

# NameFLOW-Paradise



## Annual Report 1994-1995

---

This report was compiled by Vincent Berkhout with the help of many others, whose contributions are gratefully acknowledged. Parts of the report may be freely copied as long as the source is acknowledged.

For more information please contact:

DANTE  
Lockton House  
Clarendon Road  
Cambridge  
CB2 2 BH  
United Kingdom

Tel: +44 1223 302992

Fax: +44 1223 303 005

E-mail: [V.Berkhout@dante.org.uk](mailto:V.Berkhout@dante.org.uk)

---

## Table of Contents

### [Foreword](#)

### [1. Introduction](#)

## **2. Directory facts and figures**

## **3. Operations report**

## **4. Development Activities**

## **5. Liaisons**

## **6. Promotion and publicity**

### **Annexes**

Annex 1:

[Public directory interfaces](#)

Annex 2:

[National contacts](#)

[Liaison contacts](#)

Annex 3:

[Glossary](#)

---

## **Foreword**

The last year has been turbulent for R&D directory services. The PARADISE project concluded in May 1994 and since then DANTE has been offering the service on a commercial basis. The good news is that NameFLOW-Paradise has proven to be commercially viable and is now 'up and running' without external funding.

'The bigger the better' is the first goal of NameFLOW-Paradise with the objective to be (part of) the 'universal' directory interconnecting directory services all over Europe and the rest of the world. Although many countries, universities and other organisations are already participating, the directory information base could be expanded much further.

'...but better is best', the second goal of NameFLOW-Paradise stresses the fact that quality is as vital for the success of a directory service as quantity. Quality of Service measurements should reflect the users' perception of the service as a whole and not just the system availability. The quality of directory services concerns a range of issues, such as directory performance and availability, data registration, data accuracy and data coverage. A good Quality of Service is not only useful: in most countries it is a legal obligation that distributed personal data are accurate.

A lot of work is still to be done but substantial progress can be reported as well, in particular since January 1995 when Vincent Berkhout started work for DANTE in the directories area. Expert advice has been provided by consultants Paul Barker, David Chadwick, and Colin Robbins.

DANTE will take on the further development of NameFLOW-Paradise in cooperation with the national directory managers and the operational staff at ULCC with the aim to establish the service even more firmly in the year to come.

[\[CONTENTS\]](#)

## **Introduction**

This first NameFLOW-Paradise Annual Report presents an overview of developments and activities between May 1994 and May 1995 and takes a look into the future as well. It covers daily operational issues as well as planned changes and developments for next year. The chapter on Liaisons gives an overview of the activities of organisations which are important for the future of NameFLOW-Paradise.

The PARADISE pilot project had already proved that X.500 works. The migration to the operational service now called NameFLOW-Paradise was both a technical and an organisational challenge. The technical challenge ahead is to build an international directory service based on products truly conformant to the X.500 standard and thereby building an infrastructure open for others to join. Directory software conforming to the most recent edition (1993) of the X.500 standard is rapidly becoming available: the first packages are on the shelf and many are about to follow.

From an organisational point of view the NameFLOW-Paradise directory should not be exclusive to the academic and research community. Commercial companies should be allowed to join at least as long as no alternative global public directory is available. Awareness of the benefits of directory services is no longer restricted to the research community. The first non-R&D companies are declaring an interest in the development and deployment of directories as a short term strategic goal.

In May 1995 the European Commission (DGIII) organised a conference in Brussels on the subject of electronic directories. The conference was attended by a large number of delegates from European Administrations and Institutions keen to adopt X.500. However, directory projects funded by the EC are mostly pilot projects and there is limited experience with (setting up) a large scale operational directory service such as NameFLOW-Paradise. It is in the interest of both the research community and governments to establish close cooperation towards the provision of a common directory.

Some facts: the directory currently includes information on over a million people in thousands of organisations, actually 4500, in tens of countries, actually 38, provided via 750 DSAs. The trend is a steady annual growth of 40% in the number of participating organisations and directory systems.

In the coming year a number of important changes in the organisation and provision of NameFLOW-Paradise are going to take place. The future of directory services looks promising though, and


NameFLOW- Paradise will definitely have an important role to play.

