#35

Requirements for the future NameFLOW Directory Service

Peter Gietz

This paper belongs to a series of NameFLOW Paradise papers discussing directory services' issues.

DANTE IN PRINT is a track record of papers and articles published by, or on behalf of DANTE. HTML and Postscript versions are available from: http://www.dante.net/pubs/dip

For more information about DANTE or *DANTE IN PRINT* please contact:

DANTE
Francis House
112 Hills Road
Cambridge CB2 1PQ
United Kingdom

Tel: +44 1223 302992
Fax: +44 1223 303005
E-mail: dante@dante.org.uk

Requirements for the future NameFLOW Directory Service

Peter Gietz

Introduction

The DANTE NameFLOW Service has developed as *the* international Directory. NameFLOW was based on the Paradise project and has a long experience of almost a decade. With this DANTE and its co-operating customers take a lead in worldwide deployment of directory technology. The current infrastructure is mainly based on the X.500(1988) standard with the enhancements for replication and access control as implemented in the Quipu software and as defined in RFC1276 [1]. By now this main infrastructure has been expanded by X.500(1993) servers as well as by stand-alone LDAP servers connected to the X.500 Directory via X.500 Enabler [2] of Critical Angel (now called Innosoft X.500 Connector). There also have been attempts to connect 1988 standard complying non-Quipu DSAs [3]. The backbone for the knowledge information management and replication is still based on the non standard QUIPU mode. This includes the transfer of the data entries in the so called Entry Data Block (EDB) format via a protocol specially defined for this purpose.

Most of the Quipu software is not capable of handling dates after the year 2000 in the replication process. Furthermore there is now a standardised method of replication defined in X.525 [4] as part of the 1993 version of the standard which is implemented in a whole range of X.500 products. Due to these two facts, the current NameFLOW infrastructure is obsolete.

Peter Gietz is Applications Engineer at DANTE. His e-mail address is peter.gietz@dante.org.uk .

DANTE has carried out interoperability tests of a multi vendor X.500(1993) directory service in February 1996 and in May 1997 [5, 6, 7]. Although access and retrieval worked fairly well, there were severe problems in the field of replication. The vendors of the included software were not able to completely solve these problems before the second test phase. Due to the unsatisfying outcome of the tests there was no straightforward way for the replacement of the current infrastructure at that time.

This document describes the requirements for this next step of the NameFLOW directory service. It shows the issues of building a worldwide distributed directory service in a time of competing technologies. There are various models for a new service. These are not the subject of this document, but will be described in separate documents [8, 9].

The overall task: connecting the world

Currently the main usage of directory technology in the research community is the white pages service, the best current practice of which is defined in RFC2148 [10]. Older applications in the field of email and file transfer (X.400, FTAM) have been rather unsuccessful in competing with the respective technologies without integrated directory enhancements (SMTP and FTP). In any case, new directory usages can be foreseen in the near future.

As an example Public Key Infrastructure, already defined in the 1988 version of the standard [11], begins to gain more and more importance. The Directory is a natural means

for publishing the public keys of the users. PGP keys are the most popular in the research community today. There definitely is a need for defining standard attributes to store the PGP public keys or key references in the directory. This objective has not yet been fulfilled and a respective Internet draft [12] is outdated by now. The main problem not solved by that proposal is the handling of more than one public key per directory entry. OpenPGP, the new IETF WG on PGP has not addressed this issue [13].

Other applications to come will include the storage of dynamic information for the use in, e.g., real-time multimedia communication. In this area work has been done in respect to LDAP technology [14].

The task of DANTE in the range of directories is to coordinate and connect the directory activities of the NRNs. On the one side DANTE has to take into account the directory technologies of the NRNs on the other side it has to function in a normative way to ensure consistency and interoperability. This should be done by keeping track with the standardising efforts made in the Internet community. The special task of connecting the countries involves special problems that will be dealt with in the following sections.

