NameFLOW-Paradise

Quarterly Service Report July - September 1996

TABLE OF CONTENTS

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
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.html

Operations

Introduction

This chapter summarises operations and related activities of the NameFLOW Paradise Directory service for the three monthsJuly, August, and September 1996.

1. Operations/Helpdesk

The most significant occurrence during the quarter was the removal of the country entry for Estonia/c=EE, and its FLDSA cn=Saki. Their X.500 pilot never developed, and it was decided, with the agreement of the relevant people in Estonia, that data quality was best served by deleting the entries.

The connection between JANET, the UK academic/research network, and DANTE's EuropaNET was upgraded from 4Mbs to 8Mbs. Since the root DSA Giant Tortoise and its associated information servers are hosted on JANET, this should give better access for European users.

2. Outages

Significant outages of service totalled approximately 6 hours in July; 2 hours in August; and 1 hour (scheduled) in September. Details of outages are available in the monthly reports.

3. Issues

There has been no resolution during the quarter of the IP-router problem preventing the full configuration of Malaysia/c=MY into the directory. The organisation involved is waiting for its Internet provider to sort out the router.

There are some changes to the top-level configuration expected in the near future: the FLDSAs in Australia/ c=AU, Anaconda and Bush Dog, are undergoing a reconfiguration, while the master DSA for the top-level organisation o=NATO is to be converted to an X.500(93) implementation.

A small probe utility, developed at TU-Delft in the Netherlands/c=NL, is expected to prove useful to NameFLOW-Paradise country managers. This probe can test not only that a DSA is available, but also that the organisation setup matches the "standard" as recommended in RFC 1617.

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

The fall-off in LDAP use over the Summer was not as great as expected, presumably due to increased use of the web gateway by a smaller number of people. Even more use can therefore be anticipated for the rest of the year and into 1997.

In the statistics for Giant Tortoise in appendix 2, the larger than normal numbers of local binds and "reads of DSA entries" in August correspond to testing of the probe utility mentioned in Issues.

5. Index Meeting at Brunel University

Notes: see Annex 6

Liaison

An EEMA Directory Committee Meeting was held on 17 October 1996 at the Wembley Conference Centre, London, United Kingdom

Draft minutes: see Annex 5

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 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
July 1996:          http://www.dante.net/np/report/mr/mr9607.html
August 1996:        http://www.dante.net/np/report/mr/mr9608.html
September 1996:     http://www.dante.net/np/report/mr/mr9609.html
```

The first edtion of the 2nd NP Annual Report was issued in October, see

```
http://www.dante.net/np/report/ar/ar95.html
```

APPENDIX 1 - Helpdesk summary for Jul/Aug/Sep 1996

	Country	Number of queries			
Full Name	ISO Code	July	August	September	Quarter
Australia	AU	–	1	2	3
Canada	CA	1	–	–	1
(China)	CN*	3	–	2	5
France	FR	1	–	–	1
United Kingdom	GB	2	1	1	4
Italy	IT	–	1	–	1
Korea	KR	1	–	–	1
Netherlands	NL	–	–	1	1
Norway	NO	–	1	1	2
New Zealand	NZ	–	–	1	1
Russia	RU	1	–	–	1
(Turkey)	TR*	1	–	–	1
(Ukraine)	UA*	5	–	1	6
United States	US	–	3	2	5
(Venezuela)	VE*	–	1	–	1
(South Africa)	ZA*	1	–	–	1
Total Requests		16	8	11	35

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

APPENDIX 2 - World Root DSA and LDAP summary statistics for Jul/Aug/Sep 1996

Summary of calls to DSA Giant Tortoise

From 0:05:04 on 30 June to 0:01:20 on 30 September

```
No. of binds                     July      August    September Quarter
_____

Local                             244        240         273       757
Remote                           5341       5667        5308     16316
_____

Total                            5585       5907        5581     17073



No. of operations                July      August    September Quarter
_____

Local                               6        270          26       302
Remote                          44605      47915       52142    144662
_____

Total                           44611      48185       52168    144964



System usage (calls received)    July      August    September Quarter
_____

Binds by Directory technicians   1485       1422         960      3867
Reads of DSA entries                6        221           7       234
Other ops on DSA entries            2          0           2         4
Getedb operations (inc slices)  44323      47643       52077    144043
Spot shadows                       15         24          25        64
_____

Total                           45831      49310       53071    148212
```

LDAP usage

```
LDAP usage from Jul 1 1996 to Sep 30 1996


                                 July      August    September Quarter
_____

Connections                     24333      20231       19953     64517
Total connect time (seconds)   291373     265927      397698    954998

