NameFLOW-Paradise

Quarterly Service Report October-December 1996

TABLE OF CONTENTS

Introduction

This fourth Quarterly Report reflects the NameFLOW activities and operations for the months October, November and December 1996. 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 Quarterly Report are:

- **"Integrated Internet Directory Service?"**, a positioning paper written by Thomas Lenggenhager (SWITCH, CH) involving the recent LDAP developments and future implications for an integrated Internet directory. (Appendix 9)

- **The X.500 Enabler for LDAP**, the missing link between LDAP and X.500? Short overview from the web server of Criticle Angle including the "special offer" for Academics/NameFLOW customers! (Appendices 10 and 11)

- **MIDAS:** a new and hot NEXOR product to develop your "custom made" web interfaces to X.500 and LDAP. (Appendix 11)

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

Introduction

This chapter summarises activities and operations of the NameFLOW- Paradise Directory service, and associated information services, for the three months October, November, and December1996.

1. Operations/Helpdesk

A new country, C=MY/Malaysia, was added to the DIT. This was delayed during the previous two quarters, due to a low-level network problem, which was fixed in October.

Interworking tests were performed between Giant Tortoise and the proposed new FLDSA for NATO, which is an X.500(93) implementation. Although results were positive, NATO will retain its Quipu FLDSA for the time being, while learning more about the operation of the new implementation.

With the cooperation of the Postmaster of the Department of Computer Science at University College London, the mailing list osi-ds@cs.ucl.ac.uk was removed, and forwarding was set up from osi-ds and osi-ds-request to forum and forum-request@nameflow.dante.net.

2. Outages

Significant outages of service totalled approximately 12 hours in November, and 1 hour in December. There were no significant outages in October. Details of outages are available in the monthly reports.

3. Issues

The US manager is looking at X.500(93) under the auspices of the Challenge97 project of the World Electronic Messaging Association. This should eventually lead to the US FLDSAs converting to X.500 (93) implementations, although there is as yet no timescale for this.

In Australia, Leigh Hume of Macquarie University has taken over from Mark Prior as country manager. Discussions are taking place on both the policy and operational directions of X.500 in AU.

The European Space Agency has established an X.500(93) DSA internally, and is interested in reconnecting to NameFLOW-Paradise. Rather than resurrecting the existing entry under l=Europe, because of the uncertain future of that part of the DIT, it is expected that o=ESA will be registered directly under the root, like o=NATO.

A university site in GB, running the slapd (stand-alone LDAP) server from the University of Michigan, has linked in to NameFLOW-Paradise using the recently announced X.500 Enabler DAP/DSP -> LDAP gateway from Critical Angle, Inc. (URL:http://www.critical-angle.com/). This is a very useful

interoperability tool for sites which do not wish to run a full X.500 DSA, but which want to get/keep connected to NameFLOW-Paradise.

A new telephone system is being installed at ULCC, which will require the helpdesk contact number to change in the coming quarter. ULCC's current main number will be retained, but with a recorded message advising callers to ring the new switchboard number instead.

4. Statistics

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

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

Use of the LDAP server in this quarter was much wider geographically than in previous periods, although the vast majority of use is still from the X.500/web gateway. This spread of use (visible only in the detailed stats on the info-server) presumably reflects the increased availability of LDAP clients, and general interest in LDAP.

5. [NP-93: NameFLOW-Paradise piloting X.500(93) Phase two](#)

Liaison

EEMA

Two EEMA documents commissioned by the Directory Committee have recently been finalised, firstly

Guidelines for the deployment of a corporate Directory
(subtitle: Hitchhikers guide to Directories)

http://www.dante.net/~vincent/docs/hitchhiker.doc

and secondly

Top Level Naming for Europe - Phase 1 Analysis

```
http://www.dante.net/~vincent/docs/TOPOL-1.DOC
http://www.dante.net/~vincent/docs/TOPOL-1A.DOC
http://www.dante.net/~vincent/docs/TOPOL-1B.DOC
http://www.dante.net/~vincent/docs/TOPOL-1C.DOC
```

The documents are available from the DANTE server (personal copies of VB) in electronic format, and

can be made available (on request) to NameFLOW customers.

An important topic of the EEMA Directory Committee is the Directory Challenge-97. This challenge is based on the combined effort of the Security and Directory committees. The challenge is to send a secure message, which is authenticated and encrypted from one continent to another without violating any laws. For the storage and retrieval of keys the usage of the Directory, to be specific an X.500 Directory is suggested. For more information on the challenge see: URL:ftp://ftp.ema.org/pub/challenge/ challenge-97/DIRPLN2.DOC or later version of this document.

