

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

Pilot framework document

Reference: VB(97)004

Version: 4

Date: 30 January 1997

Author: Vincent Berkhout

Introduction

The NameFLOW-Paradise Directory Service is mainly based on Quipu implementations, a specific flavour of X.500 with several enhancements to the ISO standard and therefore not fully compliant. The initial edition of the X.500 standard, published in 1988, had a few shortcomings, such as the lack of a replication mechanism. The shortcomings were acknowledged and an improved edition of the X.500 standard was published in 1993. This pilot framework document describes some of the advantage of implementations conforming to the 1993 standard and explains why NameFLOW-Paradise wants to migrate to 1993 from the current Quipu model. An overview will be given of the planned phases, starting with a root context test, an interworking/scaling test and finally a transition of the complete infrastructure. The detailed test plans for each phase are described in separate documents and include a description of one-week tests, used methods and participating organisations. If possible the test structure should remain in place after the second test/phase allowing organisations to continue testing or even performing a full transition for their part of the DIT. This document is not intended as an introduction to X.500(93) and it is assumed that the reader has some basic knowledge of X.500 Directories.

The Objectives

The NP-93 migration has two objectives:

1. Introduce a Directory system based on X.500(93)
2. Phasing out the Quipu system.

While deploying the X.500(93) system the Quipu system will be gradually phased out. Once an X.500(93) service is extensively tested and operational, Directory services based on Quipu systems can be discontinued. The long term goal is to have an X.500(93) system without the need for "backwards compatibility" (or gatewaying from X.500(93) to Quipu or X.500(88)). It is envisaged that the X.500(93) Directory infrastructure can provide a "Directory Backbone" for X.500 implementations including new Directory Servers using the LDAP protocol.

Why a 1993 pilot?

The pilot is needed to gain experience, find flaws in the software or detect defects in the 1993 edition

of the X.500 standard. The pilot will pave the way for an operational Directory infrastructure primarily based on X.500(93). NameFLOW sees the opportunity to perform the first field test as there was no and currently still limited practical experience deploying X.500(93) on a large scale. The goal of NameFLOW is to use X.500(93) to provide a multi system/multi vendor Directory service allowing other Directories (including non-R&D Directories) to connect and share information.

The X.500(93) edition of the standard was improved in the following areas:

- access control
- replication (sub tree and incremental)
- schema knowledge

The upgraded X.500(93) system will be improved with:

- new management tools
- better interfaces to the Directory (e.g. using LDAP)
- more reliable with complete data
- shorter response times
- integration of Directory Servers based on LDAP

Why not a 1988 test?

There are several reasons for testing the most recent version X.500(93) and not X.500(88). The most important reason is that X.500(88) does not deliver the required operational functionality, such as replication. Using X.500(93) provides at least functionality similar to the Quipu model and has additional enhancements e.g. in the security area. Another reason is that software development is currently focused on X.500(93), and very little attention is paid to further development X.500(88) or Quipu. Only a limited number of X.500(88) implementations exist and it is expected that this number will not increase. Most X.500(88) implementations can now be upgraded to X.500(93) without major problems. The final reason not to test X.500(88) is that it is operational and has already been successfully tested. (See [1] for further information.)

The Transition Scenarios

There are two possible scenarios for the transition:

One-step transition:

Quipu -> X.500(93) (direct)

Two-step transition:

Phase one: Quipu -> X.500(88)

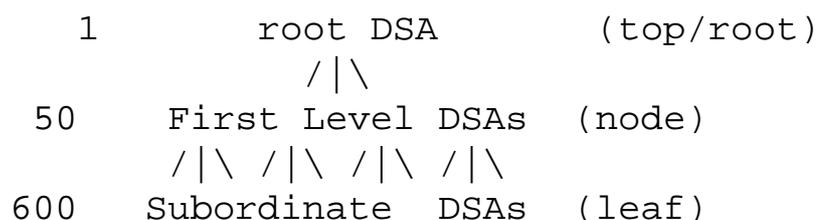
Phase two: X.500(88) -> X.500(93)

