretta Caretta (c=GR).

The LDAP figures show an unusually high number of connections in February, and large connect times in February and March; indeed the connect time for March is larger than for the entire previous quarter. As usual, most connections came from DANTE's web/X.500 gateway, and it may be that the high number of connections in February is a symptom of a web-crawler in action. Regarding connect times, it is notable that in each month only 20% of connections lasted more than 7 seconds, while a very small number of long-lasting connections contributed a significant proportion of the total connect time. These long-term connections presumably result from error conditions of some sort.

As the EuropaNET X.25 network was dismantled at the end of March, DUA figures for it are no longer shown in Appendix 3.

The FTP figures in Appendix 4 are somewhat lower than the previous quarter. It is believed that this is due to stability in the EMA and OSIdirectory mirrors.

5. NameFLOW-Paradise X.500(93) Test Report - Phase Two

Introduction

This is the test report from the NameFLOW-Paradise X.500(93) pilot (abbreviated to NP-93 in this document) testing that took place between 21st May and end of June 1997. A full description of the test plan can be found at http://www.dante.net/np/93pilot/phase2-plan.html. The actual tests performed were based on modified versions of the EuroSInet X.500(93) and LDAP Interoperability Test suites to make them more appropriate for the NP-93 environment.

The test suites are available in Word 6 and postscript formats:

X.500(93) :
http://www.dante.net/np/93pilot/np-x500(93)-tests.doc
http://www.dante.net/np/93pilot/np-x500(93)-tests.ps

LDAP :
http://www.dante.net/np/93pilot/np-ldap-tests.doc
http://www.dante.net/np/93pilot/np-ldap-tests.ps

The first part of this document gives a general description of the findings. The second part gives a summary of the problems reported and the issues raised. In an attempt to make this report as vendor neutral as possible, the names of the products are not mentioned. For further information you are advised to consult vendors or participants. Originally the testing was only intended to last for one week, but due to a slow start and the interest generated, testing effectively went on for over a month.

Overall Summary

In general the testing carried out was much more successful than during Phase 1 in 1996. All participants managed to get DAP and DSP working successfully. LDAP testing was also included this time around, which was also very successful. This all shows that the products have made a definite step forward during the past year.

The problems entering T.61 characters still seem to be present. The problems associated with configuring access control schemes, even relatively simple ones, seem to be quite large. It would appear major user interface improvements are required to allow administrators to enter Access Control Information (ACI) easily in an error-free fashion.

Only two participants attempted the shadowing tests and nobody attempted the ACI tests on shadowed data.

No significant improvements appear to have been made in support for Root Context and a definite commitment (including timescales) is sought from the vendors in this area.

DAP , LDAP and DSP

As with the Phase 1 testing carried out in 1996 (http://www.dante.net/np/93pilot/phase1-results.html), there were no real problems with DAP and DSP testing. LDAP testing was officially included this time around, which also seemed to go well. There were a few minor problems experienced with authentication and chaining, but the vendors concerned have already issued patches for these problems.

No explicit X.500(88) -> X.500(93) interworking was carried out. However, TU Chemnitz-Zwickau used an "88 DSA" for mastering its information.

DISP (Directory Information Shadowing Protocol)

SWITCH and Technische UniversitŠt Chemnitz-Zwickau attempted the shadowing tests, which seemed to work with only a couple of problems, noted below.

At this time it is unclear as to why so little DISP testing was performed, it may be that people ran short of time due problems generated by other tests. Only complete naming context shadowing with complete update was being attempted, so it should have been a lot easier than in the Phase 1 testing.

Root Context

Currently as far as can be ascertained, no vendors product supports the RFC 2120 proposal on the Root Naming Context and as such it was very difficult to validate the workability of this proposal. Root context operation was attempted by the root DSA and also by two country level DSAs.

As was noted in Phase 1, First Level DSA (FLDSA) single level search operation did not provide the outcome sought by NP-93 participants. NP-93 requires the FLDSA just to return the list of sub-ordinate DSAs and **not** for the DSA to chain to all the sub-ordinate DSAs, as it worked in the Quipu model.

Access Control Information (ACI)

Some ACI testing was carried out after some initial configuration problems, there still seems to be difficulty in participants being able to correctly configure all the relatively simple ACIs required. To quote one participant: "...ACI configuration is very tricky...".

No ACI tests were performed on shadowed data, presumably because only two participants managed to get shadowing to work, most people had problems with the un-shadowed ACI tests.

Better tools are required for administrators to be able to configure "sensible" ACIs quickly and without errors. If the experts involved in this testing have problems with the complexity of the current tools provided so will most other users.

DSA Management

In general the tools do not appear to have improved significantly over last year. The main problem seems to relate to the difficulty discussed in section 6 about the lack of good interfaces for configuring ACIs. It is difficult to come to any conclusions on the state of shadowing management as so little testing was done.

Although it is known that more advanced DSA management tools exist these were not made available for testing.

Product Stability

DSA stability is still somewhat worse than "88-Quipu", but maybe this is an unfair comparison as this was used in a largely single vendor environment.

There were a few occurrences of DSAs crashing but these were all fixed by software updates.

User Testing versus Vendor Testing

Since the Phase 1 testing performed last year there have been two EuroSInet workshops, one in Copenhagen and another in Munich (see http://www.eurosinet.org/workshops/ for more information), both of which seemed to get tests involving quite complex ACIs (not the same ones tested by NP-93) and shadowing to work.

As an aside, one must remember that the participants at these events are often the software developers or people with direct access to the software developers. They still experienced problems configuring their DSAs for ACI tests and getting the shadowing tests to run. Admittedly, the EuroSInet tests are more complex than those used for NP-93 testing, but one assumes that if developers are having problems in such a multi-user environment so will most users.

NP-93 Test Results:

A. Test Reports

The following organisations submitted test reports:

Organisation	DSA(s)
TU-Chemnitz-Zwickau	@c=DE@o=Technische Universitaet
DANTE	@c=GB@o=DANTE
TU Delft	@c=NL@o=Technische Universiteit Delft
Uniwersytet Mikolaja Kopernika w Toruniu	@c=PL@o=Uniwersytet Mikolaja Kopernika w Toruniu
Telecom PTT	@c=CH@o=Telecom PTT
SWITCH	@c=CH@o=SWITCH
Universiteit Twente [On behalf of SURFnet]	@c=NL@o=Universiteit Twente
ULCC	Root DSA @C=GB

Brunel University @c=GB@o= Brunel University

B. Issues for future testing

1. It would be useful to investigate the possibility of distributing a text file, with the required test entries in an appropriate format for "bulk-loading" into DSAs. Typographical errors in test DITs were a problem, using a "pre-configured" DIT should help avoid this type of problem and significantly expedite testing.
2. A partner assignment matrix should be displayed on the NameFLOW web pages. This would keep participants up to date on what is happening and possibly encourage more testing.
3. Complete participant information should be available on the NameFLOW web site in addition to spread sheet and "root file" information on the FTP server.
4. Production of a "cheat sheet" to help new participants over some of the common pitfalls of testing.
5. Not all participants submitted test reports and those that were received could have captured more information. This situation might be solved by better form design and easier result submission by improved use of the web in addition to existing mechanisms.
6. Clarity of the DIT configuration in the test suites could be better. This could be reviewed as part of any test suite revision.
7. There were some problems confirming that referrals had worked. (try: showentry -nocache -dontusecopy -nochain "@c=xx@o=xxx@ou=NP-93@cn=Person One")
8. FLDSAs should not have a superior DSA configured -- their only relation to the root DSA should be through their shadowing agreements with it.
9. There is a very minor issue concerning the entry for the root DSA itself. The DSA implementation running as root may require an entry in its local DIT for itself. If so, this should be prevented from being shadowed to FLDSAs by extending the chopBefore part of RFC 2120 section 4.3 to exclude the DSA entry. This isn't for security, simply to stop the DSA DN showing up in a list at the top level. This assumes that the FLDSA implementations do not need an entry in their local DIT for their shadow supplier. If this information is required, then it might just have to be lived with, relying on DUAs filtering it out, as indeed people do now.

C. General Testing Problems

1. Vendors do not appear to be committed to genuine multi-vendor interoperability. Possibly for very good commercial reasons, many seem to unofficially perceive that there is no money to be made spending too much effort working on multi-vendor interoperability. This is backed up by the fact that most directory products are being sold into homogeneous environments where interoperability is not an issue.