[\[CONTENTS\]](#)

## Directory facts and figures


The directory has been growing steadily over the past five years. In May 1995 4500 organisations participated. Commercial organisations have contributed the largest part of the growth in the number of organisations in Europe during the last year.

**Number of Organisations, European and non-European**



The number of DSAs in Europe shows the same growth trend - Europe is still ahead. In May 1995 750 DSAs were part of the directory.


### Number of Directory System Agents (DSAs), European and non-European



### Directory Enquiry (DE) service, a popular feature


The number of DE logins over the last five years shows the following trend:

Number of DE logins - Oct'91 - June'95



The graph below shows the annual growth in the number of DE logins. Between June 1994 and June 1995 the growth was still nearly 60%.

**Number of DE logins per year between October 1991-  
June 1995**




## Improving the quality

'The bigger the better...but better is best'

DANTE aims to improve the quality of directory services. But what should be measured and how? System availability is one of the quality elements that currently can be measured. The availability of first level DSAs is measured by active probing. Between January and May 1995 54 first level DSAs were probed over 10,000 times each. It turns out that most DSAs have a very high availability rate (see diagram below).

**Availability of first level DSAs**



[\[CONTENTS\]](#)

## 3. Operations Report - Ronan Flood ULCC

DANTE has contracted ULCC, the University of London Computer Centre, to operate the NameFLOW-Paradise Service, as well as an ftp/gopher information server.

The operational services comprise:

- The Directory System Agent (DSA) Giant Tortoise, which acts as the world root of the NameFLOW-Paradise Directory, and provides a replication service for root and country-level information;
- A public-access Directory User Agent, Directory Enquiries, and its supporting DSA, Ocellated Turkey;
- A public-access Lightweight Directory Access Protocol (LDAP) server, supported by Ocellated Turkey, which also offers public DAP access;
- An information server, accessible by electronic mail, and through Internet Gopher and FTP clients;
- A mailing-list expander, which hosts two lists for discussion of the service (forum@nameflow.dante.net) and coordination among national DSA managers (managers@nameflow.dante.net);
- The NameFLOW-Paradise helpdesk, providing help and information to current and potential national managers and end users of the service.

### *Hardware/software*

The services run on two dedicated Sun Microsystems computers: 'Mason', a 4/330 with 32 Mb of memory; and 'Fortnum', a 128 Mb SPARCserver-10. Both machines are connected to the UK academic IP/X.25 network, to DANTE's EuropaNET, and to the UK public X.25 network.

The network software used is SunOS TCP/IP and SunNet X.25, with OSI and X.500 facilities provided by XT-ISODE, XT-Quipu, and XT-LDAP from NEXOR Ltd of Nottingham. The Directory Enquiries interface is 'DE', developed at University College London as part of the PARADISE pilot project.

### *Personnel*

Supporting the NameFLOW-Paradise services is substantially a full-time job for one person. This role is taken on by staff in the Networks Group at ULCC. At the end of September 1994 Linda Millington, who had been providing the support function almost since the inception of the PARADISE project, left ULCC to further her career at CDC. Until a new member of staff could be found and brought up to speed, NameFLOW-Paradise was supported by Graeme Hairsine (responsible for UK national and international electronic mail relays) and John Seymour (manager of the Networks Group). At the start of January 1995, Ronan Flood and Stan Grant were appointed to join Graeme.

Gradually over the succeeding months, Ronan has taken on the tasks of NameFLOW-Paradise support, while Stan has been shadowing Graeme on mail relay support. It is expected that, in time, each team member will be able to cover another's role in that person's absence.

### *Activities*

Running the NameFLOW-Paradise service entails:

- Ensuring that the computers and software servers are kept available;
- Answering requests to the helpdesk, largely by e-mail;
- Modifying the operation of the services, at the request of, or in consultation with, DANTE;
- Where necessary, installing new versions of software, or providing new services, such as the Gopher server and updated FTP server;
- Informing national managers of scheduled downtime for the services, and of known interruptions of service for national DSAs;
- Producing reports for DANTE on the operation of the services.

