

Network Working Group
Request for Comments: 1430

S. Hardcastle-Kille
ISODE-Consortium
E. Huizer
SURFnet bv
V. Cerf
Corporation for National Research Initiatives
R. Hobby
University of California, Davis
S. Kent
Bolt, Beranek and Newman
February 1993

A Strategic Plan for Deploying an Internet X.500 Directory Service

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard. Distribution of this memo is unlimited.

Abstract

There are a number of reasons why a new Internet Directory Service is required. This document describes an overall strategy for deploying a Directory Service on the Internet, based on the OSI X.500 Directory Service. It then describes in more detail the initial steps which need to be taken in order to achieve these goals, and how work already undertaken by Internet Engineering Task Force Working Groups (IETF WGs) is working towards these goals.

Table of Contents

1.	REQUIREMENTS	2
2.	SUMMARY OF SOLUTION	3
3.	INFORMATION FRAMEWORK	3
3.1	The Technical Model	3
3.2	Extending the Technical Model	4
3.3	The Operational Model	5
4.	NAME ASSIGNMENT	5
5.	DIRECTORY INFRASTRUCTURE	6
5.1	Short Term Requirements	7
5.2	Medium Term Requirements	9
5.3	Long Term Requirements	9
6.	DATAMANAGEMENT	9
6.1	Legal Issues	10
7.	TECHNICAL ISSUES	10

7.1	Schema	11
7.2	Use on the Internet	11
7.3	Replication of Knowledge and Data	12
7.4	Presentation of Directory Names	13
7.5	DSA Naming and MD Structure	13
8.	SECURITY	13
8.1	Directory Provision of Authentication	14
8.2	Directory Security	15
9.	RELATION TO DNS	16
10.	EXTERNAL CONNECTIONS	16
11.	REFERENCES	17
12.	Security Considerations	19
13.	Authors' Addresses	20

1. REQUIREMENTS

There is substantial interest in establishing a new Directory Service on the Internet. In the short term, there is pressure to establish two new services:

- White Pages lookup of users;
- Support for X.509 Authentication for a range of applications in particular for Privacy Enhanced mail [Lin89].

In the medium term, there are likely to be many requirements for Directory Services, including:

- General resource lookup, for information ranging from committee structures to bibliographic data;
- Support of management of the Internet infrastructure, and integration of configuration information into the higher level directory;
- Support of applications on the Internet. For example:
 - o Electronic distribution lists;
 - o Capability information on advanced user agents;
 - o Location of files and archive services.
- Support for Mail Handling Systems; Be they RFC-822 based or X.400 based (IETF MHS-DS WG), e.g.,:
 - o Support for routing;
 - o Info on User agent capabilities; essential for a usage of Multimedia mail like MIME (Multipurpose Internet Mail Extensions).

For the longer term, more sophisticated usages of X.500 are possible extending it into a useful and fast yellow pages service.

2. SUMMARY OF SOLUTION

In principle, the current Internet Domain Name System (DNS) could be used for many of these functions, with appropriate extensions. However, it is suggested that a higher level of directory service is needed. It is proposed to establish an Internet Directory Service based on X.500. This provides appropriate functionality for the services envisaged and gives flexibility for future extension. This extension could be achieved either by tracking the evolution of the OSI Standard or by work specific to the Internet. In practice, it is likely to be a mixture of both.

By deploying X.500 in some form on the Internet, a truly global and universal Directory Service can be built that will provide Internet users with fast access to all kinds of data. The X.500 Directory Service in this case may range from a simple white pages service (information on people and services) to coupling various existing databases and information repositories in a universal way.

Currently, several different but cooperating X.500 Directory Services pilots are taking place on the Internet. These pilots form an important base for experimenting with this new service. Starting with these pilots, with the X.500 products arriving on the market today, and given sufficient funding for the central services described in this paper an operational X.500 Directory Service can be deployed.