2. None of the products tested conform with RFC 2120, although several vendors say they will in a future release, but will not give a timescale.
3. Products are not easy to configure, especially as far as Access Controls and Shadowing Agreements are concerned. This caused quite a few tests to fail as the ACIs were incorrectly configured.
4. There were some problems with the use of T.61 characters. Largely to do with inconsistent/confusing data entry methods. Again this caused quite a few tests to initially fail before the errors were corrected.
5. There were problems in the test suite with the incorrect use of slashes in textEncodedORaddress.

D. Overall Conclusions and Way Forward

Two major conclusions can be drawn as a result of the testing that took place:

1. X.500 is perceived to be too complex to install and maintain for the needs of the NP-93 community.
2. There is lack of significant improvement in the quality of the software provided by the vendors for testing.

As a result the NP-93 community need to investigate alternatives to pure X.500. These might take the form of LDAP only, employing some kind of ÔRoot LDAPÕ referral server or a hybrid X.500/LDAP solution.
These issues need to be discussed as soon as possible and a way forward developed

E. Appendix A - Compendium of Specific Testing Problems

```
Number: 1
Problem  Type:  First Level  DSA Configuration.
Detailed Description: Performing a search at country
level for sub-ordinate DSAa causes  all the sub-ordinate DSAs to be
contacted. Solution: This is currently a "feature" of X.500(93) and
this behaviour  will not  change until RFC  2120 is  implemented by
vendors.  See X.518  19.3.1.2.1 step  3 and  19.3.2.2.1 step  7 for
further information.

Number: 2
Problem Type: Entering Hexadecimal values in presentation
address.
Detailed Description: There  is standard way of describing
presentation address.  Various syntaxes are described  in different
places: ISO, IETF, EuroSInet, vendors and so on. Hence how to enter
the desired  presentation address is unclear.  Solution: A somewhat
counter-intuitive solution was supplied by the vendor.

Number:  3
Problem  Type: RFC  2120 "fast-track  solution" Clarity.

Detailed  Description:  It is  not  clear  from reading  RFC  2120,
exactly  what  information  a  FLDSA should  store  to  allow  this
functionality  to work.  Solution:  Seek  clarification from  David
Chadwick.

Number: 4

Problem Type:   DSA stability problems.

Detailed Description:   DSA losing information when restarted.
Solution:       Needs further investigation with Vendor.

Number:  5
Problem  Type: Use  of labeledURIObject  and labeledURI.

Detailed Description:  Not displaying correct object  and attribute
names, but  correct OIDs  are used. Solution:  System administrator
was  inadvertently using  beta  oidtables from  the previous  NP-93
testing.
```

Number: 6
Problem Type: DUA bind with incorrect credentials.

Detailed Description: DSA accepts the bind but downgrades access to "unauthenticated" without any error message. Solution: Fixed in later version of software.

Number: 7
Problem Type: Chaining.
Detailed Description: Time out
when running chaining tests. Solution: Specify "no-timelimit" option. Also investigate system clocks not being in close enough synchronisation.

Number: 8

Problem Type: DSA instability.

Detailed Description: Chained search causing DSA to crash.
Solution: Fixed by patch from vendor.

Number: 9
Problem Type: DUA bind problems.
Detailed Description:
When trying to bind with DUA error message "problem with DSA" or "Error received : Service unavailable". Solution: Not enough information to form a conclusion. Further investigation required.

Number: 10
Problem Type: Reading operational attributes.
Detailed
Description: Operational Attributes are not accessible by anyone. [test 9.4.4] Solution: Further investigation required.

Number: 11
Problem Type: Shadow update.
Detailed Description: It is
not possible to configure the DSA for updates every five minutes (Only on change, once an hour, once a day, once a week, consumer initiated). Update once an hour works properly. [test 9.4.10] Solution: The five minute update cycle was chosen as a convenience for testing. This is unlikely to be a realistic time for most live systems. This issue needs to be brought to the attention of the vendor.

Number: 12
Problem Type: DSA Instability.
Detailed Description:
SLDSA crashes when a DSP search is initiated by the FLDSA (test 7.4.8). This problem also occurred with another DSA from the same vendor. Solution: Patch from vendor fixes the problem

Number: 13
Problem Type: Displaying wrong object class and
attribute names.
Detailed Description: Certain objectClass and
attribute type names used in the testing document were replaced by different ones in the output display. Solution: This problems is a product of primary and secondary naming of OIDs in the DUA oidtable files. This is not a problems as such. Users just need to be aware of it when they get unexpected names displayed. Presumably the oidtables could be altered as required.

```
Number:  14
Problem  Type:  Problems  installing   DSA.
Detailed
Description: Software  would only install correctly  in the default
directory .  Solution: A somewhat messy  work-a-round was developed
using symbolic links.
```

Liaison

EEMA

DANTE was present with a stand at the EEMA Annual Conference. (See next section "Information" under exhibitions.IETF

The 38th IETF took place in Memphis, Tennessee. Minutes of the meetings of the Access and Searching of Internet Directories WG, the Integrated Directory Services WG and the FIND WG are attached in Appendices 7, 8, and 9.

European Commission

A workshop was attended by NameFLOW representatives that was hosted by the European Commission to determine how electronic directories should develop in the short future to allow a level playing field for all PNO once the markets are liberalised. A trip report is included as appendix 6.

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 maintained 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 half year 1997
January 1997
February 1997
March 1997
April 1997
May 1997
June 1997

Conferences/exhibition

From 12-16 June 1997 DANTE was present with a stand at the EEMA Annual Conference and Exhibition. This year's exhibition was relatively successful as we Show Floor Participants where people had to visit the stand in order to win a prize. Over a hundred people listened to what NameFLOW actually does. As an indication of what the Challenge is about a short introduction is included.

"WEMA Challenge 97 consists of a series of ground-breaking live demonstrations of the real business benefits of linking strong security measures with X.500 electronic directories to enable electronic commerce. The WEMA Challenge 97 is a global event, it was kicked off at EMA '97 in Philadelphia in April, and has now moved to this Exhibition in Maastricht. From Europe it will move onto Australia in October.

Each set of demonstrations will involve many companies at each public event, demonstrating under live conditions, how X.500 electronic directories work, their benefits, and just how secure they can be with modern security applications.

"This is not a future technology" says Chris Taper (a Senior Consultant with ICL and Chairman of the European Electronic Messaging Association's Security Privacy and Legal Committee) "With the collaboration of industry leaders world-wide, we will be able to show

that all this is possible today. Our European initiative is not being performed in isolation, through the World Electronic Messaging Association (WEMA), all the world's messaging associations are working on a global basis to significantly influence developments within the electronic commerce and messaging industry.

Challenge 97@Electronic.Commerce.Europe will demonstrate that secure messaging over different networks and computing platforms is not only technically feasible but can be used to satisfy business oriented needs for electronic commerce in a practical way."

European Electronic Commerce practitioners and solution vendors co-operating in the Challenge 97@Electronic.Commerce.Europe, at the 10th EEMA Annual Conference and Exhibition in Maastricht in June, are listed below:

Show Floor Participants:
BT, ISOCOR, NEXOR, DANTE, ICL, o.tel.o, Directory NeT Inc., JEMA, Siemens Nixdorf Informationssysteme AG , Enterprise Solutions Ltd, Lotus, SOFT-SWITCH, Grabowsky Polytechnics BV, MaXware and Utimaco Belgium"

APPENDICES

APPENDIX 1 - Helpdesk summary for Jan/Feb/Mar 1997

Country		Number of queries			
Full Name	ISO Code	January	February	March	Quarter
(Armenia)	AM*	1	-	-	1
Australia	AU	1	-	-	1
Canada	CA	-	1	-	1
Switzerland	CH	1	-	1	2
Germany	DE	1	1	1	3
Denmark	DK	1	-	-	1
France	FR	2	-	-	2
United Kingdom	GB	2	4	5	11
India	IN	-	-	1	1
(Iran)	IR*	1	-	-	1
Japan	JP	1	-	-	1
Malaysia	MY	1	-	-	1
Netherlands	NL	1	-	-	1
(Peru)	PE*	1	-	-	1
Poland	PL	-	-	1	1
Sweden	SE	1	-	-	1
Slovakia	SK	-	1	-	1
United States	US	4	4	5	13
Total Requests		19	11	14	44

Full Name	ISO Code	April	May	June	Quarter
Australia	AU	1	-	-	1
Canada	CA	-	-	2	2
(China)	CN*	1	-	-	1
Germany	DE	-	-	4	4
(Algeria)	DZ*	-	1	-	1
Spain	ES	-	-	1	1
United Kingdom	GB	-	1	1	2
Ireland	IE	-	-	1	1
Israel	IL	-	1	-	1
Iceland	IS	-	-	1	1