(954998 seconds is 265 hrs 16 mins 38 secs)
```

APPENDIX 3 - Public DUA summary statistics for Jul/Aug/Sep 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	July	August	September	Quarter
Internet	1129	1031	1458	3618
UK academic X.25 (JANET)	2	1	7	10
EuropaNET X.25	0	0	4	4
Public X.25	1	2	4	7
Total	1132	1034	1473	3639

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

Domain	July	August	September	Quarter
edu	307	368	762	1437
unresolved	171	152	189	512
uk	146	128	145	419
com	110	72	59	241
net	53	42	49	144
dz	94	13*	19	126
org	26	26	26	78
ca	36	15	24	75
nl	21*	19	31	71
de	23	20	11*	54
Total	987	855	1315	3157

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

APPENDIX 4 - WWW/FTP/Gopher summary statistics for Jul/Aug/Sep 1996

WWW server

TOTALS FOR SUMMARY PERIOD Mon Jul 1 1996 TO Mon Sep 30 1996

	July	August	September	Quarter
Unique hosts	659	534	563	1635
Number of HTML requests	1375	1440	1498	4313
Number of non-HTML requests	237	224	127	588
Number of malformed requests	77	59	76	212
Total number of all requests/errors	1692	1723	1701	5116
Total number of Kbytes requested	37096	20708	16598	74403
Average requests/day	55.5	55.7	57.5	56.0
Kbytes/day	1216	669	561	813

FTP server

TOTALS FOR SUMMARY PERIOD Mon Jul 1 1996 TO Mon Sep 30 1996

	July	August	September	Quarter
Files Transmitted	418	625	376	473
Kbytes Transmitted	89825	152770	117983	120193
Average Files Daily	15	25	14	18
Average Kbytes Daily	3326	6110	4537	4658

Gopher server

Gopher usage from Tue Jul 2 1996 to Mon Sep 30 1996

	July	August	September	Quarter
Total connections	16	35	21	72
Total files retrieved	13	16	11	40

APPENDIX 5

EEMA Directory Committee Minutes

17 October 1996
Wembley Conference Centre, London, United Kingdom

ATTENDEES:

Nick Emery   Altavista   UK   NE

```
Stephen Vercella    Barclays                        UK         SV
Keith Reid          Barclays                        UK
Joanne Ghahremani   BT North America                USA        JG
Vincent Berkhout    DANTE (Project Manager)         UK         VB
Karl Heinz Bonacker Dr Materna GmbH                 Germany
Erik Andersen       Fisher & Lorenz                 Denmark    EA
Bernard Tardieu     Global One                      France     BT
Steve Kille         ISODE Consortium                UK
David Goodman       Lotus Development (Chairman) USA           DG
Bernd Stadler       Telia                           Sweden     BS
Willibrand Lederer  Siemens AG                      Germany
Rod Ward            IBM Global Network              UK         RW
Keith Richardson    ICL                             UK         KR
Amanda Edwards      MAFF                            UK         AE
Colin Robbins       NEXOR                           UK         CJR
John Wimpole        The Radicati Group Inc.         UK         JW
Sara Radicati       The Radicati Group Inc.         UK         SR
Peter Pawlita       Siemens Nixdorf                 Germany    PP
Magnus Andersson    Sweden Post Ltd                 Sweden
David Chadwick      University of Salford           UK         DC

Roger Dean          EEMA                            UK         RD
Richard Harwood     EEMA                            UK
Jane Jobson         EEMA                            UK