Like most complex software systems, Quipu DSAs have their problems. It is possible to recover automatically from some of these problematic conditions by means of 'watchdog' programs. Where automatic recovery is not possible, the watchdog should alert support staff that manual intervention is required. In identifiable, reproducible circumstances, NEXOR has been able to supply updated versions of XT-Quipu to fix particular problems.

### *Helpdesk*

Queries to the helpdesk continue to range from the straightforward, through the complex, to the desperate and the frankly bizarre. Simple requests of the form 'Please send some information' (or just 'help'!) are answered with a copy of DANTE's online brochure and pointers to the email/Gopher/FTP information servers. The stream of requests for help in finding e-mail addresses of people in organisations which do not appear in NameFLOW-Paradise simply emphasises the need for White Pages services in general, and for publicising them more widely. People e-mail the helpdesk because they have found someone, somewhere who appears to offer to help them - we try not to disappoint.

---

Tales from the helpdesk One man wanted help 'to find the e-mail address of a gal I met at a party last weekend in Dallas, who works for ...'; we supplied the address of the general postmaster at the relevant company. A particularly surreal and baffling query in March 1995 reported 'Every time I enter the network and try to find a name, all I get is the name of fish...' - we never did clear up that one.

---

Helping a new country manager to configure a DSA, join the Directory Information Tree, and set up replication, is an ongoing job which can last for several weeks.

ULCC provides DANTE with monthly and quarterly operational reports. The reports contain the



following elements: general overview of operational/helpdesk issues, outages, development issues, and statistics. An html version of the monthly reports are included on DANTE's NameFLOW-Paradise web pages (URL: <http://www.dante.net/nameflow.html>).

*Future developments (see also Development activities)*

Since April 1995, a copy of the root naming context has been made available for anonymous FTP from the NameFLOW-Paradise information server (referred to as out-of-band replication). This is updated on a daily basis, if required by a change in the data. The root naming context appears in two formats: as a Quipu EDB file, and as a pure knowledge reference file. The latter is produced from data on Giant Tortoise using software developed by NEXOR Ltd and donated to NameFLOW-Paradise.

In the summer and autumn of 1995 it is planned to install software from NEXOR to provide replication of the root naming context using the X.500(93) standard Directory Information Shadowing Protocol (DISP). At the start of 1996, Quipu replication will be withdrawn and Giant Tortoise will cease operation. The remaining services are then likely to be consolidated onto one computer.

Currently, Giant Tortoise provides replication not only for the root naming context, but also for the entries at the next level. Some of that information is necessary to produce the complete knowledge reference file for FTP. It is not yet clear whether this country- level information will continue to be gathered and made available centrally via DISP.

[\[CONTENTS\]](#)

## **4. Development Activities**

*Transition to X.500(93)*

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

In April 1995 DANTE announced the NameFLOW-Paradise participants' decision to launch a pilot to test the 1993 edition of the X.500 standard. This international pilot will test X.500(93) implementations at various levels. The short term goal of the pilot being to gain experience with 1993 compliant implementations to pave the way for a full scale X.500(93) service next year, open to many software vendors. The improved 1993 X.500 standard incorporates the advantages of RFC1276 and adds necessary features to the 1988 edition, such as replication, security, and access control.

The pilot comprises two phases. The first phase - which has already started - is a small pilot covering three European countries, where 93 conformant services will be run in parallel with the current Quipu based service. The initial organisations involved are DANTE, UKERNA (United Kingdom Education

and Research Networking Association) who will test the software in their organisation, SURFnet (Dutch national research network) and SWITCH (Swiss national research network) who will be performing tests nationally, while ULCC will test a top level (international) service. The ISODE Consortium and NEXOR were the first two vendors to make their X.500(93) products available to the pilot participants.

The second phase will be a large scale pilot covering most of the European countries and possibly some countries outside Europe, and is planned for October-December 1995. During 1996 the pilot structure will gradually replace the current operational directory service. It is foreseen that the deployment of a full 1993 compliant Directory Service may need several operational enhancements, as it has not been tested on a large scale before. As a result the pilot will have a major impact on the current structure of the service.