The final goal of the strategy described in this paper is to deploy a fully operational Directory Service on the Internet, providing the functions mentioned in the previous section.

3. INFORMATION FRAMEWORK

The most critical aspect of the Directory Service is to establish an Internet Information Framework. When establishing a sophisticated distributed directory with a coherent information framework, it involves substantial effort to map data onto this framework. This effort is an operational effort and far outweighs the technical effort of establishing servers and user agents.

3.1 The Technical Model

By choosing the X.500 model as a basis for the information framework, it will also be part of a (future) global information framework. The key aspects of this model are:

- A hierarchical navigational system that couples distributed databases (of various kinds), which allows for management of the data by the organization/person responsible for the data;
- Each object in this information structure (called the Directory Information Tree, DIT) is represented as an entry;
- Objects are typed by an "object class", which permits multiple inheritance;
- An object is described by a set of attributes;
- Each attribute is typed. Attribute types are hierarchical;
- Each attribute type has an associated attribute syntax, which may be generic or shared with other attributes (e.g., Integer Syntax; Distinguished name Syntax); This allows for representation of simple attributes (e.g., strings or bitmaps) or complex ones with detailed structures.
- Each entry has an unambiguous and unique global name;
- Alternate hierarchies may be built by use of aliases or pointers of distinguished name syntax.

This framework allows for representation of basic objects such as users within organizations. It is also highly extensible, and so can be used for a range of other applications.

3.2 Extending the Technical Model

In the longer term, the model could be extended to deal with a number of other requirements which potentially must be met by an Internet Directory Service. Possible extensions include:

- Support of ordered attributes (needed by some applications such as message storage);
- Extensions to allow unification with management information, associated with SNMP (Simple Network Management Protocol) [CFSD90] or other management protocols;
- Handling of non-hierarchical data in a better manner for searching and retrieval, whilst retaining the basic hierarchy for management purposes. This is essentially building a general purpose resource location service on top of the basic infrastructure. It will need work on the information model, and not just the access protocols.

It is noted that although X.500 may not provide the ultimate solution to information retrieval, it has good potential for solving a lot of information service related problems.

3.3 The Operational Model

To make the Directory Service with a coherent information framework really operational requires a lot of effort. The most probable operational model is one where larger organizations on the Internet maintain their part of the DIT on their own DSA (Directory System Agent). Smaller organizations will "rent" DSA space from regional networks or other service providers. Together these DSAs will form the Internet Directory Service Infrastructure. To couple the various parts of the DIT that are contained on these Internet DSAs, a special DSA containing the Root for the naming hierarchy within the DIT has to be established and maintained.

The following tasks can be foreseen:

- Defining the naming hierarchy; See section 4.
- Creating the Directory Infrastructure; See section 5.
- Getting the Data into the directory; and
- Managing the data in the Directory. See section 6.

4. NAME ASSIGNMENT

In order to deploy the Internet Directory Service, it is important to define how the naming hierarchy will be structured. Although the basic model suggests a simple monolithic "database" containing all of the Internet's information infrastructure, with a namespace divided along geographic boundaries, this may not be the definite model that turns out to be the most appropriate to the Internet. Different models may evolve according to the needs of the Internet and the applications used on the Internet (i.e., some parts of the DIT may be assigned at the root for the Internet). Below this one can envisage several loosely coupled namespaces each with their own area of applicability. This should be handled as a part of the general operation of a directory service. An example of this might be assignment of a representation of the Domain Namespace under the root of the DIT. This is further discussed in [BHK91a].

However, the core DIT information will be nationally assigned. The parts of the DIT below country level will be managed differently in each country. In many countries, registration authorities will be established according to the OSI Standard [ISO]. This has been done in some countries by the national ISO member body representative (for example in the UK by BSI).

The lower parts of the hierarchy will, in general, be delegated to organizations who will have control over Name Assignment in that part of the tree. There is no reason to mandate how to assign this hierarchy, although it is appropriate to give guidelines. Proposed solutions to assignment of namespace are given in [BHK92].