The Root Context

In a hierarchical system of organising data there is one node that differs from all the other nodes, in having no superior node: the root node. At this point the knowledge information has to be maintained to connect the country level nodes, which are managed by the first level DSAs. Being critical to resolution of names and performing of searches, the knowledge information of the first level DSAs needs to be widely available by means of replication as demanded in RFC1275 [15]. In the current NameFLOW service this high level knowledge information is maintained at a root DSA and replicated to the first level DSA. The root DSA is not a feature of the X.500 standard, but was introduced in the framework of the Quipu knowledge model. The main advantage of this structure is the ease of management because of only one central point of administration.

In the 1988 version of the X.500 standard the terms root context as well as root naming context are mentioned, both of which should be taken synonymously [16]. Unfortunately this concept is not mentioned in the 1993 edition of the standard any more. RFC2120 [17] was dedicated to the use of X.500(1993) standard for managing the root context, with the aim to take over the advantages of the Quipu Root DSA concept. Based on the new protocols applicable for replication as defined in the 1993 standard, namely the Directory Information Shadowing Protocol (DISP) and the Directory Operational Binding Protocol (DOP), three technical solutions were offered in the RFC:

1. A fast track solution proposes the manually update of the knowledge information about the single first level DSAs at the root DSA. The first level DSAs shadow this information via DISP.

2. The slower track solution replaces the manual update of the root DSA with a set of “spot” shadowing agreements between the root DSA as consumer and the first level DSAs as supplier. These agreements should then be enhanced by added subordinate references to emulate a Hierarchical Operation Binding (HOB) as normally achieved via DOP.

3. The long term solution includes a genuine HOB via DOP.

It is important to note, that all three solutions presuppose more or less severe amendments to the standard, as reflected in the defect reports annexed to RFC2120. The proposal was based on the fact that DOP was not fully implemented in most of the software products.

A different approach to the root context problem has recently been brought forward in [18]. This contribution proposes the enhancement of the concepts of Administrative Directory Management Domains (ADDMD) and Private Directory Management Domains (PRDMD) as defined in the X.500 standard. This solution overcomes the single point root noded and replaces it with an Asymmetric Replicated Root Directory System. It requires a number of amendments to the standard concerning the concept of the First Level DSAs, the DSP as well as the operational model.

It is essential for the future NameFLOW service to have a manageable and deployable solution for the root context problem. It will be a matter of decision how much information should be maintained by the superior level, taking into account the performance issues as well as the concern of reduction of complexity and the effects of the single point of failure issue.

Support of Quipu Replication

For a smooth transition to a new directory system and for backwards compatibility with participants, still running DSAs in Quipu mode, the support of Quipu replication should be provided. A DSA as part of the X.500 Backbone should therefore fulfill the requirements defined in RFC1275 [15] by means of the methods proposed in RFC1276 [1]. This includes support of the EDB format, as well as the support of the EDB file based replication protocol.

Support of Quipu ACLs

The first edition of the X.500 standard did not include access control mechanisms. The Quipu implementation had to go its own proprietary way again, by creating the Access Control List (ACL), a specific attribute specifying the type of access allowed to single users or groups of users. These ACLs were usually inherited down a subtree into the single entries.

The 1993 standard defined a complete and complex so called Basic Access Control Scheme (BACS) [19], which provides the means to protect entries, single attributes, attribute values, as well as DNs from all possible actions by users, groups of users or users defined by their authentication levels. Entries can be grouped in respect to administrative areas. To make it even more complex BACS defines precedence levels for the sake of resolving conflicts of several Access Control Information Items (ACII). This extreme comprehensiveness of the BACS leads to an explosion of the amount of data stored in the directory. Additionally, it makes the task of defining access control very complex for the administrator.

There is a definite need for the establishment of an access control policy for the NameFLOW service to reduce the complexity and to gain homogenity in this respect.

For a smooth transition to a service based on the X.500(1993) standard, the NameFLOW service requires an easy to use and policy based administrator interface for the definition of access control.

Integration of LDAP

LDAP originated as a simplified access protocol to X.500. In its development in the Internet standardisation process it evolved to an X.500 competing directory technology. Although there is a complete set of RFCs describing the latest version 3 of this technology [20, 21, 22, 23, 24], there are still a number of unsolved problems, mainly in the fields of access control, authentication, replication and knowledge information management. In spite of this, stand-alone LDAP servers become more and more popular in the research community, and it is one of the main issues for the future NameFLOW serv-

ice to integrate this technology. There is a yet unfinished Internet draft which includes a solution for this integration [25]. It proposes algorithms for interconnection of LDAP servers with each other as well as with the global X.500 directory, based on LDAP Referral servers. The usage of the following technologies is presupposed by this Internet draft: Domain Name System (DNS), Server Location Protocol (SLP) [26, 27] and Common Index Protocol (CIP) [28, 29, 30, 31].