### *Managing the root*

The current Quipu model features a root DSA managing the top of the Directory tree. The root DSA offers amongst others the following services:

1. resolving queries.
2. relaying queries between different network domains.
3. managing and distributing knowledge information.
4. replication of the top level entries.

However, the root DSA is not a feature of the X.500 standard (1988 nor 1993 version). The first and the second service elements have been taken over by the national level. With respect to the third and fourth service elements the most practical solution is the inclusion of a root DSA function in the 1993 implementation.

### *Resolving and relaying queries*

Due to historical reasons the root DSA could be accessed arbitrarily by Directory User Agents (DUAs) all over the world. This meant that the central root DSA had to resolve many queries for many DUAs. As a result the performance of the root DSA degraded significantly. Resolving queries is not among the mandatory functions of the central root DSA, as user queries can be resolved just as well by country level DSAs. After consulting the national directory managers it was decided to distribute the functionality of resolving user queries among all the country DSAs. The access policy to the root DSA has changed from 'access for all' to a 'need for access' whereby user access to the root DSA is restricted to necessary management functions only.

Relaying queries between different network domains is another function provided by the root DSA. A User Agent in an X.25 network domain should be able to retrieve information from a DSA situated in an IP network. While the User Agent and the DSA are in different network domains, an intermediate DSA could relay the query from one domain to another. From a management perspective the easiest way to achieve this is via a centrally located DSA that is connected to many different networks: the root DSA.

In the process of analysing the problem that almost every DSA was using the root DSA as relay, it became clear that all country DSAs are accessible via Internet and the idea of an X.500 backbone between countries based on IP-only was born. This requires that every country DSA has to support all local network domains. In practice the model, whereby a few country DSAs are kindly relaying for other countries, works well. An advantage of the current model is that the X.25 domain is no longer needed which has a positive financial side effect\*.

\* UKERNA is gratefully acknowledged for providing an X.25 access to the root DSA without charge  
*Knowledge references and replication*

How does a 'random' DSA know how to contact another country DSA and find the country information? The manager of the new (or changed) first level DSA could contact all national directory managers (e.g. via the managers mailing list) and ask them if they are willing to add the information. This one-to-many approach is extremely labour intensive and error prone. A better way is to distribute this knowledge information via the central root DSA. The root DSA holds references to all country and DSA entries. Each country (first level) DSA shadows this information from the root DSA. Each first level DSA only needs to have one bi-lateral agreement, between itself and the root DSA. This agreement ensures that the first level DSA keeps the root DSA up to date with its country level information, and in turn, the root DSA keeps the first level DSA up to date. When a new first level DSA comes on line, it only needs to establish a bi-lateral agreement with the root DSA in order to obtain information of all first level DSAs. This root DSA model is a much easier configuration to manage than simply a set of first level DSAs without a root DSA.

The NameFLOW-Paradise Directory is distributed between many countries (38) and links together many national directory services. These national services are all provided by national organisations and linking them together on a many-to-many basis will be unmanageable, from a technical as well as an organisational point of view. Having only one bi-lateral agreement with the root is a much better solution. This 'Quipu root feature' is highly desirable and must be maintained in the times to come.

Beside the knowledge references to other first level DSAs the root DSA replicates all the country entries. This replication (or 'shadowing') is currently possible as the Quipu DSAs implement the replication mechanism specified in RFC 1276. Replicating the complete first level has a positive impact on performance outweighing the needed resources for replication and the dangers of data latency. A query to find 'object=country AND Name=Utopia' would result in an avalanche of queries to all known first level DSAs. If all first level information was stored locally the conclusion could be that Utopia is not (yet) stored in the Directory!

The 1993 edition of the X.500 standard includes a Directory Information Shadowing Protocol (DISP), a replication mechanism comparable to that deployed in the Quipu model. The replication mechanism has proved to be very useful in improving the performance of the Directory and replication will be an important part of the new 1993 infrastructure.

*Index Servers*

One of the major benefits of the X.500 directory is the scalability, allowing a globally distributed database. However the current X.500 implementation cannot resolve all queries or is simply too slow. A white page search for a person without knowledge of the country or organisation name is currently unfeasible because the directory as a whole has to be interrogated. A special type of directory server is needed to solve these type of queries. Paul Barker addresses this issue in his discussion paper (see page 21) produced on behalf of DANTE proposing indexed DSAs. The next step will be to design and implement an index DSA to test its operation.