In North America, there is an alternative approach being developed by the North American Directory Forum (NADF), which leverages existing registration mechanisms [For91]. It is not yet clear what form a final North American Directory Service will take. It is expected that similar initiatives will be taken in other places, such as Europe. For the Internet, the Internet Society (ISOC) has been suggested as a possible Naming Authority.

A discussion of the main issues involved with representing the Real World in the Directory Service is part of the work undertaken by the IETF OSI DS Working Group.

The core of the Internet Directory will therefore come to exist of a country based structure with different national naming schemes below the countries. It is clearly desirable that the Internet Directory Service follows any evolving national and international hierarchies. However, this should not be allowed to cause undue delay. The strategy proposed is to proceed with name assignment as needed, and to establish interim registration authorities where necessary, taking practical steps to be aligned with emerging national authorities wherever possible.

It is suggested that the Internet Directory Service does two things:

First, each national part of the Internet DIT namespace should be delegated to an appropriate organization, which will usually be in the country of question. Second, the delegated organization should assign names for that country as part of the Internet Directory Service. This should be done in a manner which is appropriately aligned with any emerging local or national service, but does not unduly delay the deployment of the Internet Directory Service. For most countries, this will fit in as a natural evolution of the early directory piloting, where operators of pilots have acted as interim name registration authorities.

5. DIRECTORY INFRASTRUCTURE

To provide access to the Internet Directory Service, an infrastructure has to be built. Although the technical components of an X.500 infrastructure are clear: DSAs (that hold the actual data) and DUAs (that allow users and applications to access the data), a lot more is needed for deployment of an Internet Directory Service.

The Integrated Directory Services (IDS) Working Group of the IETF is playing a key role in solving most of the issues that are related to the building of an appropriate infrastructure.

Many of the issues cited in this section have come forward out of interim pilots that have been established on the Internet:

PSI White Pages Pilot

This is a pilot service which is operating X.500 on the Internet. In many ways it is operating as an Internet wide pilot.

FOX

Fielding Operational X.500, a project to explore the development and interoperability of X.500 implementations.

Paradise (Piloting A ReseArch DIrectory Service in Europe)

This project has been providing the necessary glue to hold the various national activities together [Par91].

5.1 Short Term Requirements

- Central Operations. There is a need for a number of operations to be managed as a service for the whole Internet. These services are:
 - o A root DSA; containing the top-level of the DIT, has to be provided. Currently, this root DSA is managed by the Paradise project.
 - o Name assignment; Inserting names into the Directory, this has been discussed in section 4. This could be done in conjunction with the appropriate Registration Authority or by the Registration Authority. In most cases it is likely to be the former, and mechanisms will need to be set up to allow organizations to get their names installed into the directory, either direct or through the registration authority.
 - o Knowledge management; i.e., the information on "which DSA holds what part of the DIT, and how can that DSA be accessed". DSAs will be established by Organizations. There will be a need to centrally coordinate the management of the knowledge information associated with these DSAs. This is likely to be coupled to the name assignment.
 - o Knowledge and Data replication; For the Directory to perform well, knowledge and data high up in the DIT must be significantly replicated. A service must be provided to make replicated information available to DSAs that need it.

It is suggested that for the time being, Paradise should be used as the initial basis for handling the top-level of the DIT and for provision of the central services. However, the services mentioned above need to be provided at a national level for every participating country in the Internet Directory Service. Whenever an organization starts a new country branch of the DIT in the Internet Directory Service the central operations will have to help out to make sure that these services will be properly installed on a national level.