IETF

The 37th IETF took place in San Jose (US) in December 1996. 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 6, 7, and 8.

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
3rd Quarter 1996    http://www.dante.net/np/report/qr/96Q3.html

October 1996        http://www.dante.net/np/report/mr/mr9610.html
November 1996       http://www.dante.net/np/report/mr/mr9611.html
December 1996       http://www.dante.net/np/report/mr/mr9612.html
```

In September 1996 the compilation of the second NameFLOW-Paradise Annual Report 1995/1996 started. A first edition of the report was issued last October 1996. A second, slightly revised edition is availbable in electronic fomat via: http://www.dante.net/np/reports/ar/ (hard copies can be ordered from this location as well).

APPENDIX 1 Helpdesk summary for Oct/Nov/Dec 1996

Country		Number of queries			
Full Name	ISO Code	October	November	December	Quarter
Switzerland	CH	–	–	1	1
(China)	CN*	2	3	4	9
Germany	DE	–	–	1	1
France	FR	–	–	1	1
United Kingdom	GB	2	7	2	11
Malaysia	MY	2	–	–	2
Netherlands	NL	1	1	–	2
Norway	NO	1	1	–	2
New Zealand	NZ	–	–	1	1
Poland	PL	1	–	–	1
United States	US	–	2	10	12
Total Requests		9	14	20	43

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

APPENDIX 2 World Root DSA and LDAP summary statistics for Oct/Nov/Dec 1996

Summary of calls to DSA Giant Tortoise

From 0:16:11 on 30 September to 0:04:53 on 31 December

No. of binds	October	November	December	Quarter
Local	286	283	268	837
Remote	4956	4121	4868	13945
Total	5242	4404	5136	14782

No. of operations	October	November	December	Quarter

Local	37	11	2	50
Remote	45468	39769	37750	122987

Total	45505	39780	37752	123037

System usage (calls received)	October	November	December	Quarter
Binds by Directory technicians	996	953	959	2908
Reads of DSA entries	18	5	17	40
Other ops on DSA entries	7	0	0	7
Getedb operations (inc slices)	45352	39475	37500	122327
Spot shadows	23	17	22	62

Total	46396	40450	38498	125344

LDAP usage

LDAP usage from Oct 1 1996 to Dec 31 1996

	October	November	December	Quarter
Connections	22050	21271	21952	65273
Total connect time (seconds)	532971	800520	469214	1802705

(1802705 seconds is 500 hrs 45 mins 5 secs)

APPENDIX 3 Public DUA summary statistics for Oct/Nov/Dec 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	October	November	December	Quarter

Internet	1650	948	795	3393
UK academic X.25 (JANET)	2	1	0	3
EuropaNET X.25	2	0	0	2
Public X.25	1	0	0	1
Total	1655	949	795	3399

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

Domain	October	November	December	Quarter
edu	742	337	289	1368
uk	308	198	94	600
unresolved	284	124	104	512
com	59	43	76	178
net	39	35	43	117
dz	13*	20	46	79
org	20	29	22	71
ca	26	22	9*	57
us	19	17	10	46
nl	19	16*	9*	44
Total	1529	841	702	3072

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

APPENDIX 4 WWW/FTP/Gopher summary statistics for Oct/Nov/Dec 1996

WWW server

TOTALS FOR SUMMARY PERIOD Tue Oct 1 1996 TO Tue Dec 31 1996

	October	November	December	Quarter