### *Non-X.500 Directories*

As X.500 is not the only directory protocol, ways for co-existing and inter-working with other systems, such as Whois++ are being examined. Whois++ is a new IP directory services protocol which started to evolve in mid -1992 and is based on the Whois protocol. It was originally designed to be an extension to the Whois protocol, but the functionality increased and backwards compatibility with Whois was no longer possible. One of the key differences between X.500 and Whois++ is the naming structure: in X.500 it is based on a hierarchical tree, while Whois++ is based on a space of index servers, called an index mesh. In X.500 the geographical location and organisation names are used as prefix and are (partially) needed to resolve the query and returning a so-called distinguished name. Whois++ does not use a distinguished name but uses a token as unique identifier. The WHOIS++ architecture and technology seems to stabilise and will become an Internet Standard at the end of 1995. The first software, called Digger (Bunyip), is now available and the number of Whois servers is growing rapidly.

### *Locality Europe*

It is unclear who should manage locality Europe (if needed at all), and who would be allowed to register under it. The management of the node locality=Europe will partly depend on the outcome of the EEMA TOPOL-project. DANTE will ensure that the interests of NameFLOW- Paradise will be taken into account in the project.

### *Quality of Service and Service Level Agreements*

One of the goals of DANTE is to improve the quality of the directory service offered. As part of a directory management framework Service Level Agreements between 'participants' are needed. One of the deficiencies is that currently the measured quality level is based on DSA availability, being only a part of the total directory system. Adequate management of the directory service requires monitoring and controlling the quality of the total system. The requirements for monitoring the quality of the service have been defined, however, measuring the system against these requirements requires tools. DANTE will work on an initial software specification for development of these QoS-tools and consider further development.

[\[CONTENTS\]](#)

## **5. Liaisons**

## EEMA - Directory Committee

The Directory Committee of the EEMA (European Electronic Messaging Association) is working on several important projects. The on-going projects within the EEMA are:

1. X.500 Product Guide.

The second version of the X.500/Directory Synchronisation Product Guide will be produced. A revised version of the questionnaire used for the first version will be sent to the vendors. The scope of this second report will be more global, in terms of vendors as well as intended audience.

2. Guidelines for Top Level Naming in Europe.

The objective of the project is to publish a document which will provide the basis for common agreements on naming and DIT structuring across European countries in order to achieve a single European namespace. The first phase of the project involves Issues and Requirements related to structuring of the Directory Information Tree (DIT). The second phase will examine all existing pilots, and existing services such as NameFLOW-Paradise.

3. Guideline for Corporate Directory Deployment.

The project will identify and provide guidance upon the issues associated with the migration of a corporate directory to an X.500 based Directory.

4. EEMA WWW Service.

The server will provide access to all EEMA's documentation. An important objective of this project is the establishment of a Web- X.500 Gateway.

5. EEMA X.500 Service.

The EEMA has a 'corporate' X.500 service with information on EEMA and its members. Access below organisational level will be restricted to members only, meaning that non-members can only access organisational information.

The EEMA Directory Committee showed at the annual EEMA Conference in June 1995 in Amsterdam that X.500 electronic directories are real and that they work. The interworking demo given at the conference showed a multi-vendor directory based on the products of ten vendors. An overview of all companies providing X.500 products is given in the EEMA X.500 Product Guide.

DANTE had a booth at the exhibition of the Annual Conference to promote NameFLOW-Paradise, and to analyse the need for and development of directories outside the research community. A frequently asked question was: "How can I connect my company to the global directory?" The answer was that the only operational public directory at the moment is NameFLOW-Paradise. The need for a commercial international public directory service is clear but it will take time for such a service to develop and to reach the size and quality of NameFLOW-Paradise.

Liaising with EEMA is important because it could facilitate closer cooperation between NameFLOW-Paradise and commercial service providers. In addition it could offer the opportunity to provide a

service to commercial companies.