- An effective service will need to have sufficient implementations, in order to give full coverage over different hardware and software platforms, and to demonstrate openness. The recent Directory Information Services (pilot) Infrastructure Working Group's (DISI) Survey of Directory Implementations suggests that there will not be a problem here. This provides a list of available X.500 implementations and their capabilities [LW91].
- An executive summary, necessary to convince the management of computer centers to invest manpower into setting up a X.500 Directory Service. This is provided by DISI [WR92].
- Due to the possible different and rather independent structured namespaces that can be envisaged in the DIT for different purposes, DUAs will have to be "tuned intelligently" for the applications that they are used for.
- To allow users easy access to the Internet Directory Service even from low powered workstations, a lightweight protocol has to be developed over TCP/IP. Already two private protocols that do this have been developed: The Michigan DIXIE protocol [HSB91] and the PSI Directory Assistance Service [Ros91]. The IETF OSI Directory Services Working Group (OSI-DS WG) is currently working on a standard lightweight protocol called LDAP.
- Although the Internet Directory Service does not have to make any mandatory requirements about the use of lower layers, it is noted that the use of STD 35, RFC 1006 to allow use of OSI applications on top of TCP/IP is essential for deployment in the Internet. Other stacks like the ones using CLNS, CONS and X.25(80) will probably also be deployed in parts of the Internet. DSAs with different stacks will be linked through use of either application level relays (chaining) or Transport Service bridges.
- There are multiple issues that are not dealt with (properly) in the X.500 standard and thus prevent the building of an Internet Directory service. Intermediate solutions for these issues have to be established in an "open" way. The results will have to be

deployed as well as to be fed back into the relevant standard committees. The IETF OSI-DS WG deals with these issues. Section 7 describes several of these issues.

- Site support. The IETF IDS WG is looking at providing the necessary documentation to help with the provision of support for Directory users at participating sites.

5.2 Medium Term Requirements

- Enhanced performance is necessary to allow for a real global usage;
- The schema has to be extended to allow for various kinds of data, e.g., :
 - o NIC data;
 - o Resource location;
- Support for Internet Message Handling services (RFC-822, MIME and X.400). This work is already undertaken by the IETF MHS-DS WG.

5.3 Long Term Requirements

- To make sure that X.500 evolves into an operational service, it is essential to track its evolution, and to feed back into the evolution process.
- Interface existing RDBMS into the Directory Service.
- To increase the performance of the directory, and thereby making it useful for an even wider range of applications (e.g., policy based routing), a lightweight protocol for access and system usage is needed.

6. DATAMANAGEMENT

The whole of the Directory Infrastructure won't stand much chance without proper datamanagement of the data contained within the DIT. Procedures need to be established to assure a certain Level of Quality of the data contained in the DIT.

Due to the very nature of X.500, the management of the data is distributed over various sources. This has the obvious advantage that the data will be maintained by the owner of the data. It does however, make it quite impossible to describe one single procedure for datamanagement.

For the Internet Directory Service, guidelines will have to be developed (by the IETF IDS WG), to help organizations that start with deployment of X.500 on how to manage data in their part of the DIT. The guidelines should describe a minimum level of quality that has to be supplied to make the service operational. The IETF OSI-DS WG will initiate a pilot on Quality of Service parameters in the Directory, that will be of use.

Pilot datamanagement projects will have to be done (e.g., existing databases should be connected to the Internet Directory Service). Tools that are developed to achieve this should be made available to the Internet community for possible future use.

6.1 Legal Issues

Most countries connected to the Internet have some sort of law that dictates how data on people can and cannot be made available. These laws deal with privacy and registration issues, and will differ from country to country. It is suggested that each of the national organizations within the Internet that manages the Internet Directory Services master for that country, undertake some research as to the applicability of laws within that country on data made public through use of X.500.

In the mean time, a general "User Bill of Rights" should be established to indicate what the proper use of the Internet Directory Service is. This "Bill of Rights" could be drafted by the IETF IDS WG. As a basis, the NADF "User Bill of Rights" [For92] can be used.