```
Unique hosts                             585       494       470      1419
Number of HTML requests                 1403      1432      1225      3981
Number of non-HTML requests              200       237       208       635
Number of malformed requests             123        53        40       228
Total number of all requests/errors     1726      1722      1487      4858
Total number of Kbytes requested       28330     28006     38206     96451
Average requests/day                      55        57        48        53
Kbytes/day                               915       934      1233      1027
```

FTP server

TOTALS FOR SUMMARY PERIOD Tue Oct 1 1996 TO Tue Dec 31 1996

	October	November	December	Quarter
Files Transmitted	539	641	485	1665
Kbytes Transmitted	140973	166607	122786	430367
Average Files Daily	22	23	17	20
Average Kbytes Daily	5873	5950	4385	5403

Gopher server

Gopher usage from Tue Oct 1 1996 to Sat Dec 21 1996

	October	November	December	Quarter
Total connections	19	28	25	72
Total files retrieved	21	27	39	87

APPENDIX 5

NP- 93: NameFLOW-Paradise piloting X.500(93) - Framework document

APPENDIX 6

IETF - Access and Searching of Internet Directories WG
Meeting

Draft Meeting Minutes
Wednesday, December 11, 1530-1730
Reported by: Tim Howes

- Agenda review/changes

        The proposed agenda was slashed quite a bit, with some items
        punted to the list, in an effort to make room for LDAPv3,
        which was anticipated to require a lengthy discussion. Items
        dropped were: pgp draft (to the list), domains draft
        (discussed already in IDS), cip and ldap discussion
        (discussed already in FIND). Items cut down in time were:
        whois++, rwhois.

- application/directory MIME type drafts

        - application/directory framework

                Tim reported that a new application/directory
                framework draft had been produced which addressed all
                outstanding comments received. A brief discussion
                revealed several more issues with the draft that
                people raised.

                These issues were:

                - Example is wrong in how it does line breaks.

                ACTION: Tim to fix this in the draft.

                - Using MIME vs. BEGIN: END: sentinals to carry
                  multiple parts.

                ACTION: Discussion to take place on the list.

                - Change the "proto" parameter to "context"

                ACTION: Tim to change this in the draft.

                - Reference to RFC 1123 time/date formats should be
                  change to reference an I-D describing the ISO 8061
                  time/date format. Chris Newman volunteered to write
                  up this draft.

ACTION: Tim to fix references in the draft.

                ACTION: Chris to write up the 8061 draft.

                  - Ned Freed and Kevin Jordan both had comments that
                    they agreed to send to the authors and/or bring
                    up on the list.

                ACTION: Ned and Kevin and others with issues to bring
                        them up on the list, and/or to give feedback
                        directly to the authors.

      - vcard profile

                Frank Dawson reported that the vcard profile draft
                had been revised to address all known problems and
                issues raised at the last meeting. One additional
                issue was raised at this meeting: the use of MIME
                media types for audio and photo types. The group felt
                this would be better than devising a new scheme.

                ACTION: Frank to revise the draft to reference the
                        MIME media type registry.

- WHOIS++ drafts

      New WHOIS++ drafts have been produced which address
      various problems found during implementation of the
      drafts. These include:

      - Multi-language handling

      - Separate INCHARSET and OUTCHARSET parameters.

      - New templates for X.509, PGP, etc.

      ACTION: Tim to ask the ADs to re-issue these
              documents as proposed standard.

- RWhois

      Network Solutions is working on a meta-directory service
      that will map organization and domain names to directory
      services. Version one supports RWhois. The next version
      will support more general access from other protocols and
      the ability to refer to arbitrary directory services via
      URLs.

- LDAP API

        Tim and Mark produced a new draft updating RFC 1823,
        describing the LDAP API. The updates include preliminary
        support for the changes expected in LDAPv3, support for
        threading, better data encapsulation, etc.

        The group discussed the future of this draft, whether