A one-step transition (Quipu -> 93) is preferred over a two-step- transition (Quipu -> 88 -> 93), even though it will be more challenging as not all implications can yet be foreseen. The one-step transition is expected to require less effort and time. Currently there are only a few pure X.500(88) implementations connected to NameFLOW-Paradise and introducing this third X.500(88) flavour on a large scale would be undesirable. A major benefit of the one-step transition is that once the X.500(93) model is supported, it should allow both Quipu and X.500(88) implementations to connect (Quipu -> 93 <- 88). The reason for this is that the 1993 version of standard supports at least all 1988 features.

The major change from the current Quipu model and the an X.500(93) model is that the root DSA can be minimised as a new First Level DSA can have its own root entry. The minimum requirements for such a root DSA is knowledge reference coordination (how to reach another FLDSAs?) and a form of information replication to other FLDSAs (see [2] and [5]).

Bottom-up vs. Top-down transition

The current functional model (depicted below) is based on one root DSA at the top level, at the next level are called First Level DSAs (50) and finally other DSAs (600) at the lower levels.



The bottom-up approach requires that all subordinate DSAs must be replaced before a superior DSA can be replaced. Bottom-up will mean to replace the greatest number (600) of "leaf" DSAs first. In practice this will be unfeasible as a number of leaf DSAs are poorly maintained. In addition it would need a great effort to co-ordinate and some organisations might be reluctant to migrate. To minimise the impact on the current operational service a separate NP- 93 tree ("DIT") will be used, whereby the test tree is a partial shadow copy of the real tree. This will avoid some of the migration problems but requires synchronisation (conversion tools) between the two trees. So why the top-down approach? As explained above the root DSA functionality could be reduced, in theory the root DSA could even be considered obsolete. In practice this means that every FLDSA must have its own root entry. All that is necessary for a FLDSA are knowledge references to the other FLDSAs. In short: the root DSA as known in the Quipu model could be replaced with a simple file containing the references making the Root DSA obsolete. This step could be repeated for each node (level) of the tree. The implication is that each branch of the tree can be replaced at a time convenient to the National Service provider managing their branch of the DIT, and repeat this for the subordinate level. Using this strategy does not require a "switch day" where one complete level or even the complete tree will switch to a X.500(93) based infrastructure, but will allow a node to run two DSA in parallel until the subordinate level fully supports X.500(93). A simple reference file would suffice initially, however some root DSA functionality should be maintained as the first experiences with X.500(93) learned that some functionality of the current Quipu model/root DSA must be maintained to make a global system manageable. The two "root DSA functions" in the 93 model are:

1. distribute all knowledge references to First Level DSAs,
2. some form of replication. (See [2] and [5] for full details).

Time table

So far the time frame has strongly depended on the availability of X.500(93) software and the speed that required updates have been implemented. Although software vendors state that their software is fully compliant to the 1993 edition of standard, it still has to be seen how much is implemented. Being compliant with the standard does not mean that implementations can interwork. To get all the organisations to work together it was proposed to have dedicated weeks or several consecutive days to perform the tests. The Experience Test (Phase 0: done in August/September 1995) was followed by:

1. Root Context Test (Pure X.500(93) replication) (February 1996)

The next two phases are:

2. Inter working & Scaling Test (using X.500(88) and Quipu)
3. Operational Transition

Phase Zero: Experience Test (August/September 1995)

The Initial Experience Test, referred to as phase 0 was a small sized test done by a limited number (three) of countries and several severe problems were encountered: unavailable '93 software, Root Context replication (DISP agreements, replication errors) and unsupported EDB conversion tools.