7. TECHNICAL ISSUES

The IETF has established the OSI-DS WG. The major component of the initial work of this group is to establish a technical framework for deploying a Directory Service on the Internet, making use of the X.500 protocols and services [CCI88b]. This section describes the work already done by this working group, which has been implicitly focused on the technical infrastructure needed to deploy the Internet Directory service.

The OSI Directory Standards do not yet contain sufficient specifics to enable the Internet Directory Service to be built. Full openness and interoperability are a key goal, so we may need Internet specific agreements, at least until the ISO standards are more complete. This section notes areas where the standards do not have sufficient coverage, and indicates the RFCs which have been written to overcome these problems.

The work is being limited to (reasonably well) understood issues.

This means that whilst we will attempt to solve a wider range of problems, not all potential requirements will necessarily be met.

The technical work is done in conjunction with the RARE WG on Network Application Support WG (formerly RARE WG3). The IETF WGs and the RARE WG have a common technical mailing list. It is intended that this will lead to a common European and North American technical approach.

7.1 Schema

A Directory needs to be used in the context of an Information Framework. The standard directory provides a number of attributes and object classes to enable basic operation. It is certain that the Internet community will have requirements for additional attributes and object classes. There is a need to establish a mechanism to register such information.

Pilots in the European RARE Community and the US PSI White Pages Pilot have based their information framework on the THORN and RARE Naming Architecture. This architecture should be used for the Internet Directory Service, in conjunction with COSINE based services in Europe. A revised version of the Naming Architecture, with a mechanism for registration of new attributes and object classes, has been released as RFC 1274 [BHK91a].

7.2 Use on the Internet

It is a recognized policy on the Internet to deploy OSI Applications over non-OSI lower layers (such as STD 35, RFC 1006) [RC87]. This policy allows deployment of OSI Applications before an OSI lower layer infrastructure has been deployed. Thus, the Internet Directory Service will decouple deployment of the OSI Directory from deployment of the OSI lower layers. As the Internet Directory service will extend into the far corners of the Internet namespace, where the underlying technology is not always TCP/IP, the Internet Directory Service will not make any mandatory requirements about use of lower layers. When configuring the Internet Directory Services, variations in the lower layers must be considered. The following options are possible:

- Operation on top of TCP/IP using a lightweight protocol.
- Operation over TCP/IP using STD 35, RFC 1006. This is a practical requirement of deployment at very many Internet sites, and is the basis of the existing services. It is highly recommended that all participating DSAs support this stack.
- Use of OSI Network Service (Connection Oriented or Connectionless).

- X.25(80) will probably not be used in the core infrastructure of the Internet Directory Service, but is the basis of some European activities. It may be needed later to interconnect with US commercial systems not on the Internet. There will be a practical need to interwork with DSAs which only support this stack.

This approach has the following implications:

1. There is a need to represent TCP/IP addresses within OSI Network Addresses. This is specified in RFC 1277 [HK91a].
2. It will be desirable to have a uniform method to present Network Addresses of this style. Therefore, a string representation of presentation addresses is specified in RFC 1278 [HK91d].
3. This approach leads to the situation where not all DSAs can communicate directly due the different choice of lower layers. This is already a practical result of many European sites operating DSAs over X.25. When the Internet Directory Service is deployed, the issue of which DSAs operate which stacks must be considered in order to achieve a coherent service. In particular, there may be a need to require DSAs that serve parts higher up in the DIT to serve multiple stacks. This will be tackled as an operational issue.
4. There may be a requirement to extend the distributed operations, so that there is no requirement for full connectivity (i.e., each DSA supports each stack). A solution to this problem, by defining "relay DSAs" is specified in RFC 1276 [HK91b].

7.3 Replication of Knowledge and Data

There are a number of requirements on replication, both of data (the actual information on objects in the directory) and knowledge (the information on where do I find what data) information, which must be met before an Internet Directory can be deployed. The 1988 standard cannot be used as is, because it does not deal with replication or caching. This leads to serious problems with performance. There is a partial solution available in the 1992 version of the standard, however there are no products available yet that implement this solution. These issues are discussed in more detail in RFC 1275 [HK91c].