Italy	IT	-	-	1	1
Japan	JP	-	1	-	1
Netherlands	NL	-	1	2	3
New Zealand	NZ	1	-	-	1
Sweden	SE	-	1	-	1
United States	US	6	5	2	13
Total Requests		9	11	15	35

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

APPENDIX 2 - World Root DSA and LDAP summary statistics for Jan/Feb/Mar 1997

Summary of calls to DSA Giant Tortoise

No. of binds	January	February	March	Quarter
Local	312	305	356	973
Remote	4545	3598	3664	11807
Total	4857	3903	4020	12780

No. of binds	April	May	June	Quarter
Local	387	406	443	1236
Remote	2851	2693	2614	8158
Total	3238	3099	3057	9394

No. of operations	January	February	March	Quarter
Local	10	3	2	15
Remote	38609	36692	36005	111306
Total	38619	36695	36007	111321

No. of operations	April	May	June	Quarter
Local	12	4	8	24
Remote	38108	38454	37183	113745
Total	38120	38458	37191	113769

System usage (calls received)	January	February	March	Quarter
Binds by Directory technicians	852	305	356	1513
Reads of DSA entries	35	6	13	54
Other ops on DSA entries	1	2	1	4
Getedb operations (inc. slices)	38442	36555	35612	110609
Spot shadows	81	81	69	231

System usage (calls received)	April	May	June	Quarter
Binds by Directory technicians	387	406	443	1236
Reads of DSA entries	23	24	7	54
Other ops on DSA entries	12	0	2	14
Getedb operations (inc. slices)	37841	38213	36876	112930
Spot shadows	72	118	176	366
Total	38335	38761	37504	114600

LDAP usage

	January	February	March	Quarter
Connections	21648	46798	22170	90616
Total connect time (seconds)	867914	1449296	1884540	4201750

(4201750 seconds is 1167 hours 9 minutes 10 seconds)

	April	May	June	Quarter
Connections	26350	17736	18901	62987
Total connect time (seconds)	1681843	681726	262203	2625772

(2625772 seconds is 729 hours 22 minutes 52 seconds)

APPENDIX 3 - Public DUA summary statistics for Jan/Feb/Mar 1997

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	January	February	March	Quarter
Internet	927	578	639	2144
UK academic X.25 (JANET)	3	2	4	9
EuropaNET X.25	0	0	0	0
Public X.25	15	0	0	15
Total	945	580	643	2168

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

Network	April	May	June	Quarter

Internet	513	478	424	1415
UK academic X.25 (JANET)	3	2	1	6
Public X.25	5	0	0	5
Total	521	480	425	1426

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

Domain	January	February	March	Quarter
edu	431	220	239	890
unresolved	125	72	105	302
uk	101	99	60	260
com	39	41	47	127
dz	25	28	63	116
net	45	15	17	77
ca	18	18	15	51
nl	25	10	7*	42
us	21	10	11	42
de	13*	14	9	36
Total	843	527	573	1943

Domain	April	May	June	Quarter
edu	193	133	113	439
unresolved	63	59	92	214
uk	42	73	39	154
com	35	46	35	116
net	32	41	18	91
dz	29	26	33	88
de	18	8	7*	33
nl	7*	13	12	32
us	13	5*	14	32
org	12	13	4*	29
Total	444	417	367	1228

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

APPENDIX 4 - FTP/Gopher summary statistics for Jan/Feb/Mar 1997

WWW server

	January	February	March	Quarter
Unique hosts	713	677	789	1964
Number of HTML requests	1923	1940	2449	6312

Number of non-HTML requests	346	432	440	1218
Number of malformed requests	79	66	134	279
Total number of all requests/errors	2348	2444	3035	7827
Total number of Kbytes requested	32146	60540	48329	141056
Average requests/day	75.9	87.4	98.0	87.0
Kbytes/day	1038	2164	1560	1568

	April	May	June	Quarter
Unique hosts	971	888	859	2440
Number of HTML requests	2801	2502	2324	7627
Number of non-HTML requests	1088	720	1028	2836
Number of malformed requests	181	110	132	423
Total number of all requests/errors	4074	3335	3485	10894
Total number of Kbytes requested	80142	82809	85657	248609
Average requests/day	136.0	107.8	116.4	119.9
Kbytes/day	2676	2676	2861	2735

FTP server

	January	February	March	Quarter
Files Transmitted	385	347	466	1198
KBytes Transmitted	105881	100637	101093	307612
Average Files Daily	13	14	19	15
Average KBytes Daily	3529	4193	4212	3978

	April	May	June	Quarter
Files Transmitted	231	165	196	592
Kbytes Transmitted	63251	43176	49720	156147
Average Files Daily	10	7	7	8
Average KBytes Daily	2635	1799	1841	2091

Gopher server

	January	February	March	Quarter
Total connections	12	15	13	40
Total files retrieved	10	6	8	24

	April	May	June	Quarter
Total connections	10	12	6	28
Total files retrieved	14	1	1	16

APPENDIX 5 - European Electronic Messaging Association

EEMA Directory Committee (Part I - 25 February, 1997)

Minutes of the Meeting 25 February 1997 (9:00 - 13:00), Maritim proArte Hotel, Berlin, Germany

Source Vincent Berkhout - Project Manager

1. Opening

VB opened the meeting and welcomed the members. The chairman David Goodman (DG) DG was not able to attend the meeting due to personal circumstances and sent his apologies. Colin Robbins (CR) was asked to act as a vice chairman and this was anonymously accepted. Later during the meeting this was formalised with a show of hands for the record. CR is now officially the Vice Chairman of the Directory Committee. Minutes of previous meeting accepted and attendee list was circulated.

2. Liaisons/Education

EWOS (Keith Richardson)

It seems that the interest in X.500 profiling and in particular the 1997 edition of X.500 is declining slightly. As a result EWOS is redefining its role (see previous minutes).

(VB Note: EWOS has a new web site, see)

KR explained the LDAP hausse is giving X.500 a bad name. But even with the LDAP popularity there will be enough room for X.500 to exist and develop. KR has send out a mail to find X.500 vendors to form a forum, not with a true X.500 vs. LDAP group, but more a subgroup of the LDAP vendors with an interest in X.500. So far he has received five positive responses to his mail within a week. If all goes well he is planning a meeting to start the forum in London.

ISO (Colin Robbins and Keith Richardson)

A short overview was given of the current status of X.500(97) which is now in the concluding phase. The document has improved in the following areas:

- management control via DOP

- storage of "secured" information in the DIB (signed and if needed encrypted) attribute can have a context, e.g. language, or time ("this telephone number is reachable from 9:00 to 17:00 on working days").
- several problems were resolved of X.500(93)

The 1997 edition of X.500 will be the official standard.

NameFLOW-Paradise (Vincent Berkhout)

VB reported on the NP-93 interworking test with X.500(88)/Quipu directories and also scaling within the NameFLOW community. Other organisations/vendors will be allowed to participate in the tests. The pilot has three phases. This is second phase is where coexistence with X.500(88) and LDAP will be tested. In addition NP-93 will look at large data sets and the consequences. They are using a tailored version of EuroSInet test suite that will fit the NameFLOW general requirements. The test suite will be extended to test First Level DSA replication (X.500(93) DISP).

NADF/EMA (Colin Robbins)

CR gave an update on NADF/EMA. Two documents have been put out recently on how to build an X.500 directory and business issues. The NADF is expanding into the World Directory Forum/ EMA chapter, EEMA chapter, etc. Issues that go beyond the Challenge will be dealt with by this group.

EuroSInet (John Horton)

JH reported on the recent workshops to do X.500 conformance session in Munich. One workshop was held in Brazil in co-operation with BRISA. The next test writing workshop will be in Dublin (IE) 10-11 March 1997. One workshop is planned in May in Brussels (BE) to do X.500 and LDAP and allow some testing for the Directory Challenge/Security work. The Brussels workshop will also focus on secure e-mail and if possible some extended DISP testing.