Phase One: Root Context Test (February 1996)

The Root Context Test was performed in February 1996 by a group of ten participants. The actual test were defined at the EuroSInet test writing workshop and the target was to test the top level/Root Context, in particular DISP. The test was extended to DAP and DSP. A detailed description of Phase One are available via the WWW:

Plan: <http://www.dante.net/np/93pilot/phase1-plan.html>

Results: <http://www.dante.net/np/93pilot/phase1-results.html>

Phase Two: Inter Working & Scaling Test (or The Small Transition)

On the precondition that the Root Context works properly the next phase will be to test interworking (including Quipu and X.500(88) coexistence) and scaling. The test is planned for April/June 1997 and will allow involvement from organisations with an interest in X.500(93) and the new infrastructure. This will be a large scale test, focusing on ACI, coexistence between X.500(93) and Quipu/X.500(88). In X.500(93) ACI is differently modelled (outside an entry) and this could have an impact on performance. This large scale test is intended for big organisations with large data sets.

During this phase the first non-Quipu (derivative) DSAs can be introduced. A major part of access to the Directory is via LDAP gateways and clients and should therefore be tested. If there is an interest the coexistence with an LDAP servers can be tested.

Phase Three: Operational Transition (or The Big Transition)

All will depend on the outcome of the previous interworking & scaling test. If previous tests are successful and the participating organisations are convinced that a transition can be done, then the Big Migration to an X.500(93) infrastructure can start. This will be done after consultation of the NameFLOW customers.

The Tests

The EuroSInet test suites (scenarios) have an emphasis on X.500 inter working and are well suited as a basis for the NP-93 tests. For phase one NameFLOW-Paradise sent four representatives to the EuroSInet workshop where the X.500(88) test was enhanced to an X.500(93) test. The test suite started with basic DAP and DSP and now include the first DISP and ACI test. For the second phase a revised document will be produced with a condensed set of tests. In addition, the test document will be extended with "Root Context" specific DISP tests (see [4]) and LDAP tests.

Documents

The outcome of the test needs to be documented.

- Test Experiences (This plan/Tests/Problems/Evaluation report)
- Possibly a first draft of technical manual "How to upgrade a Quipu (or 88) DSA to an X.500 (93) DSA"

Challenges Ahead

Problems may occur in the following areas:

- Conversion of the Quipu EDB format to X.500(93) format and vice versa.
- Effective list and search operations using the 1993 model
- Access Control Information is differently modelled, conversion could be problematic (automatic conversion could create very long ACI)
- Implementing the Root Context, how to manage knowledge references and what level of replication will be needed.
- How to support unsupported Quipus (prior to ISODE version 8.0) and pure X.500(88) DSAs in the NP-93 infrastructure.
- Will current LDAP clients work, in particular those using University of Michigan LDAP servers 3.1 and 3.2.
- Use of the X.500 Enabler to integrate LDAP Directory Servers.

Mailing List

A dedicated NP-93 mailing list is available:

Name list: NP-93@dante.org.uk
Subscription: NP-93-request@dante.org.uk
Archive: not available

The list np-93 has a moderator for subscription.

Acknowledgements

Many thanks to the following organisations and people who participated and made valuable contributions to the tests and documents.

Brunel University (UK), DANTE (UK) Delft University of Technology (NL), EuroSInet (UK), NEXOR (UK), Nicholas Copernicus University (PL), SURFnet (NL), SWITCH (CH), SUNET (SE), UKERNA (UK), ULCC (UK), Isode (UK) and many other organisations.

References:

- [1] OIFP reports <<ftp://ftp.nameflow.dante.net/paradise/oifpfinal.txt>>
- [2] "Managing the Root Context", David Chadwick, "DANTE in Print #18", <<http://www.dante.net/pubs/dip/18/root.txt>>
- [3] "Root Context Test Plan", Vincent Berkhout, DANTE Docs VB(96)006, <<http://www.dante.net/np/93pilot/phase1-plan.html>>
- [4] "Root Context Test Results", Vincent Berkhout, DANTE Docs VB(96)015, =<<http://www.dante.net/np/93pilot/phase1-results.html>>
- [5] "X.500 Shadowing Profiles", David Chadwick, draftPietfPidsPx500P shadprofP00.txt (Expired January 1997, copy available via <http://www.dante.net/np/ds.html>)