As it took too long for 1992 implementations to arrive to be of any help to the already rapidly growing pilot that urgently needed a solution, an option was chosen to use a simple interim approach as defined in RFC 1276. It will be clearly emphasized that this is an interim approach, which will be phased out as soon as the appropriate standards are available and stable implementations are deployed. The

interim approach is based on the approach used in the QUIPU Implementation and it is widely deployed in the existing pilots.

7.4 Presentation of Directory Names

The standard does not specify a means to present directory names to the user. This is seen as a serious deficiency, and a standard for presenting directory names is required. For Distinguished Names, a string representation is defined in [HK92a]. However, as the distinguished name is not very friendly for the user, a more user oriented specification of a standard format for representing names, and procedures to resolve them is chosen on the Internet, and is specified in [HK92b].

7.5 DSA Naming and MD Structure

There are some critical issues related to naming of DSAs and the structure of Directory Management Domains. The main issues are:

- It is hard to achieve very high replication of knowledge information as this is very widely spread;
- There is a need to give DSAs more reasonable names, which will contain an indication on the role of the DSA; This is necessary for DSAs high up the DIT.
- There is too much DIT clutter in the current pilots;
- There is no real concept of a DMD (Directory Management Domain) authority.

These will be significant as the directory increases in size by orders of magnitude. The IETF OSI-DS WG is working to develop a solution in this area.

8. SECURITY

A Directory can be an important component in the overall provision of security in a distributed system environment, especially when public-key cryptographic technology is employed. The directory can serve as a repository for authentication information, which, in turn, forms the basis of a number of OSI Authentication Services (e.g., X.400) and non-OSI Services (e.g., privacy-enhanced mail, PEM). The directory may also use this and other stored authentication information to provide a wide range of security Services used by the Directory system itself.

8.1 Directory Provision of Authentication

The directory will be used to provide X.509 strong authentication. This places minimal requirements on the directory. To use this infrastructure, users of authentication services must have access to the directory. In practice, this type of authentication can be deployed only on a limited scale without use of a directory, and so this provision is critical for applications such as Privacy Enhanced Mail [Lin93]. The PEM development is considering issues relating to deploying Certification Authorities, and this discussion is not duplicated here.

PEM defines a key management architecture based on the use of public-key certificates, in support of the message encipherment and authentication procedures defined in [Lin93]. The PEM certificate management design [Ken93] makes use of the authentication framework defined by X.509. In this framework, as adopted by PEM, a "certification authority" representing an organization applies a digital signature to a collection of data consisting of a user's public component, various information that serves to identify the user, and the identity of the organization whose signature is affixed. This establishes a binding between these user credentials, the user's public component and the organization which vouches for this binding. The resulting, signed, data item is called a certificate. The organization identified as the certifying authority for the certificate is the "issuer" of that certificate. The format of the certificate is defined in X.509.

In signing the certificate, the certification authority vouches for the user's identification, in the context specified by the identity of the issuer. Various types of organization may issue certificates, including corporate, educational, professional, or governmental entities. Moreover, these issuers may operate under different certification policies, so that not all certificates may be equally credible (i.e., some certificates may be more trustworthy as accurate identifiers of users, organizations, mailing lists, etc). The PEM certificate management design allows for this diversity of certification policies, while ensuring that any certificate can be traced unambiguously to the policy under which it was issued.

The digital signature is affixed on behalf of that organization and is in a form which can be recognized by all members of the privacy-enhanced electronic mail community. This ability to universally verify any PEM certificate results because the PEM certification design is a singly rooted tree, in which the Internet Society acts as the root. Once generated, certificates can be stored in directory servers, transmitted via unsecure message exchanges, or distributed via any other means that make certificates easily accessible to

message originators, without regard for the security of the transmission medium.