IETF (Colin Robbins)

Work on LDAP is rapidly progressing towards a scaled down version of the LDAPv3 draft. (LDAPv2 has been implemented at the Univ. of Michigan and is distributed as release 3.* which causes some confusion). LDAPv2 is a proposed standard at this time and will go to a full standard, if not superseded by LDAPv3 prior to this. A draft of LDAPv3 came one year ago. Issues being addressed in the new revised draft:

- extension mechanism
- internationalisation (beyond ASCII character set)
- ecurity within LDAP and support for certificates
- referrals

The comments/changes have all been progressed and the updated version is already available and expected to be accepted at the next meeting as a draft standard. LDAP servers are moving in the direction of not using X.500 to tie them together and an ad hoc solution, e. g. using DNS in some way is expected. Another proposal is to replace the "traditional" distinguished name with an RFC822 name (enables finding the LDAP server and name on either side of the @). Most people use the directory to look up the e-mail address - so there is an anomaly with this approach. Another proposal is only addresses on the right hand side of the e-mail address use the DNS structure as Directory Information Tree. There are two CIP terms currently used at the IETF, one for the Common Indexing Protocol and one for the Common Internet Person (or LIPS: Lightweight Internet Person Schema). As an alternative Directory GOLDD pops up which stands for - Good Old Directory on DNS Domain names.

3. Projects

Completed: Top Level Naming In Europe (TOPOL) - Phase 1

Top Level Naming in Europe (TOPOL) Phase 1 report completed by Logica is in an almost final form and soon will be a completed project. A few minor editorial mistakes have to be corrected and is than ready for redistribution. The document basically reviews work that is going on in the directory area; however it focuses on X.500 and does not address LDAP implication, which appeared on the marketplace during the period of study. The success of LDAP and Stand Alone Directory Servers will have an impact on the document, but it was decided that this could be addressed in a short paper rather than a long document.

Proposed Project: (1)Top Level Naming In Europe (TOPOL) - Phase 2

The TOPOL-2 project is put on hold until the EEMA BOD has decided on the budgets. TOPOL-2 will need budget for an editor of the document. The document will reflect the recommendations of the committee members and a subgroup will be formed to do the work/ discussions. An editor will be selected/appointed and will take care of the layout and other wordings. The turnover time for a document should be three months, although this sounds too ambitious. The TOPOL-2 document will now be a series of short documents each dealing with one of the (7) issues. The first document would be, "a recommended DIT structure for multinationals". The second potential document that was discussed is the Multiple Service Providers problem. Petter (PL) and Pipsa (PN) said they could do a page or two on their planned solution and we could than take this as a basis for the EEMA recommendation. A suggestion was made that this has an overlap with the "High-level Naming" project. Unfortunately, this project is planned for a two year period and will not be ready before that. The group decided that this would take too long.

Action: Send mail to assess TOPOL interest for break-out session (VB by default)

Action: PL submit a two page description of the proposed MSP solution

Action: PN submit a two page description of the proposed MSP solution

Proposed Project: Understanding LDAP (David Goodman/Colin Robbins)

The first draft of "Understanding LDAP", initially written by CR with DG as co-author, was tabled, however not yet ready for redistribution. A discussion followed what the document should look like and according CR it was already getting to big (8 pages maximum). The intend was a quick Q&A (question and answer) document to explain the misconceptions and what LDAP actually is. The paper was intended to look at LDAP from an X.500 perspective. It was decide to split the document in two and start with the simple and short Q&A paper, finish and publish it and move onwards from that. The second paper would be something like, " LDAP, protocol and beyond" addressing the usage of LDAP and Directory Servers. A brief moment was spend on asking some questions that should be in the Q&A paper. It was decided that recent developments of LDAPv3 and LDAP servers should be briefly mentioned.

Action: VB to review next draft before distribution.

Action: CR split document and distribute to list. (deadline: before next meeting.)

Action: CR & DG Proposal for second document

Proposed Project: Business Drivers for Directories (Stephen Vercella)

The business drivers project is on hold for this meeting and will be removed from the project ladder next meeting if no organisations offers to help Barclays progress this project. SV has asked the EEMA User Committee (and UK User Group) if people were interested in the project but received no response. SV will try to progress the project within Barclays and will look for possible contributions from other business peers. Good Business Drivers were discussed - one was directory support for an up-to-date printed telephone book. The conclusion was the review was too much from the Users' perspective. It was felt that the project should be a "Review of Directories" including real figures for business cases. The group felt the emphasis should be on an enterprise directory concept. There should also be a distinction between intra-company directories and inter company directories.,p>

Action: PP - Write "Directory Business review" proposal to see if it is possible to get a few business cases together, and incorporate actual costs. These could then be given to potential customers.

Proposed Project: World Directory Forum - Europe

DG has put forward a proposal for the WDF, the perspective/buy-in looks good. JG indicated that the EMA NADF, when looking at renaming from the *American* to *Atlantic* would only include the Europe/NA groups. The perspective is to go to a World group with regional chapters. The point was also made that the Directory Committee should make the transition to the WDF and the focus should be on projects that have current implementation perspective rather than written documents. The Challenge environment will form the basis from which the WDF would work to evolve into a public directory service.

The terms of reference (ToR) for the WDF arrived during the meeting and so were unable to be discussed in great detail. CR read out the ToR and a few comments were made. The great question was how the WDF would it be organised, would everybody attend a once a year "world meeting" [no - only world co-ordination anticipated and maybe a meeting with representatives only]. Would the EEMA DirComm be split into two groups? [no - the DirComm would continue its work with a bigger audience and possibly with more contributors]. In general people were curious how WDF would work - the logistics and benefits. An example may be the TOPOL document with appendices specific for each regional chapters.

Proposed Project: EEMA Directory Challenge 97

As a special challenge meeting was held the day before this topic was not further discussed.

4. Next Committee Meeting

The next committee meeting will be at the Annual Conference in Maastricht 10-00 to 12.30 on 15 June 1997

5. AOB

None.

Acknowledgements

Thanks to Joanne for producing the first version of the minutes.

Attendees

Attendees

Colin Robbins NEXOR CR (Vice Chairman)
Dr. Mirko LŸdde SAP AG ML
Freddie Dawkins Media Training FD
Jens Ramsb¿l Tele Internet JR
Joanne Ghahremani BTNA JG
Setsaas MaXware AS JS
John Horton EuroSInet JH
Richardson ICL KR
Michio Kato JEMA MK
Per Hagero WM-data PH
Peter Pawlita Siemens Nixdorf PP
Petter Ljunggren Sweden Post Ltd. PL
Pipsa Ylœ Nononen Telecom Finland Ltd PN
Roger Dean EEMA RD
Stuart Battersby Boldon James SB

Vincent Berkhout DANTE VB
(Project Manager)

INTERNATIONAL

Apologies

David Goodman Lotus DG (Chairman)
Bernd Franzke c/o Dr. Materna GmbH
Daniel Langeron Alcatel
David Chadwick University of Salford
David Ray SmithKline Beecham John Erik
Pharmaceuticals
David Walker Enterprise Solutions Keith
Limited
Eddie Gillis REUTERS
Emil Mayer Siemens.
Erik Andersen Fischer&Lorenz
Ernie Pope Boldon James
Esko Kippo TeamWARE Group
Gertjan Vaessen Fenestrae
Greg Wynne Axsis
Jacqueline Pasquerea EDF GDF
John Shillitoe IBM Global Services
Keith Reid Barclays Bank
Mark Vandenwauver ESAT/COSIC
Martin Helin Finland Post Ltd.
Martin Reber Lotus
Michael Frank Deutsche Telekom AG
Mick Reid SYSTEC
Mr John Heaney Reuters Limited
Niels K. Jensen Kommunedata
Per Cedell LM Ericsson Data AB
Ralph Ehlers F. Hoffmann-La Roche Ltd
Rod Cole Telstra
Rod Ward IBM Global Services
Sergio Galliano CSELT
Stephen Gilbertson MAFF
Stewart Bridge Boldon James
Tomislav Micic URSA MAJOR

Vincent Cobes Global One

European Electronic Messaging Association (Part II - 15 June 1997)

Draft Minutes

Sunday - 15 June 1997 (10:00 - 12:30)

MECC Conference Centre, Maastricht, Netherlands

1. Opening

The chairman David Goodman (DG) opened the meeting and welcomed the members at the early hours of a Sunday morning (Appendix 1). The minutes of the previous meeting were accepted and the action items checked. (Apologies, Appendix 2). Included in these minutes for the first time is a glossary of terms used (Appendix 3).