        they wanted it brought within the working group, and
        if so, what track should it be put on (standard,
        informational, experimental). The group consensus was to
        bring the draft into ASID so it would get the careful review
        it deserves. The group decided to try to push the draft
        along the standards track initially, with informational
        as a fall-back.

        An issue was raised about draft ownership and perceived
        credit, should the draft become an informational RFC.
        The concern was that an informational document that was
        essentially the product of a single company rather than
        the working group, not be presented as the work of the ASID
        group. Only if the group has consensus on the draft and
        feels it has had sufficient input to it, should the draft
        be advertised as a product of the ASID working group.

        ACTION: Tim to re-issue the next version of the draft
                to the working group.

- LDAPv3

        The LDAPv3 discussion began with Mark Wahl summarizing
        the outstanding issues with the current drafts. These
        issues and others raised during the first part of the
        meeting were:

        - The relationship between SSL authentication and the
          LDAP Bind operation needs cleaning up.

        - Compliance - What does it mean to be LDAPv3 compliant?
          The current drafts are not clear.

        - Normalized matching - Do we really want to make this
          optional, as stated in the current draft?

        - Paged searching - When can the server discard result
          sets from searches? Some discussion that this no longer

matters, since each paged search request now contains
enough information to reconstruct the original search.

- Bind as DN w/out password - The semantics of this
  operation need clarifying.

- Mapping onto LDAPv2 - Needs clarifying.

- Mapping onto DAP - Needs clarifying. Should this be
  throughout the document, in a separate document, or
  in an appendix?

- X.500 93 subentries on search - This is believed to
  be covered by doing an explicit search for the proper
  object class.

- Relationship of the X.500 93 contexts feature and
  the current multi-language support - This needs to
  be reexamined and clarified to see if 1) there is
  more valuable stuff we can steal from X.500 and 2)
  there are small changes we can make to be more
  compatible with X.500 93 without increasing complexity.

- Additional SASL mechanisms - Should we define some.

- X.500 97 user requirements - [[can someone explain
  the issue here?]]

- Mapping of strong authentication - How does this
  map onto DAP? What does it mean?

- General direction of LDAPv3 - Some people feel it
  is too complex.

- LDAPv2 revisions - should this be progressed or dropped
  in favor of LDAPv3 entirely?

- LDAPv2 coexistence strategy - We need one.

Discussion very shortly centered around two related topics:
The future of the LDAPv2 drafts, and the general feeling
that the current LDAPv3 proposal represents an overly complex
revolutionary rather than evolutionary change to LDAPv2.

Harald emphatically stated that LDAPv2 could not be
progressed past draft standard since it has the following
known fatal deficiencies:

- No referrals

- No internationalization support

- Broken handling of certificates

- Generally insecure password-based authentication

- No extensibility mechanism

There was much discussion about the best approach to take
to fix these deficiencies in LDAPv3. The debate soon
centered around two options, the final form of which are
presented below:

1) Start with the LDAPv2 RFCs and add support for referrals,
   i18n, extensibility, and better authentication. Fix
   the broken certificate support.

2) Start with the LDAPv3 drafts and do a brutal feature
   review and cut with the following criteria: Anything
   that's in must solve one of the problems above. Other
   features to be added later via the extensibility
   mechanism.

A third option that involved bludgeong Harald into letting
the group progress LDAPv2 as is was quickly dismissed,
much to Harald's relief.

There was much debate and an initial straw poll showing the
room pretty evenly divided between the two options. After
much "concensizing", the group actually came to a miraculous
concensus view that approach 2) was the way to go, provided
there was a way to ensure that the feature review and cut
would actually happen.

Tim proposed and the group agreed that a small group of
motivated volunteers should be tasked with going off and
doing the feature review and cut, which would then be
brought back to the group. The group agreed that this
task must be completed by January 31, 1997.

ACTION: Tim to organize the feature review and cut posse.

- Any Other Business