Due to the increasing popularity of LDAP the future NameFLOW architecture has to integrate this technology. Possible scenarios would be either an LDAP backbone solution [8] or an integration of LDAP servers by an X.500 backbone [9].

Recently work on LDAPv3 profiles and conformance testing has been done outside the IETF, which can be helpful in including LDAPv3 products into the NameFLOW service: Work on the definition of special task LDAP servers [32, 33, 34, 35, 36], on the means of evaluation of LDAP software [37], as well as on interoperabilty testing [38].

Schema and name mapping

The original X.500 Directory schema was defined in X.520 [39] and X.521 [40] providing a standard set of attributes and object classes. This set gives the means for interoperability between independently managed servers, in respect to information retrieval as well as to the naming schema, including the rules how to build a DN. Today more attributes are required to solve the current directory needs, e.g., for the mobile telephone number. In the scope of white pages services new proposals for standard attributes have been made outside the X.500 standardising process. RFC2218 [41] as outcome of the IETF working group IDS defines an Internet White Pages Schema (IWPS). A competing proposal was made by the Network Applications Consortium, the Lightweight Internet Person Schema (LIPS) together with a mapping proposal for LIPS and IWPS [42]. Additionally there is a proposal for the inetOrgPerson object class [43] to meet the requirements of today's Internet and Intranet directory services.

The conventional X.500 approach to naming involved not yet completely solved problems in respect to name registration. Another naming and registration structure, which is working and widely deployed in the Internet community is the DNS. A new proposal has been made [44, 45] to use this existing structure for directory service, introducing a new attribute named Domain Component (DC). A solution for the integration of standard X.500 naming schema into this new naming schema is included in [25].

To acquire the best possible flexibility in respect to future developments the forthcoming NameFLOW service should be able to handle the above mentioned extensions to the standard attributes and object classes. Another task for the future NameFLOW directory service will be to include the DC naming plan and to provide a mapping to the conventional X.500 naming schema.

Integration of Indices

As Directory use continues to grow and as new demands for future directory services (e.g., public key server) are coming up, there is a need for enhanced mechanisms of information retrieval beyond the currently available search facilities and replication mechanisms. Such demands were already addressed in the year 1993 [46] and the first solution was brought forward deploying classical X.500(1993) means of replication for indexing purposes [47]. An implemantatin of this service can be found at [48]. A more promising approach to an index mechanism for directories [49] deploys the centroid technology of whois++ [50].

The above mentioned CIP serves as a complete and complex system for indexing directories. The Tagged Index Object defined in CIP [31] includes all the information needed to be indexed. In addition several transport mechanisms are defined, as well as a retrieval interface.

An important aim of the future NameFLOW directory service is to include mechanisms for indexing and for distribution of index information. Although
CIP provides a good solution, a first step in reaching this aim is to only deploy a subset of it.

Summary

The following requirements are essential for the future NameFLOW directory service:

- Building up a new infrastructure that fulfills the requirements of the NameFLOW customers but also takes into account the standardising work going on in the internet community.
- A deployable solution for the root naming context problem most probably compliant to RFC2021 [17].
- Integration of the current Quipu software, including the replication methods described in [1].
- Tools for policy based management of X.500(1993) BAC.
- Integration of LDAPv3 server technology as described in [25].
- Integration of the DC naming structure as proposed in [44, 45].
- Construction of an index service for directory information on tracks of [49].

The special task of the future NameFLOW service will be to include backwards compability on the one hand by supporting standard X.500(1988) as well as non-standard Quipu DSAs. On the other hand the future technologies have to be supported as well, namely X.500(1993) with a solution for the root context problem, as well as LDAPv3 stand-alone servers. In addition policies for replication and access control have to be defined.