DG started the meeting with a presentation on the history and achievements of the Directory Committee. The presentation dealt with what has happened since it was proposed in Montreux (June Ô93) and inaugurated in Maastricht (September Ô93). DG related that the primary concerns at the time were about providing public services (1), with worries about the respective roles of PARADISE, Eurescom and the NADF and how they would all co-exist (2). There was also an assumption that wide-scale deployment of X.500 was just around the corner, despite perceptions about lack of products (3) and complexity of deployment (4). The subsequent work of the Directory Committee was to address those issues and DG enumerated the projects that have been successfully completed:

1. *Multiple piloting: infrastructure and interworking demonstrations (Amsterdam '95, Brussels '96, WEMA Challenge '97) culminating in the formation of the WDF (World Directory Forum)*
2. *Co-existence: TOPOL (Top-level naming in Europe)*
3. *Lack of Products: X.500 Product Guides (V1 and V2)*
4. *Complexity of Deployment: Hitchhikers Guide to Corporate Directory Deployment*

A more recent co-existence issue has been addressed in the latest document, Understanding LDAP and X.500.

2. Liaisons/Education

EWOS (Erik Andersen)

EA gave a presentation on recent changes within EWOS. EA made it clear that EWOS had to change with the times and there had been a desire to dissolve EWOS. It looks now that EWOS will be Òabsorbed Ó in CEN EBCES and ICITC. There are currently three projects running on the following subjects: Profiling, Directory Deployment and High Level Naming and all three projects are funded by the Commission. Some of the X.500(93) profiling has been completed and there is interest from the US in this. A new project will be started on profiling X.500(97). Another new project is based on the extensions to the Directory (which have been accepted by ISO/ITU). There are a few problems with the current strategy, such as slow implementation and adoption by vendors and service providers. EWOS is not widely recognised, and in some cases ISO/IEC are claiming intellectual property rights with the experts and companies not gaining credit for their work. This compares unfavourably with the Internet standardisation process where these problems do not exist through the widespread publication of RFCs i.e., unlimited distribution and proper accreditation with the result that these standards are now becoming the industry standards.

Other work is being done in collaboration with EIDQ, mostly on X.500 extensions, DMD interconnections, corporate directories catalogues of schema definitions (stored in web pages), OSITOP liaison, European use of Internet specifications and the use of databases to act as an information centre for all kinds of directory information. The work of high-level naming is part of EPII (European Project on Information Infrastructure) from ETSI and is an area where there could be beneficial collaboration with EEMA.

European Directory Forum (no representative)

The European Commission is investigating how to set up a true European Directory and has hired consultants to do a study. The key objective for the Commission is to use it for telephone numbers to allow fair competition after the deregulation of 1998. This is a potentially very large Directory with entries of over 250M people. It looks like a paper exercise, as the Commission is not willing to fund parts of the infrastructure. Bernd Stadler (BS) said that Coopers & Lybrand have been employed to do a second study. David Ferris (DF) suggested that it maybe a way of breaking the monopoly situation and telephone numbers, were a starting point. The Commission has said this has to be solved by the marketplace (service providers) on a co-operative basis or they would legislate on a European or national level to make a European Directory happen.

Swedish Directory Indexing Group (Bernd Stadler)

A Swedish Directory Indexing Group, has been awarded a project, by the Swedish Government. The key players are Roland Hedberg and Patrick Falstrom who are working on a common indexing protocol to support a directory of e-mail addresses in Sweden by the end of 1997 **(BS to provide summary for minutes)** based on a combination of Whois++ and X.500.

NameFLOW-Paradise (Vincent Berkhout)

The main current activity is the NP Ô93 pilot where there are about twenty interested parties, although only ten are actively participating. The tests focus on the operation of the first level DSAs, access control and replication. There is a lack of confidence that the root context is working. Access control is proving difficult with a lot of problems coming to the surface, having used the Eurosinet test suite. Replication was working a year ago, and should be working again soon. The group is looking for co-operation with the WEMA Challenge and the October event in Australia. Steve Kille (SK) observed that NP Ô93 was pushing harder than Challenge Ô97 on the critical X.500(93) issues. He suggested it would be useful to look into some of the problems in more detail and see how these problems relate to Challenge Ô97 issues and feed edited NP stuff into the infrastructure. A break-out session was agreed on operational perspectives on NP Õ93 problems and how to move the process of co-operation between NP Õ93 and Challenge Ô97 forward.

X.500 Vendor Forum (Keith Richardson)

The current profiling work is diminishing so a new driver to do profiling outside standardisation process needs to be found.

Earlier this year, KR was promoting an initiative to start a forum primarily for X.500 vendors. This came in response to the meteoric rise in profile of LDAP and an attempt to push X.500 DSAs as powerful LDAP servers. A brief (marketing) meeting was co-ordinated at EMA Õ97 in Philadelphia in April, but (suggested SK) nothing positive came out of it. In the next one-two months ICL and others will put out an invitation for a meeting London to gauge the real interest in such a group - if this doesn't happen KR will abandon the initiative.

This posed the question as to what X.500 vendors are going to do individually and

collectively to respond to the LDAP challenge? DG supported KRÕs initiative suggesting that X.500 vendors should get together to make access controls and replication work with the consequence that if they donÕt do this convincingly and coherently X.500 is going to suffer. The next meeting of the ITU study Group 7 on Directories will be looking at LDAP (July 1997). Ainis Noritis (AN) suggested that Directory Committee attendees might benefit from a resume of what is going on in the LDAP/X.500 world. DG replied that that is one of the reasons people come to the meetings.

Clive Betteridge (CB) said that EuroSInet have passed the LDAP test suite to the Internet Mail Consortium (IMC), (as they do not have any of their own) to use as a starting point for the LDAP interworking workshop. SK suggested the first event is really just a hand waving exercise and not really technical. EuroSInet will be holding an event in the Autumn 1997 which will be mainly LDAP-centric. He also informed the Committee that vendors have reached a plateau where their own products work fine in a single vendor environment, but only work in a mixed environment with the support of the people who wrote the code. SK suggested that a lot of X.500 vendors were treating the EuroSInet workshops as a marketing tool rather than a true interworking event.

A discussion followed on the role of X.500 and where it is going and it appears that there are problems within organisations with this question. DG suggested vendors need to get together and sort out the problems, not only for servers but also for desktop and back-end systems.

SK said that the EuroSInet workshop had demonstrated the perception of LDAP as being a highly interworkable technology, was shown to be flawed - including the reference implementation.

What can be done for users who get LDAP 'by accident' - such as in an email package - and do not know what LDAP actually does? BS expressed disappointment on the way that the Challenge Õ97 activities do not appear to have taken X.500 forward, as the same problems on top level, replication and security are still re-occurring. ÒThe marketÓ is looking for an available form of certification of products.

SK suggested the need for a new model for the Challenge. According to DG there is still a lot of interest in X.500 amongst user organisations (even more than previously) and a lot of opportunities/possibilities. In order to get a sense of real interworking, SK suggested that any participating user organisation should have products from at least two independent vendors.

Ferris Research (David Ferris)

Ferris Research has two papers available: an X.500 tutorial and a Practical Implementation Guide. Both papers are available at the Ferris stand.

EMA (Joanne Ghahremani)

The Challenge Ô97 event at EMA in April established a country level infrastructure interconnecting 31 DSAs. Three applications using the directory infrastructure were

demonstrated:

- white and blue (for government) pages
- open EDI scenario;
- voice profile for mailing by taking telephone numbers and converting them to email addresses (VPIM).

The Challenge achieved a high level of interoperability between DSAs. The event pushed the idea that the Directory is here today. SK asked for positive comments to be circulated (already available from EMA). CR suggested that the negative articles were published by trade magazines and the positive articles by government magazines.

3. Projects

Completed:

Top Level Naming In Europe (TOPOL) - Phase 1

The Top Level Naming in Europe (TOPOL) Phase 1 report was made available in its final format (paper and on diskette) and was distributed among the committee members. Each EEMA member could have one free copy.

Ongoing Project:

Understanding LDAP and X.500 (David Goodman/Colin Robbins)

This paper is in a ÒFrequently Asked QuestionÓ format with the goal of explaining the merits and de-merits of both LDAP and X.500, including questions about the interworking and co-existence between the two as well as a way forward. It will be a ÒlivingÓ (regularly updated) document on the EEMA Web site. More comments would be taken during the week of the conference with the intent of publishing as soon as possible afterwards. A one page executive summary was suggested to help promote the document and gain acceptance more easily with trade magazine and newspaper publishers. DF offered to help with this. SK noted that the paper at times took a defensive position, which needed to be reviewed and put in a more positive light.

David Goodenough (DGA) and others said that the document was too technology-focused and doesn't adequately explain why anybody should want either LDAP or X.500 - what as a user can it do for me? CR replied that there was no point in changing scope at that time, but it could serve as the basis of a new document. VB said that originally there was an idea to have a second document anyway (see previous minutes)

Ongoing Project:

World Directory Forum - Europe

A WDF was originally proposed a year ago, taken to WEMA in October and finally given the go-ahead in April in Philadelphia. From a European perspective this is an opportunity to

work with the NADF in North America on an equal, global footing.