```
        The meeting concluded with consensification, almost on time.
        The next ASID meeting will be in April in Memphis, TN,
        USA.
```

APPENDIX 7

IETF - Integrated Directory Services Working Group

Draft Minutes Meeting San Jose
Wednesday December 11th 1996 / 09:00-11:30
Reported by : Linda Millington

1. Liaison Reports

An overview of the Internet Nomenclator Project is available at :

http://cm.bell-labs.com/cm/cs/what/Nomenclator

The software is complete and is currently being tested by three organisations and demos are available on request. Comments have been received on draft-ietf-ids-snqp-00.txt and a new draft will be circulated.

The Long Bud project has finished. It provided a demonstration of the technology but it is no longer clear that this is a problem that needs to be solved.

It is business as usual for Nameflow and the following documents are available :

      http://www.dante.net/np/report.html Annual Report
      http://www.dante.net/np/meeting.html Minutes of last Nameflow meeting

There has been an extension to the WHOIS++ Project and the final report is due in March. There are currently two independent implementations and experience is showing that incremental indexing is neeeded. Patrik will send a message to the list on the effect of size on indexes.

2. Status of Drafts submitted to AD for progression

The X.500 root Naming Context is being progressed as an Experimental RFC. The X.500 catalog is being progreesed as Informational.

There are some issues with the DNS Aliases Draft and Harald will send it back to the Group with comments.

3. Schema Requirements Draft

After some discusiion on sorting it was decided that we were not trying to solve the sorting problem in this Draft.

There was some concern expressed that unique id could cause problems across protocol contexts and the decision was that unique id should not be included.

There was rough concensus that creator and modifier should be changed from DN to URI.

The bio label will be changed to description.

A revised Draft will be circulated by Friday 13th December and this will be progressed as Proposed standard.

4. BCP on Directory

It was reiterated that the purpose of this paper is to say what is best now and not to predict the future. The following revisions were agreed on and once they are made the Draft will be progressed :

> a. Add in a reference to the Internet White Pages Schema
> b. The naming structure section needs clarifying
> c. Add in a glossary of terms, possibly using http://www.dante.net/np/ds/glossary.html or RFC 1983

5 Directory Naming

a. Naming Plan for an Internet Directory Service

An overview of draft-ietf-ids-dirnaming-01.txt was given to the group. The basic premise of this Draft is that X.521 Annexe B and NADF have proved to be cumbersome in practice. The Draft suggests using DNS names to build a directory tree.

b. An Approach for Using Domains in LDAP Distinguished Names

draft-ietd-asid-ldap-domains-00.txt which is a successor to RFC 1279. DC attributes are currently rooted under o=Internet and in order to be useful the rest of the organisational attributes would need to be in the DC entries.

c. Conclusions and Actions

The two drafts will be combined and distributed by the end of 1996. The schema will need to be extended and discussion about whether naming based on domains ahould be rooted under o=Internet or not will be taken to the mailing list. Hoyt Kesterson will solicit feedback from ISO on their opinions about o=Internet and expanding the registration authority for DC.

6. CCSO Nameserver (Ph) Architecture

A revised Draft is imminent at which point it will be ready for progression.

7. Finding Stuff

There are a number of outstanding issues in this area. Should this work be part of IDS or the DNS work? The issues will be taken to the SRV WG and a possible conclusion will be that a new WG is needed.

8. LDAP Gateway to RWhois

Michael Mealling gave a brief description of changes being made to rwhois inclusing an LDAP gateway. A Draft will be available by the Memphis IETF and this project will become part of the liaison/service reports.

APPENDIX 8

IETF: Find Working Group

Draft Meeting minutes 37th IETF, San Jose,
Date: December 11, 1996

Chairs: Patrik Faltstrom, Roland Hedberg

Patrik opened the meeting with a set of Wormholes which he suggested we avoid: the Common Indexing Protocol (CIP) works between servers. There will be a separate protocol for clients which we will not discuss today. We will discuss what goes on between servers and we might discuss what goes in an index object and what goes in the protocol.

Patrik suggested that we need to divide the current CIP draft into an architecture document and a protocol document. The discussion on access protocol is deferred to the list.