## **EURESCOM**

EURESCOM (European Institute for Research and Strategic Studies in Telecommunications) is a collaboration between European Public Network Operators to perform research and development in telecommunications. Coordination and cooperation in the area of X.500 is perceived as vital for the future development of a pan-European X.500 infrastructure, in particular with NADF in the US and NameFLOW- Paradise.

The above objectives led to two projects:

- Pan European Directory Service - Phase II - pilot project (P309) - July 1993-March 1995;
- Europe-wide Directory Services - pilot phase (P416) January 1995 - April 1996.

The main objective of these projects is to establish a prototype of the future X.500 directory. From a technical point of view such a prototype directory could pave the way to the provision of future X.500 services by permitting:

- operational/performance evaluation of a real heterogeneous directory;
- test/validation of X.500 in near operational configuration;
- test/validation of possibly national - specific X.500 architecture with emphasis on interworking equipment such as X.500 first level DSA or X.500 gateway to existing electronic directory.

From an awareness point of view the pilot could serve the purpose of:

- promoting X.500 awareness among PNOs, private domains and private users;
- demonstrating PNOs capability and willingness to provide operational X.500 services;
- providing one of the necessary building blocks for developing future European teleservices;
- providing a representative forum for analysing and solving technical and non-technical problems of the directory.

Eurescom's X.500 activity is important for NameFLOW-Paradise because the PNOs have the potential to provide one of the largest directories possible, by storing their telephone directories in the X.500 directory. Fourteen major European PNOs are represented in the pilot project. Eurescom's activities will be monitored closely by DANTE and cooperation will be sought where and whenever possible.

## **EWOS - EGDIR**

### **Directory Expert Group - David Chadwick**

The European Workshop for Open Systems (EWOS) was established in 1987 with the aim of being an open forum to discuss OSI interworking issues, and to produce documentation that would enable different OSI implementations to interwork. One of EWOS's original objectives was to produce

functional profiles of OSI International Standards. The base International Standards contain many optional features and different levels of functionality, with the consequence that interworking between different implementations is often impossible to achieve without predefining the exact functionality of an implementation.

Another EWOS objective on the road to interworking is the definition of conformance test suites. Test suites, when implemented by conformance test centres such the NCC in the UK or CSELT in Italy, determine if a particular implementation correctly obeys the directives of the relevant functional profile(s).

### *Collaboration*

EWOS has around 20 expert groups, covering the whole range of OSI standards from lower layer protocols to MHS, EDI and X.500. EGDIR is the EWOS expert group responsible for profiling and related work in the X.500 area. It started meeting in early 1989 and has collaborated with its sister group, ETSI TE.6, since inception. ETSI - the European Telecommunications Standards Institute - has a similar though broader remit to EWOS, and was formed by the European members of ITU-T. TE.6 is the ETSI expert group responsible for X.500 interworking and service related issues, and it makes eminent sense that TE.6 and EGDIR should work together - otherwise we might end up with conflicting X.500 profiles, which is clearly in no-one's interest.

### *Scope*

Many, though not all, of the functional profiles for the 1988 edition of the X.500 standard have been completed long since by EWOS. However, in the world of international computer communications, purely European focused profiles are not sufficient for a global marketplace. Consequently this has led to the creation of International Standardised Profiles (ISPs), which are the amalgamation of the profiles originally created by EWOS, AOW (Asia- Oceania Workshop) and OIW (the US based OSE Implementor's Workshop). Table 1 lists the ISPs for the 1988 edition of X.500, cross referencing these to the original EWOS publications where appropriate.

The X.500 profiles fall into two general categories: profiles that control the functionality and procedures of the protocol exchanges between X.500 components, and profiles which govern and define X.500 schema elements e.g. object classes and attribute types. The first category is the most extensive, containing profiles for accessing the directory by DUAs, interworking between DSAs, procedures for distributed operations and the use of strong authentication. The second category contains schema definitions and limitations for general (or common) use of the directory, and schema definitions for particular applications' use of the directory (e.g. MHS). In addition to the functional profiles, EGDIR has also produced a series of technical guides, denoted by the assignment ETG in Table 1. Technical Guides are not European Standards as such, but rather, as their name suggests, give guidance to implementors and administrators of the most appropriate way of tackling various technical issues that may cause interworking problems.