Learning from the experience of the NADF and NameFLOW-Paradise, the intent of the WDF is neither to architect a solution for a single global directory nor to determine how one should be deployed. Neither is it to pre-judge protocol solutions, but to understand and cater for the evolution and co-existence of multiple extranet and public directory services over the next few years. There is an emerging opportunity for a critical mass of multi-vendor/multi-protocol directories involved in electronic commerce of one form or another. An initial approach is to review the Challenge Õ97 and see if it can be taken a stage further.

The terms of reference were distributed and were further discussed in a break-out session.

Ongoing Project:

EEMA Directory Challenge 97

Unfortunately there was no time to discuss this in full detail but there were two conference sessions allocated for this. A special thank-you was extended to CR for co-ordinating activities in Europe.

Proposed Project:

Directories in Use (Per Hagero)

PH produced and distributed a proposal for a questionnaire looking at the decision-making process involved in setting up and running directory projects in user organisations in different vertical industries. As the meeting was running late, the project discussion was taken up in a break-out session.

4. Next Committee Meeting

The next Directory Committee meeting will be at the EEMA Autumn Conference at the Nova Hilton Hotel, Geneva, Switzerland on the 5 November 1997 from 14h00-17h00. (The focus of the conference is Business Quality Messaging.)

5. AOB

VB said that as he was moving jobs within his company, he would have to resign as a project manager for the Directory Committee. Anyone wishing to apply for the position as EEMA project manager for the Directory Committee they should contact either VB or EEMA. John Horton (JH) is now acting as the temporary co-ordinator for NameFLOW-Paradise. Permission was given to advertise for a NameFLOW directory expert on the Challenge mailing list.

Acknowledgement:

Thanks to John Horton for taking the minutes.

Attendees

1. Ainis Noritis \ ABB Infosystems (CH)
2. Ann Brown \ Nortel Technology (CA)
3. Astrid Lamberts \ o.tel.o (DE)
4. Bernd Stadler \ Telia InfoMedia (SE)
5. Bruce Hunt \ Cascade Communications (UK)
6. Clive Betteridge \ EuroSInet (UK)
7. Colin Robbins \ NEXOR (vice-chair) (UK)
8. David Chadwick \ University of Salford (UK)
9. David Ferris \ Ferris Research (US)
10. David Goodenough \ DGA (UK)
11. David Goodman - Lotus Development (chair) (US)
12. Debbie Burkert \ GE Information Services (US)
13. Diana Dwamena \ Cable & Wireless Business Networks (UK)
14. Eric Andersen \ Fischer & Lorenz (DK)
15. Gerhard Schmied \ INFONOVA (AT)
16. Ian Clark \ Texas Instruments (US)
17. Jane Jobson \ EEMA (UK)
18. Jan van der Weele - Dupont Nemours (DE)
19. Jens Ramsb \ Tele Denmark Internet (DK)
20. Joanne Ghahremani \ BT North America (US)
21. John Erik Setsaas \ MaXware (NO)
22. John Horton \ DANTE (UK)
23. John Wimpole \ The Radicati Group (UK)
24. Karel Mak - Shell Information Services (NL)
25. Keith Dixon - Origin (NL)
26. Keith Hadden \ British Gas (UK)
27. Keith Richardson - ICL (UK)
28. Lena Johansson - Consultant
29. Magnus G. Panderson \ PostNet Infra (SE)
30. Marko Kaittola \ Telecom Finland (FI)
31. Mathias Bjarme \ Telia (SE)
32. Michael Dufva \ Enator Communications (SE)
33. Michael Kipp - Europol (NL)
34. Niclas Karlsson \ Enator Communications (SE)
35. Nicolas Pougnet - Europol (NL)
36. Niels K Jensen \ Kommunedata (DK)
37. Nuno Moura \ SIBS (PT)
38. Per Cedell \ L M Ericsson Data (SE)
39. Per Torlof \ University Hospital Lund (SE)
40. Peter Farrington \ British Gas (UK)
41. Phil Ratcliffe \ Boldon James (UK)
42. Pipsa Yla. Monorjen \ Telecom Finland (FI)
43. Ralph Ehlers - Roche (CH)
44. Richard Bailey \ Boldon James (UK)
45. Steve Kille \ Isode (UK)
46. Susan May \ Concert Management Services (US)
47. Vincent Berkhout - DANTE (project manager) (UK)

Apologies:

1. Rod Ward - IBM (UK)

2. Holger Wosnitza - Dr Materna (DE)
3. Henriques Cesar de Conti - BRISA (BR)
4. Greg Wynne - Axsis Consultants (UK)
5. Jose Eduardo Pina Mirana - SIBS (PT)
6. Herbert Blankenburg - o.tel.o (DE)
7. Yukio Saito - NTT (JP)
8. George B. Huitema - KPN Research
9. Larry Guyot - Sionnest - BARTS
10. Vicky Ellis - Oracle InterOffice (UK)

(Note 1: Not all attendees registered.)

Glossary for some of the terms used is from http://www.dante.net/np/ds/glossary.html

CEN ?
DMD Directory Management Domain
EBCES ?
EDSG European Directory Services Group
EEMA European Electronic Messaging Association
EIDQ European International Directory Enquiries
EPII European Project on Information Infrastructure
ETSI European Telecommunications Standards Institute
Eurescom An organisation sponsored by many of the European PNOs, responsible for initiating research and development projects on behalf of the PNOs
EuroSInet ?
EWOS European Workshop for Open Systems
ICITC International Construction Information Technology Conference
ICT Information and Communication Technology
IEC ?
ISO International Standardisation Organisation Responsible for a wide range of standards, including networking
ITU-T International Telecommunication Union - Telecommunication, a standards-making body for telecommunication operators (PTTs)
NADF North American Directory Forum
NP NameFLOW-Paradise
ISOTOP a European user organisation
PARADISE a European pilot for X.500 directories
RFC Request For Comments
WEMA World Electronic Messaging Association

APPENDIX 6 - Trip note: European Commission

European Directory Forum

Workshop on the barriers to and the problems with
the implementation of a European Directory
Brussels, 14 March 1997 by Vincent Berkhout;

A workshop focusing on the scope of the European Directory and on identifying potential barriers to its implementation. The workshop was well attended and the participants came

from various Directory areas. The objective of the workshop was not completely clear, but the EC seems to have a double barrel approach to promote directories. The first track is via official communications with the objective that users only need access to one directory to get all information and to abolish current monopoly rights. The problem here is property rights and privacy issues. The second track is via ECTRA and the ENF to promote "city and mobile users" after the liberalisation in 1998. From these approaches it is clear that the real European Directory the EC targets for is a telephone directory, including fax and mobile users and not really the general purpose directory holding other types of information. To promote directories the EC sees two possibilities, firstly via the TelCos and secondly upwards via upwards pressure e.g. via Internet usage. A general conclusion of the workshop is that this is a study for the EC to break the current directory monopoly of TelCos, there is no plan afterwards and there will be no funding.

The invitation to the workshop:

"The European Commission services has recently appointed Coopers & Lybrand to undertake a study on the barriers to implementing a European Directory. The study will also address issues such as objectives, benefits and scope of a European Directory. The Commission view this as an important study which will provide a basis of reference to help the Commission services, together with Directory Providers, Regulators, IT providers and other interested parties, decide on the necessary actions for the development of a European Directory.

To ensure that the analysis, conclusions and recommendations of the study are founded on a broad and representative range of views and interest, Coopers & Lybrands will be organising a number of workshops and an interview programme over the next four months. Since the success of the study is heavily dependent on the active participation of all interested parties, I would appreciate it if you could afford the Commission your time to contribute to this study."

A conclusion of this meeting for NameFLOW is that this is way out of our depth as we are competing with PNOs and other commercial players. Furthermore, the scope of any developments is too restricted in the sense of Directory content diversity (telephone numbers only) and a too broad/large a target base of customers, that is all telephone subscribers in Europe aiming at 500 Million entries in 1998.

It is strongly advised to tread careful in this arena.

APPENDIX 7 - IETF: ASID WG

Access and Searching of Internet Directories WG Meeting
Meeting Minutes
Wednesday, April 9, 1545-1800, 2000-2100
Reported by: Tim Howes and Patrik Faltstrom

- Introduction of new co-chair

Patrik Faltstrom was welcomed as the new co-chair of ASID.