8.2 Directory Security

A number of security services are possible with the directory:

Peer Authentication at Bind

Authentication (one or two way) between DUA/DSA and DSA/DSA, established during the bind operation. This authentication may be provided using simple passwords (not recommended), one-way hashed passwords (more secure), or via public key cryptography (most secure). The various authentication options are specified in X.500(88), but most existent implementations implement only simple password authentication.

Per-operation Authentication and Integrity

This is usually used to identify the DUA originating an operation to the Directory (e.g., to authenticate prior to data modification). It may also be used to verify the identity of the DSA which provided data in a response to the user. In both examples, the integrity of the data also is ensured through the use of digital signatures. This is specified in X.500(88), but not yet widely implemented.

Single Entry Access Control

This is used to control which users (DUAs) can access and modify data within an entry. This is specified in X.500(92) and most DSA implementations provide this function.

Multiple Entry Access Control

This is used to control search and list operations, in order to allow location of information by searching, but to deter "trawling" of information and organizational structure. Usually, these access controls are limited in their ability to prevent trawling because of the conflicting goal of allowing a certain level of legitimate browsing in support of "white pages" functionality.

Service Authorization

This allows DSAs to control service in a data independent manner, based on peer authentication. For example, one might prevent students from making non-local queries, while permitting such queries by faculty and staff.

Security Policy

This term encompasses the security goals for which data access control, service authorization, and authentication mechanisms are used to implement. For example, a local security policy might require that all directory database modifications employ strong authentication and originate from a computer at a known (local) location.

Data Confidentiality

The directory does not include explicit features to protect the confidentiality of data while in transit (e.g., between a DUA and DSA or between DSAs). Instead, it is assured that lower layer security protocols or other local security facilities will be employed to provide this security service. Ongoing work on adaptation of the Network Layer Security Protocol (NLSP) is a candidate for provision of this security service with directories.

There is no specification of any Internet-wide security policy for directories, nor are there currently any security mechanisms required of all directories. Deployment of a directory could be based on a variety of policies:

- Read only system, containing only public data and restricted to local modification.
- Use of X.509 authentication, and private access control mechanisms (this will not allow open access control management, but this is not seen as a fundamental problem).

It will be important to understand if global Internet requirements for minimum essential directory security mechanisms will be required to promote widespread use of directories. We recommend that an informational RFC be written to analyze this issue, with an operational policy guidelines or applicability statement RFC to follow.

9. RELATION TO DNS

It is important to establish the relationship between the proposed Internet Directory, and the existing Domain Name System. An Experimental Protocol RFC (RFC 1279) proposes a mapping of DNS information onto the Directory. Experiments should be conducted in this area [HK91e].

10. EXTERNAL CONNECTIONS

It will be important for this activity to maintain suitable external liaisons. In particular to:

Other Directory Services and Directory Pilots

To ensure a service which is coherent with other groups building X.500 services. e.g.,:

- Paradise
- NADF
- FOX
- PSI White Pages

Standards Bodies

To feed back experience gained from this activity, so that the next round of standards meets as many of the Internet requirements as possible. e.g.,:

- CCITT/ISO
- RARE WG-NAS
- EWOS/OIW
- ETSI

11. REFERENCES

- [BHK91a] Barker, P., and S. Hardcastle-Kille, "The COSINE and Internet X.500 Schema", RFC 1274, Department of Computer Science, University College London, November 1991.
- [BHK92] Barker, P., and S. Hardcastle-Kille, "Naming Guidelines for Directory Pilots", RFC 1384, Department of Computer Science, University College London, ISODE Consortium, January 1993.
- [CCI88a] The Directory --- authentication framework, December 1988. CCITT Recommendation X.509.
- [CCI88b] The Directory --- overview of concepts, models and services, December 1988. CCITT X.500 Series Recommendations.
- [CCI90] The Directory --- part 9 --- replication, October 1990. ISO/IEC CD 9594-9 Ottawa output.
- [CFSD90] Case, J., Fedor, M., Schoffstall, M., and J. Davin, "A Simple Network Management Protocol", STD 15, RFC 1157, SNMP Research, Performance Systems International, MIT Laboratory for Computer Science, May 1990.