Whilst not all of the 1988 profiles are yet completed (those marked draft in Table 1), work is now simultaneously under way to produce the profiles for the 1993 edition of X.500 (which incidently

was itself only recently published!). Table 2 lists the planned profiles for the 1993 edition of X.500. Given the added complexities of the 1993 base standard, plus the additional functionalities of shadowing, operational bindings, and access controls, one could be excused for thinking that EWOS would have its work cut out well into the next century before all of these documents were completed. However, help is at hand. The three regional workshops, OIW, AOW and EWOS are now collaborating effectively and sharing the profiling work between themselves, thereby reducing their individual workloads.

## **IETF - ASID - Tim Howes**

The Access, Searching, and Indexing of Directories (ASID - formerly Access and Synchronization of Internet Directories) working group of the IETF (Internet Engineering Task Force) was formed in the wake of the realisation that a single directory service is unlikely to take over the Internet. Instead, ASID's goal is to provide an even playing field for multiple protocol development, in the hope that the strongest and best protocols will be selected naturally and will be able to work together to provide a harmonized hybrid directory service system for the Internet.

ASID is currently developing the following directory service protocols:

- LDAP, the lightweight directory access protocol, both a front-end to X.500 and a stand-alone service. Work on LDAPv2 has begun.
- Internet X.500, including defining a transition path to X.500(93), defining new object classes and attribute types for various purposes, and defining algorithms such as UFN.
- WHOIS++, a relatively new directory service based on a text oriented protocol and a distributed indexing model. Currently at proposed standard status.
- SOLO, the simple object lookup service, another new directory service drawing parts from WHOIS++, LDAP, and X.500. Currently at Internet draft status.

In addition, ASID is developing the common indexing protocol (CIP), which will allow directory services to exchange distributed indexing information which can be used to perform efficient wide-area searches. CIP is currently at Internet draft status.

mailing list subscription:

*ietf-aside-request@umich.edu*

## **IETF - IDS**

The IETF has a second X.500-related working group, the Integrated Directory Services (IDS) Working Group. The objective of this group is to facilitate the integration and interoperability of current and future directories (WHOIS++, X.500, CCSO) into a unified Internet directory service.

In addition to specifying technical requirements for the integration, the IDS Working Group also



addresses administrative and maintenance issues by publishing guidelines for users and administrators on directory data integrity, maintenance, security, schema, and privacy and legal issues.

Current activities:

- Track emerging directory service protocols in order to identify the need for specifying standards for interworking with other service protocols;
- Liaise with groups working on deployment and development of directory services to locate and fix interoperability problems;
- Identify unfilled needs of directory service offerers, administrators, and users.

The following Internet Pilot Projects reporting to this group are:

- The Long Bud Project: Internet Pilot Project for the Deployment of X.500 Directory Information in Support of X.400 Routing (RFC 1802);
- The Internet Nomenclator Pilot Project;
- The Internet WHOIS++ Pilot Project;
- The Schema Registry project - identifying and publishing X.500 schema elements used on the Internet;
- Deployment of the X.500(93) Directory System on the Internet, in cooperation with NameFLOW-Paradise.

IDS is a very productive working group which has already published two RFCs (RFC1803 Recommendation for an X.500 Production Directory Service and RFC1632 A Revised Catalog of Available X.500 Implementations.) Currently the group has submitted three Internet- Drafts and another 13 drafts are planned to be submitted before the end of this year.

mailing list subscription:

*ietf-ids-request@umich.edu*

## **NADF - Ted Myer**

NADF (North American Directory Forum) is a group of service providers and users who joined forces in early 1990 to build an X.500-based public directory infrastructure for North America. Our principal activities are continued development of the technical foundations and operating plans for the directory, and the operation of an experimental pilot directory.

We view the pilot, which began operation in late 1992, as evolving continuously into the full scale commercial service that constitutes our mission. As we work with the pilot, we have been experimenting with different applications, initially straight white pages style listings focused on E-mail, during the last year additional experimental listings in support of:

- Public health service information. These are listings of information sources organized by topic.