- Agenda review/changes

The proposed agenda was accepted without change.

- Hour 1 1545-1700: LDAPv3 core documents

The following documents were discussed, with the goal of making any minor changes needed, issuing a last call, and putting the documents forward for proposed standards status.

draft-ietf-asid-ldapv3-protocol-04.txt
draft-ietf-asid-ldapv3-attributes-04.txt
draft-ietf-asid-ldapv3-dn-00.txt
draft-ietf-asid-ldapv3-filter-00.txt
draft-ietf-asid-ldapv3-url-00.txt

There was some discussion of the master/slave designation added to the LDAP URL draft draft-ietf-asid-ldapv3-url-00.txt. The group felt that this was not a good thing, in the absence of a replication model, and since it did not fit well into the general URL concept.

ACTION: Tim to revise the URL draft to remove this feature.

There were no revisions proposed to the filter draft draft-ietf-asid-ldapv3-filter-00.txt.

Chris Newman commented that the string dn format draft draft-ietf-asid-ldapv3-dn-00.txt needed some revision of its ABNF.

ACTION: Chris to send proposed ABNF revisions to Mark (DONE).

ACTION: Mark to produce a new DN draft.

The following comments were made on the attributes draft draft-ietf-asid-ldapv3-attributes-04.txt. The draft needs a keyword index, to make it easier to find things. The keyword index, though thought generally useful, was not a must-have at this time. The userPassword syntax should be deprecated and discussed in the security considerations section. The audio syntax, which currently refers to the "SunOS 4.x format" should either be updated to reference audio/basic, or removed.

ACTION: Mark to produce a new attributes draft with these changes incorporated.

There was a fair amount of discussion on the protocol draft draft-ietf-asid-ldapv3-protocol-04.txt. Most of the discussion centered around the use of SASL in the document and whether it was correct. The following consensus was arrived at:

- The SASL credentials should be OPTIONAL
- No SASL mechanism name should be included on the BindResponse
- Some language was needed advising that changing the SASL authentication layer during an LDAP session is allowed, but changing the SASL security layer is not allowed.

An additional point was raised by Nick Emery about the handling of unknown filter components in searches. The current draft says such components are to be treated as though they match no

entries. This leads to the undesirable result that a filter searching for (!(unknown component)) returns all entries. After some discussion, the group agreed to adopt the tri-state logic used by X.500.

Jeff Hodges suggested a number of clarifications in the text describing subschema subentries and extensible matching.

Another discussion topic involved the lack of protocol encoding examples in the draft. Examples of on-the-wire client/server interactions were thought to be very useful, and a pretty standard part of many other Internet protocol specifications.

Mark Wahl stated that he is working on a separate document explaining the basics of BER encoding needed by LDAP. This document will contain somewhere between 8 and one hundred examples. Since BER encoding is (unfortunately) more complicated than many of the simple text encoding used by other protocols, a separate document (or perhaps appendix to the protocol document) was thought to be the best approach.

To avoid holding back LDAPv3 at this time, the group agreed that the document should go forward, with examples being added when the document goes from proposed to draft.

The group also discussed the fact that the LDAPv3 drafts referenced the SSL specification, which is not an Internet standard. The group agreed that the goal is to reference the TLS specification when it becomes available. The status of TLS was not known, but it was thought that it should be moving forward soon. The group agreed to change the LDAPv3 reference to TLS in anticipation of the TLS specification being progressed along with LDAPv3.

There was further discussion of the relationship between SASL and TLS, and that the lack of clarity on that relationship is contributing to the slow progression of SASL. The group agreed that this relationship should be clarified ASAP.

There was also discussion about the fact that the LDAP documents reference the X.500 standard in many places, and that X.500 is not freely available on the net. The question was whether this would prevent the drafts from progressing. Based on precedents set by SNMP and other IETF work, the group did not feel this should be a problem.

ACTION: John Myers to work with the security area director and the TLS group to clarify the SASL and TLS specs.

ACTION: Mark to produce a new protocol draft with these changes incorporated.

ACTION: Tim and Patrik to ensure that examples are included before LDAPv3 goes to draft standard.

The group agreed that after the documents were revised (estimated time to revise: 1 week), last call should be issued on the ASID list. At the conclusion of a successful ASID last call, the documents should be given to the area directors for approval of the IESG.

ACTION: Tim and Patrik to issue last call to ASID on revised

documents.

ACTION: Tim and Patrik to request progression of the documents
pending successful ASID last call.

- Hour 2 1700-1800: MIME-DIR and WHOIS++

- WHOIS++ drafts

Patrik reported that new WHOIS++ drafts have been produced
with no protocol changes, only revisions and clarificationÕs
from operational experience implementing the protocol.

One example of a such clarification is the addition of a
grammar for the output from a Whois++ server to the existing
grammar for the input to a Whois++ server.

The group agreed that a last call should be issued on the
revised documents, after which they should be put forward
for draft standard status.

ACTION: Patrik and Tim to issue last call on the revised
WHOIS++ documents and progress to the area directors.

- application/directory framework

There was much fruitful discussion of the application/directory
framework document. The first issue discussed was whether the
content-type should be changed to text/directory. The argument
for is that the information is primarily textual in nature and
the desired behaviour is to have the content-type displayed to
users even if the type is unknown. After a brief struggle, the
group agreed to change the content-type to text/directory.

A more contentious issue surrounded the use of the "lang"
parameter defined in the draft. The values of the "lang" parameter
are currently defined to be language tags from RFC 1766. Patrik
argued that an additional tag (he proposed calling it "default")
was needed to support the use of text/directory in WHOIS++.
The tag is needed so that applications (like WHOIS++) may
determine which (if any) attributes to return in the event
that the language requested by the client is not present.

After much debate, misunderstanding, and confusion, it was
decided not to add this to the spec. The argument was made
that 1) the default solution is not entirely correct, since
it does not also provide a way to determine the type of the
default language returned, and 2) the parameter set is already
extensible, so something could be defined later to solve
this problem.

The final issue discussed was the use of the "charset"
attribute parameter, allowing the character set to be set on
individual values within a text/directory content-type.
This was felt to be a Bad Thing and contrary to the MIME
way of doing things and contrary to the IAB character set
proclamation (thou shalt use UTF-8) by several people.

After a bit of debate, the group decided to remove the "charset"
attribute parameter and require the UTF-8 "charset" MIME header
parameter be specified by default on text/directory content-types.
The result of this is the loss of ability to switch charsets

on a per-value basis within a text/directory component, but this was thought to be a good thing.

ACTION: Tim to produce a new MIME-DIR draft with the agreed on changes.

ACTION: Tim and Patrik to progress the revised draft to the area directors.

- vCard profile

One comment received prior to the meeting was the use of English as the default language in the vCard profile. The group agreed that this statement should be removed, and that there should be no default language associated with the vCard profile.

Two additions to the vCard profile had been proposed to the ASID list by the MOPA consortium (Mobile Office Promotion Association). The additions were for a CLASS attribute and a PCS property on the TEL attribute.

The CLASS attribute would identify the class of the information contained in the profile (e.g., PRIVATE or PUBLIC).

The PCS property would identify a TEL attribute as referring to a PCS telephone.

The group considered these additions useful, but there was some discussion of where we should draw the line before putting vCard forward as a proposed standard. After a bit of discussion, the group agreed to allow these two additions and then progress the draft to proposed standard.

ACTION: Frank Dawson to revise the vCard draft with these changes.

ACTION: Tim and Patrik to progress the revised draft to the area directors.

- Hour 3 2000-2100: Various LDAP documents

This hour began with the daunting task of listing the 15 (count them 15) documents submitted for the group's consideration. The documents were:

draft-ietf-asid-ldapv3-lang-00.txt
: use of language codes in LDAPv3

draft-ietf-asid-ldapv3-strong-00.txt
: SASL authentication mechanism for X.500 authentication

draft-ietf-asid-ldapv3schema-x500-00.txt
: LDAPv3 definitions of X.500 schema

draft-ietf-asid-schema-pilot-00.txt
: LDAPv3 definitions of pilot schema

draft-ietf-asid-ldapv3-simple-paged-01.txt
: LDAPv3 extension for simple paged results

draft-ietf-asid-ldapv3ext-03.txt
: LDAPv3 extension for dynamic directories

draft-ietf-asid-replica-selection-00.txt
: How to use SRV records to select LDAP servers

draft-ietf-asid-ldapv3-sorting-00.txt
: LDAPv3 extension for sorting results

draft-ietf-asid-ldapv3-referral-00.txt

        referrals and knowledge references in LDAPv3
draft-ietf-asid-ldapv3-api-00.txt
        Update of RFC 1823 for LDAPv3, etc.
*draft-ietf-asid-ldap-mult-mast-rep-00.txt
        Multi-master replication proposal
*draft-ietf-asid-ldif-00.txt
        LDAP text directory interchange format
*draft-ietf-asid-changelog-00.txt
        LDAP text changelog format
draft-ietf-asid-email-routing-su-00.txt
draft-ietf-asid-email-routing-ns-00.txt
        Two email routing using LDAP proposals

The group started by agreeing not to try and discuss
all the documents, but rather to spend the first part
of the meeting deciding what to discuss. The group first
decided that the two email routing documents were outside
the scope of ASID, so they were crossed off the list.