The group worked on CIPv3 open issues. CIP moves MIME messages back and forth. The payload section is still undefined. Open questions include: What is a Data Set Identifier (DSI)? There is a difference between globally unique IDs which aggregate for indexing and those which are used to break referral loops. We have a solution for aggregation but not for referrals.

HERE THERE BE DRAGONS: A virtual dataset could reside on more than 1 server. The Base-URI of a server is in a referral, but it is not clear what this is for a virtual dataset. When 2 disparate datasets are aggregated they become 1 thing. Referrals must bounce through any server which has aggregated them. This is easier to understand when the servers all speak the same protocol(s). There needs to be a method of pruning. Jeff needs to lay out these rules in the document. Since it is easier to do for homogeneous networks, the group decided to address these first and try to make sure none of the rules for

homogeneous networks would bomb heterogeneous networks later. There need to be rules for two processes: referrals and resolution. For aggregation - if you do not aggregate you can have skip level referrals. If you're part of aggregation you must be part of the referral. This increases processing and response time. Trade offs of these 2 must be made clear in the draft. The decision on the trade off will probably be application or domain driven.

Mime: Instead of creating new mime types, the group has elected to use application/cip-request. Jeff would like the body null and to pack the parameters. This entire discussion was deferred to the list. There was a question asked about whether the spec should be more persuasive about using mime to format the payload? The answer was that we have no idea of what mime will cost. There may be appropriate uses, so the draft should have examples of when to use it.

Version: this was retained for backwards compatibility. This seems to be a problem on TCP transport and a Whois++ issue and should be moved somewhere else. There was a suggestion to look at the SNXP Mime Object transport at http://www.fv.com. This issue was also deferred to the list.

Policy for indexing objects: This is what to do when an indexing server gets an index object it doesn't understand. This can happen during an index-push. If a server accepts the object it may re-export the object intact without storing - especially if it didn't understand the payload. If it rejects it, it should return an error code.

Three more issues were deferred to the list: Mesh Management, which really needs help; the abstract definition of matching semantics also needs help and HTTP URLs. Should 2 different URLs be used for aggregation?

CIP and LDAP Draft (draft-ietf-find-cip-ldap-00.txt)
The draft defines the ldap schema for index information kept by the server and defines the ldap index object. The ldap schema allows mesh management and allows ldap to query the index. Futures will include defining the incremental index mechanism and to discuss the implications of aggregation and filtering. (These will be different depending on payload.) Suggestions are welcome to the list.

CIP and WWW
Indexing the web makes for interesting scaling problems. Do we think we're willing to write index objects for each application or should we use a core SOIF? The Query Referral group from MIT found the context of the index objects more interesting than the rules for passing them around. Negotiation seems to be the key issue - This is my index object, can you handle it? These rules should be as strict as the MIME registry rules. This may cut out proprietary indexing schemes because the public mesh can't deal with it. We should also distinguish between indexing "the web" and indexing "html documents on the web."

Integrating Hierarchical data:
This is an attempt to fit hierarchical data into the index. Hierarchical attributes do not compress well. There are lots of unique tokens. This tweak allowed that compression by having schema information which told which attributes were hierarchical, and also had integrity checks.

RWhois v. 2.0 is out soon. It will use CIP to index information from other NICs.

Patrik wants anyone working on CIP to communicate with him so he can manage progress.

Our Charter needed revision since we had clarified that indexing objects were different from the transport protocol. Patrik updates the charter with new milestone:

> Dec 95 fir CIP as an RFC
> Dec 95 Whois++ navigation as an RFC
> Feb 97 three papers (client interface, server interface and engine) as I-D
> Mar 97 CIP and Whois++ I-D
> Mar 97 CIP and LDAP I-D
> Jun 97 three papers (client interface, server interface and engine) as RFCs
> Jul 97 CIP and Whois++ as RFC
> Jul 97 CIP and LDAP as RFC
> Aug 97 Interoperability between LDAP/X.500 and Whois++ as RFC

Patrik ended the discussion by reminding us we needed clear distinctions between routing and fulfillment.

APPENDIX 9

Integrated Internet Directory Service?