- EDI/Electronic Commerce. These listings provide EDI capabilities of organisations organised by business category.
- Residential telephone listings. We are experimenting with naming, DIT design, and entry content issues. The test subjects are two communities whose listings will appear at the lowest geographic level of the DIT.
- Toll free telephone directory listings. This will be an X.500 replica of a directory currently maintained by one of the US service providers.

Although the main focus of NADF is on North America, we view this as just one part of the emerging global directory. Accordingly we maintain close liaison with related efforts such as NameFLOW- Paradise, and Eurescom. We have coordinated with the Canadian and US branches of NameFLOW to create a single namespace, and are working to physically merge the NADF and NameFLOW DITs for c=CA and c=US. When that effort is complete, the NADF pilot listing will appear in the North American portion of the NameFLOW-Paradise DIT. The past year US work on technical development has focused on developments of a model and standards for accounting and settlements in a public directory environment. The resulting document (NADF Document SD-13) was submitted to the ITU and the model it contains is reflected in the current ITU draft documents on this subject. The NADF standing documents are available at <ftp://ftp.gte.com/pub/nadf/nadf-docs>

#### **Current NADF members include:**

- Advantis
- Arinc
- AT&T
- Bankers Trust
- Bell Canada
- BT North America
- Canada Post
- Digital Equipment Corp.
- GE Information Services
- GTE
- Hitachi Computer Products
- MCI Communications Corp
- Novell
- Performance Systems International
- Premenos Corp.
- Sprint International
- SITA
- Texas Instruments
- US Postal Service

[\[CONTENTS\]](#)

## 6. Promotion and publicity

In the area of promotion and publicity a number of activities have been carried out/developed in the past year.

### *NameFLOW-Paradise WWW server*

In addition to the ftp/gopher information server maintained at ULCC since January 1995 a WWW server has been in place at DANTE. The web pages contain general information on NameFLOW-Paradise, contact details of the National Directory Managers, an overview of and links to public directory interfaces, the operational reports prepared by ULCC, and meeting information.

### *NameFLOW-Paradise publications*

Papers\* on the following development topics were produced by external advisors Paul Barker, David Chadwick, and Colin Robbins during the first months of 1995:

Paul Barker:

'X.500 Index DSAs'  
(DANTE IN PRINT, No.12)

David Chadwick :

'Phasing Out/Opening Up the Root DSA'  
'Shadowing First Level DSAs'  
'Multiple Service Providers - Various Solutions'  
(DANTE IN PRINT, No.13)

Colin Robbins:

'Monitoring Quality of Service in a Global Information Service' (DANTE IN PRINT, No.14)

### *Production of publicity material*

In addition some documents were prepared for publicity purposes

- 'NameFLOW-Paradise launches X.500(93) Pilot', press release, 4 May 1995, on 93 pilot (and announcement in news bulletin 'The Works of DANTE');
- Introductory brochure for general audience "Find and be found in NameFLOW-Paradise", December 1994;
- Flyer with overview of number of European DSAs and Orgs, June 1995. This will become a

standard feature to be updated when appropriate;

- Production of contact details list National Directory Managers to be updated regularly;
- Annual Report 1994-1995, a low budget version for 'internal' use initially - possibly a glossy version for wider distribution at a later stage.

### *Presence at events*

DANTE has been present at four exhibitions/conferences where NameFLOW-Paradise information material has been distributed:

- EEMA Annual Conference, Amsterdam, 6-9 June 1995.
- Electronic Directories for European Administrations - an initiative of the European Commission - DG III, Brussels, 3-4 May 1995. (see also Introduction)
- CeBIT 1995 - Hannover, 8-15 March 1995
- OnLine/CD ROM 1994 - London, 6-8 December 1994

\* These papers have been included in DANTE's publications series, [\*\*\*DANTE IN PRINT\*\*\*](#), a track record of all papers and articles published by and on behalf of DANTE .

[\[CONTENTS\]](#)

---

## **Annexes**

### **Annex 1:**

[Public directory interfaces](#)

### **Annex 2:**

[National contacts](#)

[Liaison contacts](#)

### **Annex 3:**

[Glossary](#)

---

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

Contact: [Peter.Gietz@dante.org.uk](mailto:Peter.Gietz@dante.org.uk)