The group next decided that replication was the topic of most
interest, with replica selection a close second. So, discussion
began on replication with an attempt to answer three questions:

1) Should we be working on directory replication?
2) What group should do the work?
3) What should be the scope of the work?

There was much debate on these topics, mixed together with
debate about the future of the ASID group. The latter topic
was raised with the idea (put forward off-line by the area
directors) of splitting ASID into two groups: one for LDAP
and one for text directory stuff (WHOIS++, MIME-DIR, etc.).
The group thought this was basically a good idea.

Debate on question 1) quickly consensized on a decision
that replication is definitely a problem that we should be
tackling.

Debate on question 2) was somewhat tangled up with the scope of
the work. Some people felt that replication was a big and
separable enough problem that a separate working group was
required. Others felt that replication would soon drag in other
problems (e.g., access control, schema) that really need to be
considered by the entire group. Consensus on this issue was
rough, at best, but the group seemed to be leaning towards
handling replication in ASID (or the as-yet-unformed LDAP
group), for the reasons given above.

Question 3) generated lots of discussion. The proposals
ranged from a general replication solution, to a general
LDAP-only solution, to an LDAP-only solution specific to
either multi-master or single-master models. There was much
concern that to make any progress we should try to focus
the problem as narrowly as possible. After much discussion,
the group consensificated that we should narrow our focus to
LDAP replication, and not try to solve the more general
problem.

After much further debate with little progress, the group
decided to switch gears and discuss one of the other documents.
The replica selection document was chosen. Paul Leach
described the draft briefly, which defines a method of locating
LDAP servers, given a domain name, based on SRV records. The

group thought the concept useful, but very similar to a general procedure outlined in a draft from the DNS group for using SRV records. The group agreed that Paul should talk to the DNS group to ensure that they two documents did, in fact, overlap, and that everything needed in Paul's document was present in the DNS SRV record document.

ACTION: Paul Leach to follow-up with the DNS group.

ACTION: Tim and Patrik to chase the group-splitting issue with the area directors.

- Any Other Business

Noting the lateness of hour, the general glazed look of the participants, and the fact that another meeting's attendees began filing into the room, the ASID meeting concluded, about on time.

The next ASID meeting will be in August in Munich, Germany.

APPENDIX 8 - IETF: IDS WG

IDS WG Meeting Minutes

Wednesday, April 9, 9-10am

Reported by: Linda Millington

1. Liaison Reports

The liaison reports will be posted directly to the mailing list. Barbara Jennings will forward details about the EMA directory Challenge to the list.

2. Documents Status

Managing the X.500 Root Naming context has been published as RFC 2120 (Experimental)

Use of DNS Aliases for Network Services will be submitted to the next IESG meeting as a BCP

A Common Schema for the Internet White Pages Service will be updated then submitted to the IESG as Proposed Standard

The X.500 Catalogue is currently progressing as Informational

Best current Practice for the Internet white Pages Service received very little comment during Last Call and will be submitted to the IESG as a BCP

3. Work Items Outstanding

The two Nomenclator Drafts will be updated and published as Informational

Progress on the PH Architecture Draft is sought in the near future or this item will be dropped from the work list

4. Naming Plan for an Internet Directory Service

The Approach for Using Domains in LDAP Distinguished Names and the Naming Plan for an Internet Directory Service Drafts were merged and distributed to the mailing list as agreed in San Jose. The Group discussed this Draft and recommended that the following changes be made in order to make the intent clearer and that the future intent was to forward this as Proposed Standard :

- Crisper requirements are needed
- What problems are being solved needs to be clarified
- The minimum criteria necessary to comply with this scheme must be defined
- Wording needs to be tightened up on the implied finding of LDAP servers

There was also discussion on which properties of DNs should we depend on with them currently being used for uniqueness, search constraint bases and actual information held in the entries. Further discussion on this topic will continue on the mailing list.

Migration to and implementation of this naming scheme belong in a separate Draft and it was suggested that the WHOIS++ deployment experience should be heeded and that operational experience (12 months + at least) should be sought before finalising the scheme.

APPENDIX 9 - IETF: FIND WG

Summary of the Find Working Group

38th IETF, Memphis<

Chair: Patrik Faltstrom

Minute Taker: Sally Hambridge

SUMMARY: The Find Working Group discussed the problems the group has had making progress. Patrik has decided to change the grouping of the documents to:

+ Framework

+ Mimetype/Base Functionality

+ Index Objects

+ Transport Mechanisms

We discussed the CIP-soif-00.txt draft, the ldap-01.txt draft, and the draft-weider-cip-hierarchy-00.txt

Patrik emphasised the need to make progress of face extinxtion.

Minutes of the Find Working Group

Chair: Patrik Faltstrom

Minute Taker: Sally Hambridge

Patrick opened the meeting by talking about the lack of progress and the need to keep discussions going on the mailing list and on the documents or the group will be shutdown.

The documents have been re-grouped:

+ Framework - to be authored by Michael Mealling and Paul Leach

+ Mime Types/Base Functionality

+ Index Objects

+ Transport Mechanisms

Documents should be authored by May and offered to the WG in last call by June.

We still see lots of problems on how to administer the MESH and with incremental updates. We also need operational experience in order to know what's going to work. But Patrik would like to see the drafts go forward as Proposed standard rather than Experimental.

The framework doc will *not* contain:

- Incremental indexing - since we are unsure as to whether it should be part of the protocol or specific to the index type

- What protocol is required for access to the servers - whether we have to force all the servers in the same mesh to have the same access protocol.

Some open Issues:

- Base URI/DSI - some need both, some need only Base URI

- Security

- Parameters in mime-specification - which are general and where to put the index parameters in the mime-type

Chris Weider said there are implementors waiting for the docs and that implementation experience would come soon after publication.

Ted Hardy presented his cip-soif document. Soif originated from the work done in the Harvest project and is envisioned as a query referral for the web. Ted went through the payload format, and grammar. He also talked about template types. His matching semantics were that a query identified attribute should be considered to match a Soif attribute when a case-insensitive character by character comparison matches that portion of the Attribute identifier prior to the hyphen. So if a person searched for author, author-1 author-2, etc. would match the attribute requested. There was a question of whether the attribute side would match US-Ascii only, and Harald suggested that that made the most sense for debugging.

Ted also talked about matching of data, which is also problematic. Should we match only exactly, which would give fewer responses but no false negatives (a requirement of the protocol) or should we allow fuzzier matching which would give more responses but perhaps allow false negatives. Ted suggested that a strategy would be to limit that which would be queried to those attributes which most likely foster correct referrals.

Ted also presented a template type of Dublin Core attributes as an example. He then talked about CIP-Hints, which he hoped would foster discussion. The hint template would give a weigh list and a threshold (under which no values would be reported) as a way of

aggregating data to be passed up the tree. There was some discussion about the difference between aggregation and compression.

Bruce Greenblatt presented the ldap-01.txt draft. This is a tagged index object which can be exchanged by index servers. It would also allow directory servers to find information more quickly in their own large trees. It also allows sharing of index information among administrative domains. Attribute values are tagged to have the ability to tie the attribute-values back to the DIT. This change also allows a mechanism for allowing incremental updates by using ldif changes.

Chris Weider presented his draft: draft-weider-cip-hierarchy-00.txt He said this would move the find hierarchy today. This draft defines a centroid component necessary for hierarchical compression of data. He showed the template type and the rules. The prototype server is up and he is working with Michael Mealling to interoperate with a RWhois server. The Hierarchical data type helps give fewer false positives with no false negatives. The schema has a template, attributes and tokenization types. You can specify a right or left hierarchy.

Patrik ended the meeting with a final reminder that we need to make progress or the working group will be closed. Drafts should be published in May with working group last call in June.