Thomas Lenggenhager, SWITCH
lenggenhager@switch.ch

Introduction

The global NameFLOW-Paradise directory on the Internet is based on X.500(1988) with mostly QUIPU based servers, i.e. Directory System Agents (DSA). SWITCH operates the X.500 node for Switzerland within this directory since 1991, and several universities maintain their own servers or provide the information to SWITCH to load it into a central server.

Users access this global directory today mostly via LDAPv2 [RFC 1777] using clients based on the implementation of the University of Michigan. The WWW to LDAP gateways are the most popular access method e.g. URL: http://nic.switch.ch: 8888/.

LDAP - Lightweight Directory Access Protocol

The lightweight directory access protocol seems to become the access protocol for Internet directories. A new version with extended features is currently being discussed and should be available in software applications in 1997.

History of LDAP

LDAP was designed as a 'cheaper' and easier to implement access method to X.500 directories, mainly for browsers and simple management tasks. Compared to the full featured standardised X.500 Directory Access Protocol (DAP), LDAP has less weight, since it sits directly on top of TCP instead of using full OSI layers 4 to 6. In addition, LDAP mostly uses string encoding versus ASN.1 encoding in DAP. A connection-less LDAP version [RFC 1798] was added mainly for very short directory lookup requests. The well known WHOIS and ph servers are standalone servers who have no knowledge about other servers. Therefore, the user has first to find the right server by other means. The hierarchical X.500 model is based on distributed servers. The knowledge of the contents of the servers is part of the model. However, X.500 servers are more difficult to manage than standalone servers and they have to follow the hierarchical naming structure from the standard. Several people started to think about new solutions. One proposal is known as WHOIS++, building on and extending the old WHOIS approach. These activities influenced other people to design solutions, which are easier to manage, still based on the X.500 standards concept. The result is the extension of LDAPv2.

New LDAP Version 3 under Development

The IETF ASID working group currently discusses the third draft document of LDAPv3, which integrates the functionality of LDAPv2 and CLDAP. Some of the major changes and extensions are:

- Referrals
  Up to version 2, a LDAP server could not refer a client to another server to continue a search, but it had to completely resolve the query on behalf of the client. With the new LDAP version it will be possible to have standalone directories which must not be able to query other directories. Asked for non-local information they will not act themselves but tell the client, based on the knowledge configured, where further information could be found. It is then up to the client to complete the request. That way, a LDAP accessible server is e.g. not required to support DAP or DSP access to the X.500 directory, it can just refer to other LDAP servers. These referrals use the newly defined LDAP URL encoding [RFC 1959].

- X.500(1993) and X.500(1997) Additions
  The new version will include changes/extensions from the X.500 standards revisions from 1993 as well as the upcoming 1997 version. Examples are the distinction between user attributes and operational attributes, the new subschema subentries or collective attributes. It will also be possible to access long result lists page by page.

- Extensible Design
  The new LDAP version can be extended by bilaterally defined extensions. One such example currently under discussion is supported for dynamic directories. The X.500 model assumes that the entries and its values are not changing very often, so it is no problem to cache or replicate them for better performance. As soon as the directory is used to store often changing information, the model must be extended. A list of the active video broadcasted lectures at a university might be one application of a dynamic directory extension. Others might be a list of active discussion channels on a server or the current phone number of a person, which is modified when this person is on the road, in the office or at home.

LDAP integrated into Clients

In spring 1996 Netscape together with several other companies, Microsoft and Novell announced that they will support LDAP in future products. We will see LDAP support integrated into the WWW browsers as well as into e-mail user interfaces in the next generation of these applications.

LDAP integrated into DSAs

The newest release of the ISODE Consortium DSA (IC 3.1) supports not only DAP and DSP but also LDAPv2. This makes the separate LDAP gateways obsolete and helps to improve the overall performance of directory queries via LDAP. Comparing figures 1 and 2 shows the reduced number of components involved for an LDAP directory access.