References

[1] S.E. Hardcastle-Kille, “Replication and Distributed Operations extensions to provide an Internet Directory using X.500”, RFC1276, November 1991.

[2] M. B. Caines, “Experiences using two LDAP server implementations”, <url:http://www.dante.net/np/ldap/wolver.html>, March 1997.

[3] B. Koechlin, K. Treca, and P.-A. Pays, “Strangers in Paradise”, INET94 Proceedings, June 1994.

[4] ISO/IEC 9594-8 | ITU-T Rec X.509 (1988) “The Directory: Authentication framework”, 1989

[5] V. Berkhout, “NP-93: NameFLOW-Paradise piloting X.500(93). Pilot framework document”, DANTE Docs VB(97)004, <url:http://www.dante.net/np/93pilot/framework.html>, January 1997.

[6] V. Berkhout, “NameFLOW-Paradise X.500(93) Test Results - Phase One”, DANTE Docs VB(96)15, <url:http://www.dante.net/np/93pilot/phase1-results.html>, March 1996.

[7] J. Horton, "NameFLOW-Paradise X.500(93) Test Report - Phase two”, DANTE Docs JH(97)001, <url:http://www.dante.net/np/93pilot/phase2-results.html>, August 1997.

[8] J. Horton, “Draft Plan for a NameFLOW LDAP Pilot”, <url:http://www.dante.net/np/LDAP-Pilot-Plan-01.txt>, February 1998.

[9] P. Gietz, “A NameFLOW plan for a Hybrid X.500-LDAP directory service”, work in progress.

[10] H. Alvestrand, P. Jurg, “Deployment

of the Internet White Pages Service", RFC2148, BCP 15, September 1997.
[11] ISO / IEC 9594-9 | ITU-T Rec X.525 (1993) "The Directory: Replication", 1994.
[12] R. Hedberg, "Definition of X.500 Attribute Types and a Object Class to Hold public PGP keys", Expired Internet draft, <draft-ietf-asid-pgp-02.txt>, February 1996.
[13] "An Open Specification for Pretty Good Privacy (openpgp)", IETF working Group Charter, <url:http: / / www.ietf.org / html.charters / openpgp-charter.html>, March 1998.
[14] Y. Yaacovi, M. Wahl, and T. Genovese, "Lightweight Directory Access Protocol (v3): Extensions for Dynamic Directory Services", Internet draft, <draft-ietf-asid-ldapv3-dynamic-07.txt>, December 1997.
[15] S.E. Hardcastle-Kille, "Replication Requirements to provide an Internet Directory using X.500", RFC 1275, November 1991.
[16] D. Chadwick, "Phasing out / Opening Up the Root DSA", DANTE in Print 14a, <url:http: / / www.dante.net / pubs / dip / 14 / root.html>, February 1995.
[17] D. Chadwick, "Managing the X.500 Root Naming Context", RFC2120, March 1997.
[18] A. Lloyd, "Global Directory Service Backbones and Multi Master Operation - a contribution", ISSS / DIR-WS Documents N. 22, April 1998.
[19] ISO / IEC 9594-2 | ITU-T Rec X.525 (1993) "The Directory: Models", 1994.
[20] M. Wahl, T. Howes, and S. Kille, "Lightweight Directory Access Protocol (v3)", RFC2251, December 1997.
[21] M. Wahl, A. Coulbeck, T. Howes, and S. Kille, "Lightweight Directory Access Protocol (v3): Attribute Syntax Definitions", RFC2252, December 1997.
[22] M. Wahl, S. Kille, and T. Howes, "Lightweight Directory Access Protocol (v3): UTF-8 String Representation of Distinguished Names", RFC2253, December 1997.
[23] T. Howes, "The String Representation of LDAP Search Filters", RFC2254, December 1997.
[24] T. Howes, M. Smith, "The LDAP URL Format", RFC2255, December 1997.
[25] A. Brown, C. Weider, and M. Wahl, "Practical guidance for naming and interconnectivity", Internet draft, <draft-ietf-lsd-nandi-00.txt>, April 1997.
[26] R. Moats, M. Hamilton, and P. J. Leach, "Finding Stuff (How to discover services)", Internet draft, <draft-ietf-svrloc-discovery-06.txt>, April 1998.
[27] J. Veizades, E. Guttman, C. Perkins, and S. Kaplan, "Service Location Protocol", RFC2165, June 1997.
[28] J. Allen, M. Mealling, "The Architecture of the Common Indexing Protocol (CIP)", Internet draft, <draft-ietf-find-cip-arch-01.txt>, November 1997.
[29] J. Allen, M. Mealling, "MIME Object Definitions for the Common Indexing Protocol (CIP)", Internet draft, <draft-ietf-find-cip-mime-02.txt>, January 1998.
[30] J. Allen, P. J. Leach, "CIP Transport Protocols", Internet draft, <draft-ietf-find-cip-trans-00.txt>, June 1997.
[31] R. Hedberg, B. Greenblatt, R. Moats, and M. Wahl, "A Tagged Index Object for the use in the Common Indexing Protocol", Internet draft, <draft-ietf-find-cip-tagged-06.txt>, March 1998.
[32] "Read-Only LDAP Server Profile Definition", draft for LDAP Profile Specification Working Group, <url:http: / / www.critical-angle.com / test / profile / pr_ro2.pdf>, January 1998.
[33] "Read-Write LDAP Server Profile Definition", draft for LDAP Profile Specification Working Group, <url:http: / / www.critical-angle.com / test / profile / pr_rw2.pdf>, January 1998.
[34] "White Pages Application LDAP Server Profile Definition", draft for LDAP Profile Specification Working Group, <url:http: / / www.critical-angle.com / test / profile / pr_wp2.pdf>, January 1998.
[35] "Certificate Application LDAP Server Profile Definition", draft for LDAP Profile Specification Working Group, <url:http: / / www.critical-angle.com / test / profile / pr_cert2.pdf, January 1998.
[36] "Single Sign On Application LDAP Server Profile Definition", draft for LDAP