Apologies from:
A Barnes            ACB Associates
Paulo Malinverno    ISOCOR                          Ireland
Roger Molesworth    Logica                          UK
G Weckman           TeamWare Group
Leena Save          Telecom Finland                 Finland
Mark Ballan         Tele Denmark                    Denmark
Jeroen Houttuin     Unisource                       Switzerland
```

1. OPENING

Word of welcome by the chairman and agenda bashing. Minutes of
previous meeting accepted.

2. Liaisons/Education

EWOS (Erik Andersen and Keith Richardson)
Introduction to the way EWOS is changing. EWOS are currently
involved in three activities:
a. They are currently working on International Standardised Profiles
   (ISPs) for X.500 (93), and are planning to work on the 1997
   edition.
b. EWOS is working with ETSI and taken the lead in developing
   technical guidelines for High Level Naming (which involves more
   than just the directory). It is intended to include Internet
   naming and Domain names in support of a European infrastructure.
c. EWOS are also involved in Directory Deployment (EIDQ), which is a

replacement for the telephony-based TPH 28, and are working on a
common European schema. Documentation can be found at ftp://
ftp.ewos.be/pub/egdir.

Sweden (Bernd Stadler)
BS reported on two aspects of government activity in Sweden:
a.   The Swedish government is publishing a tender for electronic
trade at the end of December 1997, including the provision of X.500
and EDI. It was targeted that 70% of commercial transactions in the
government will be done electronically in the year 1998. The project
has been awarded to three organisations: Telia, Swedish Post/ENATA
and VMdata.

b.   The Swedish Government wanted one operator for addresses in
Sweden, but were opposed. There has been an offering process at the
end of 1995 and, Telia was chosen as one of three preferred suppliers
about May/June -1996 . The first customers will be "plugged in" by 1
January 1997. To the end of 1997 the service is completely
"detached" between Telia and the Government. A Swedish Directory
Forum was set up with the requirement to provide a telephone
catalogue and to provide a working service between competitors (e.g.
Global One, BT and SUNET) by the end of 1996. The Forum is there to
discuss cooperation, a common schema and content of the "mutual"
directory. The contract says that a minimal set of information must
be shared to have at least e-mail and security directories. The
suppliers which wants to be chosen has to "open up" its directory for
interconnection and perform searches for each customers, so that one
company needs to be a customer to one service provider. We have also
chosen to manage jointly the SE root out of the competitive situation
in Sweden. The Swedish Directory Forum has accepted this approach and
now we start to set up a common tree out of the recommendations.

c.   For Telia it is in a way to early to show the "TradeBase"
(registered trade mark) mid 1997 because we are in an evaluation and
development process. Telia could participate in a limited scale by
showing up the TradeBase application in Maastricht 1997. Telia is
interested to provide (parts) of the infrastructure for the Directory
Challenge. They have set up an ftp site (in English) [Details not
provided]

NameFLOW-Paradise (Vincent Berkhout)
The new NameFLOW-Paradise Annual report is available and presented to
the committee. (Get a free copy via
http://www.dante.net/np/report.html). VB gave a presentation on the
test results of the NameFLOW-Paradise X.500(93) test experiences in a
real operational environment. A complete overview of the NP-93 pilot
can be found at http://www.dante.net/np/93pilot.html. The intention
is to move the 777 DSAs connected to the service to the new
(X.500(93)) set of protocols. The next test will be made before the
end of 1996. Full migration is targeted for 1997.

EuroSInet (Peter Pawlita)

EuroSInet organised another workshop for X.500 (and X.400) implementors, hosted by Siemens Nixdorf in Munich. The number of participants (Data Connection, ISODE Consortium, ICL and SNI) was "acceptable", but could have been better. Some organisations participated remotely in the tests. The next workshop will be held in Brazil to get the South Americans involved as they are about to start similar activities.
ACTION VB: Distribute EuroSInet Press Release to the mailing list.

EMA/NADF et al (Joanne Ghahremani)
JG gave a brief word on the Challenge 97 which was discussed at much greater length in the break-out session.

3. Projects

Finished Projects
The project "Guidelines for the Deployment of a Corporate Directory" has been published as a Black Book (with a red cover) and in electronic format.   Hard copy is only available on request from the EEMA Office.

Ongoing Projects
The project, Top Level Naming In Europe (TOPOL) - Phase 1, will be finished before the end of the year. Agreement has been reached with Logica on the outstanding items.

Proposed Project: Understanding LDAP (David Goodman/Colin Robbins)
CR has drafted an initial document "Understanding LDAP" which will be edited by DG. The objective is to put LDAP in perspective to address some of the misconceptions around what LDAP can and canÕt do. LDAPv3 which is currently being designed through an IETF committee is far from a "done deal" and is significantly more complex than the current version LDAPv2. Another aspect of the document is to generate a discussion on the consequences of an LDAP-based global directory on naming issues, an topic which is seen to be outside the scope of the current TOPOL document. It was agreed that EEMA would look into a dedicated mailing list for discussion, but for the time being the general dircomm@eema.co.uk list will be used.

Proposed Project: Business Drivers for Directories (Stephen Vercella)
The terms of reference for the project was tabled. Background information: Barclays has about 90.000 users in numerous autonomous units. So far promoting directories has been very much a push rather than a pull. To promote directories requires business drivers. The project will try to answer important frequently asked questions such as: how to build a directory when no funds are available? (Answer: seems to rest on volunteers operating directories on a best effort basis.) How to organise a corporate versus a public directory?

SV argued that a net cost/benefit analysis is not enough.

It was agreed that the project should be two part:
i   some generic guidelines and analysis
ii   case studies (Those by Telstra, the MAFF and SNI were mentioned)
The project would look to the user committee for support in
understanding the issues as well as possibly some case studies.
Suppliers could also be approached for similar reasons. SR offered
to help SV with the project once it was more clearly defined.
ACTION SV: Definition of the project as discussed, together
with budget proposal

Proposed Project: Top Level Naming In Europe (TOPOL) - Phase 2
This was not discussed as it is intended to make this an important
aspect of the work of the WDF (see below).

Proposed Project: World Directory Forum - Europe
DG reported that it had been agreed at the EEMA Board meeting that RD
would submit a proposal for the adoption of a World Directory Forum
as a WEMA activity at the next WEMA Board meeting in Japan at the end
of October.

Proposed Project: EEMA Directory Challenge 97
Presentations by DG and JG. [Sheets available via .... or text
appended.]

The Challenge 97 is not a US effort but a world effort. The emphasis
is on cooperation between the *EMA's making it a true WEMA event. DG
announced that a proposal was going to be presented to the WEMA Board
at the end of October to endorse the Challenge 97 as a WEMA project.
CR started the discussion by saying that the Challenge 97 objective
had moved from its original goal, which was to have demonstrated
secure communications between *EMA's member organisations and thereby
raised the need for a secure directoriesÕ infrastructure. The
Challenge 97 now goes far beyond secure communication. One of the
problems is that a Directory Challenge is not interesting enough for
marketing departments which can sell secure communications but not
directories on their own. Directories are only a tool.

There will shortly be a press release, which will include a list of
participants. Being mentioned in the initial press release will be
free of charge but after the first there will be a charge of about
US$ 5,000.

For EEMA members who wish to participate in the Challenge but need to
convince their management to put some effort in, JG kindly offered
her help.

It was agreed that LDAP servers would participate as long as they
could be visible in the X.500 infrastructure.

CR was willing to consider taking the role of challenge project
leader. He will define four-five roles for the EEMA part of the
challenge. The suggested roles were:

```
  Project coordinator
  Directory Infrastructure coordinator
  Secure applications
  Demonstration manager
  Marketing coordinator
```

The following people want to be involved in further development of
the European part of the Challenge: CR, BS, DG, VB, AE, JW, EA, JG,
BT, PP, NE, DC, MA and KR. RW was interested longer term.

Further discussion and work will be done by this group at a meeting
proposed for 4 December in Amsterdam. A core group will meet to work
out a proposal on the 3 December. Members attending the pre-meeting:
CR, DG, VB, JG and JW.

4. Next Committee Meeting
The next committee meeting will be in Berlin on the 25 or 26 February
1997 (to be confirmed).

5. AOB
None.

APPENDIX 6

```
Trip note:   Index Meeting at Brunel University
Date:        21 October 1996
Time:        10h00 - 16h00
Place:       University of Brunel, United Kingdom
```

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

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

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

RH continued after this. RH is co-chair of the FIND working group at
the IETF. David Chadwick describes a solution and he actually
implemented something before writing it down. RH has a strong Whois++
background and also built a Ph->X500 gateway. His index server works
for all three directory types and is based on a SLAPD server. He
build a special web client for index querying. He uses SLAPD to look
at the performance with about 90.000 -100.000 entries (indices). The
current source is in PERL but he has plans to rewrite it in C.
The Swedish government is forcing "service providers" to set up a
public directory and if they fail there will be legislation for this.
They believe that index servers can be used as "common solution"
between different directory types. This could also be a solution for
the Multiple Service Provider problem (see DIP #13).

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

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

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

RH continued as people were interested what he had actually built. A
rather technical discussion followed:
each index has a dSI, a dataSetIdentifier. Within the dSI there is a
baseURI with a URL referring to a Whois++/Ph/LDAP server.
   dSI= Index for Sweden
   baseURI= ldap://ldap.umea.se/c=SE
The index is stored in a SLAPD server. Underneath the country node
there is one entry called dSI=Index for Sweden. Under this node

there is a complete one-level list with all indices. A typical index
would be:
  idx = "Roland Hedberg+user+name"
  dSI = "LDAP"
  dSI = "Whois++"
  baseURI = "ldap://some.server/some/base/reference"
  description = language:swedish
Each index like this is one entry.

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

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

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

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

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

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

What next?
VB: Contact UKERNA to see if they are interested in participation.
VB: Distribute URLs for documents
    (http://www.dante.net/pubs/dip.html)
RH: Work on Internet Draft for indexing and this will be used for
    further discussion in this group.
PJ. needs a search engine as aprt of the DESIRE project and will
    invite Roland to work on this (?)
VB: set up mailing list for this group?