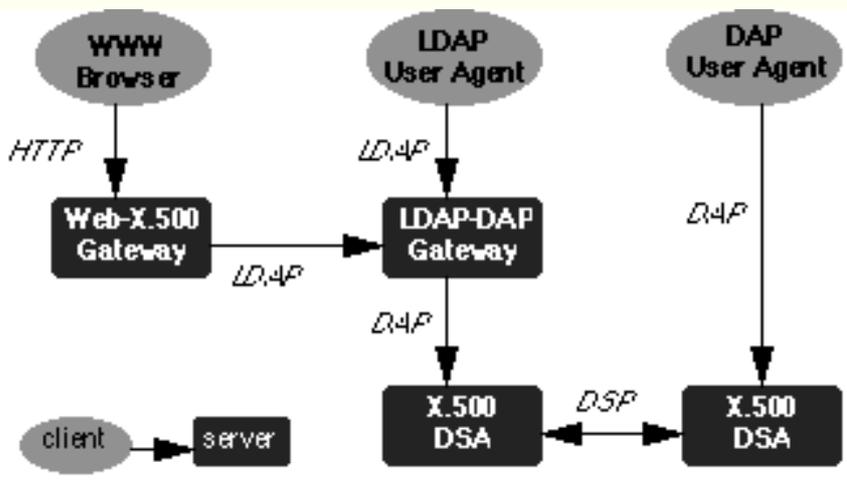
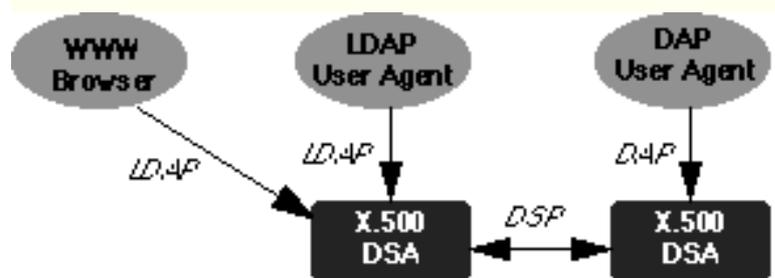
Figure 1. Current LDAP based Access to X.500

Figure 2. Future LDAP-X.500 Interworking

Standalone Directory Servers

In spring 1995, the University of Michigan added a standalone LDAP server, named slapd, as an experimental part to its LDAP distribution. Many organisations started to experiment with it since it was much easier to set up than a full featured X.500 server. We can expect to see more standalone directory servers supporting LDAP as their access protocol in late 1996 or early 1997, supplied by different vendors. One of the reasons of the increased speed of these developments is, that the most important people who designed and implemented LDAP at the University of Michigan work for

Netscape since this summer.

A recently announced product enables the X.500 directory users to access LDAP based standalone directories by acting as a gateway. It can be assumed that in the future DSAs will be extended with this functionality.
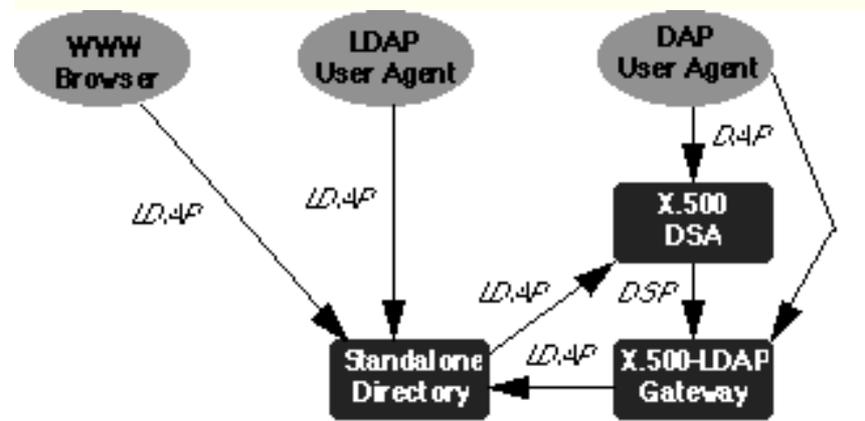

Figure 3. Internet Directory with X.500 and Standalone Servers

Index Servers

In analogy to the well known WWW search engines, the referral support in LDAPv3 will allow to build special directory index servers. WWW search engines index the loosely structured contents of web pages. Directory index servers index the highly structured directory entries. Therefore, they should be capable of answering user queries much more specifically than the WWW search engines.

Index servers will be very helpful to find information which is usually hard to discover in a strictly hierarchical structure. By collecting the information from the hierarchically structured directory and indexing it, e.g. on the name of an organisation or person, an index server can find these entries independent of the country and state or canton level they are listed in the directory.

In the future, index servers will offer business opportunities as WWW search engines are already. Once micro payment services are widely available in the Internet, their role will even grow. It will be a business to add value to the base information by indexing it for potential customers use.

Extended Usage of Directories

Up to now, the X.500 directory in the Internet was mainly used to store white pages information, i.e. address information like in the telephone book. In the future it will be very important to retrieve public keys, certificates and certificate revocation lists from the directory. Many future applications rely on public keys, and today there is no infrastructure available to distribute them Ð a good chance for directories in the future. With the integration of LDAP into available user agents and index servers, directories will hopefully be much easier to use for the end user than they are today in the Internet. Further development needs to be done for better performance and redundancy as well as for the best search strategies based on the users main queries.

The next year will show which of the promised products really make it to the market and whether they can finally give the push needed to get the directories widely deployed. The future will also tell whether an Internet directory will mainly consist of LDAP based standalone servers, as we know it with the percentage of SMTP servers compared to X.400 servers, or whether the X.500 based servers can defend a bigger portion of the installed base.

Glossary

ASID	Access, Searching and Indexing of Directories, an IETF Working Group
CLDAP	Connection-less LDAP
DAP	Directory Access Protocol (used between DUA and DSA)
DSA	Directory System Agent, a X.500 server
DSP	Directory System Protocol (used between DSA and DSA)
DUA	Directory User Agent, a X.500 client
IETF	Internet Engineering Task Force
LDAP	Light Weight Directory Access Protocol
X.500	ISO/ITU specification for a hierarchically modelled, worldwide distributed directory

References

CLDAP	A. Young, "Connection-less Lightweight Directory Access Protocol", RFC 1798, June 1995. ftp://sunsite. cnlab-switch.ch/doc/standard/rfc/17xx/1798
IETF	http://www.ietf.org/
LDAPURL	T. Howes, M. Smith, "An LDAP URL Format", RFC 1959, June 1996. ftp:// sunsite. cnlab-switch.ch/doc/standard/rfc/19xx/1959
LDAP-FAQ	http://www.critical-angle.com/ldapworld/ldapfaq.html
LDAPv2	W. Yeong, T. Howes, S. Kille, "Lightweight Directory Access Protocol", RFC 1777, March 1995. ftp://sunsite.cnlab- switch.ch/doc/standard/rfc/17xx/1777
LDAPv3 Draft	ftp://sunsite.cnlab-switch.ch/mirror/internet-drafts/draft-ietf-asid- ldapv3-protocol-03.txt
NameFLOW-Paradise	http://www.dante. org.uk/nameflow.html
SWITCHdirectory	http://www.switch.ch/switch/info/directory. html

---

APPENDIX 10

X.500 Enabler for LDAP Directories

APPENDIX 11

[PARADISE X.500 Enabler](#)

APPENDIX 12

[MIDAS](#)