Profile Specification Working Group, <url:http://www.critical-angle.com/test/profile/pr_sso2.pdf>, January 1998.
[37] A. Sundermann, P. Fantou, “LDAP V3: Level of Support of an LDAP Server”, Draft for CEN/ISSS Directory Workshop, May 1998.
[38] Chris Apple [Ed.], “Basic LDAPv3 Interoperability Test Suite (BLITS)”, Issue 1.0 Draft 3, <url:http://www.imc.org/imc-ldap-test-suite/blitsi1d3.html>, November 1997.
[39] ISO/IEC 9594-6 | ITU-T Rec X.520 (1993) “The Directory: Selected Attribute Types”, 1994.
[40]ISO/IEC 9594-7 | ITU-T Rec X.520 (1993) “The Directory: Selected Object Classes”, 1994.
[41] T. Genovese, B. Jennings, “A Common Schema for the Internet White Pages Service”, RFC2218, Oktober 1997.
[42] Network Applications Consortium, “Lightweight Internet Person Schema”, <url:http://www.netapps.org>, May 1997.
[43] M. Smith, “Definition of the inetOrgPerson LDAP Object Class", Internet draft, <draft-smith-ldap-inetorgperson-00.txt>, March 1998.
[44] S. Kille, M. Wahl, A. Grimstad, R. Huber, and S. Sataluri, “Using Domains in LDAP/X.500 Distinguished Names", RFC2247, January 1998.
[45] A. Grimstad, R. Huber, S. Sataluri, and M. Wahl, “Naming Plan for Internet Directory-Enabled Applications”, Internet draft, draft-ietf-ids-dirnaming-04.txt>, March 1998.
[46] J. Postel, C. Andersen, “White Pages Meeting Report”, RFC1588, February 1994.
[47] P. Barker, “X.500 Index DSAs”, DANTE in Print 13, <url:http://www.dante.net/pubs/dip/13/idsa.html>, June 1995.
[48]<URL: http://homes.ukoln.ac.uk/~lisap/ccsap/Directory/demos.html>
[49] D. Chadwick, “IndeX.500”, DANTE in Print 19, <url:http://www.dante.net/pubs/dip/19/19.html>, May 1996.
[50] C. Weider, J. Fullton, and S. Spero, “Architecture of the Whois++ Index Service”, RFC1913, February 1996.