- [For91] The North American Directory Forum, "A Naming Scheme for C=US", RFC 1255, NADF, September 1991. Also NADF-175. (See also RFC 1417.)
- [For92] The North American directory Forum, "User Bill of Rights for Entries and Listing in the Public Directory", RFC 1295, NADF, January 1992. (See also RFC 1417.)
- [HK91a] Hardcastle-Kille, S., "Encoding network addresses to support operation over non-OSI lower layers", RFC 1277, Department of Computer Science, University College London, November 1991.
- [HK91b] Hardcastle-Kille, S., "Replication and distributed operations extensions to provide an internet directory using X.500", RFC 1276, Department of Computer Science, University College London, November 1991.
- [HK91c] Hardcastle-Kille, S., "Replication requirement to provide an internet directory using X.500", RFC 1275, Department of Computer Science, University College London, November 1991.
- [HK91d] Hardcastle-Kille, S., "A string encoding of presentation address", RFC 1278, Department of Computer Science, University College London, November 1991.
- [HK91e] Hardcastle-Kille, S., "X.500 and domains", RFC 1279, Department of Computer Science, University College London, November 1991.
- [HK92a] Hardcastle-Kille, S., "A string representation of Distinguished Names", Department of Computer Science, University College London, Work in Progress.
- [HK92b] Hardcastle-Kille, S., "Using the OSI directory to achieve user friendly naming", Department of Computer Science, University College London, Work in Progress.
- [HSB91] Howes, R., Smith, M., and B. Beecher, "DIXIE Protocol Specification", RFC 1249, University of Michigan, July 1991.
- [ISO] Procedures for the operation of OSI registration authorities --- part 1: general procedures. ISO/IEC 9834-1.

- [Ken93] Kent, S., "Privacy Enhancement for Internet Electronic Mail: Part II - Certificate-based Key Management, RFC 1422, BBN, February 1993.
- [Kil88] Kille, S., "The QUIPU Directory Service", In IFIP WG 6.5 Conference on Message Handling Systems and Distributed Applications, pages 173--186. North Holland Publishing, October 1988.
- [Kil89] Kille, S., "The THORN and RARE Naming Architecture", Technical report, Department of Computer Science, University College London, June 1989. THORN Report UCL-64 (version 2).
- [Lin93] Linn, J., "Privacy Enhancement for Internet Electronic Mail: Part I - Message Encryption and Authentication Procedures", RFC 1421, February 1993.
- [LW91] Lang, R., and R. Wright, "A Catalog of Available X.500 Implementations", FYI 11, RFC 1292, SRI International, Lawrence Berkeley Laboratory, January 1992.
- [Lyn91] Lynch, C., "The Z39.50 information retrieval protocol: An overview and status report", Computer Communication Review, 21(1):58--70, January 1991.
- [Par91] Paradise International Report, Cosine. Paradise project, Department of Computer Science, University College London. November 1991.
- [RC87] Rose, M., and D. Cass, "ISO Transport Services on top of the TCP", STD 35, RFC 1006, Northrop Corporation Technology Center, May 1987.
- [Ros91] Rose, M., "Directory Assistance Service", RFC 1202, Performance Systems International, February 1991.
- [WR92] Weider, C., and J. Reynolds, "Executive Introduction to Directory Services Using the X.500 Protocol", FYI 13, RFC 1308, ANS, ISI, March 1992.

12. Security Considerations

Security issues are discussed in Section 8.

13. Authors' Addresses

Steve Hardcastle-Kille
ISODE Consortium
PO box 505
SW11 1DX London
England
Phone: +44-71-223-4062
EMail: S.Kille@isode.com

Erik Huizer
SURFnet bv
PO box 19035
3501 DA Utrecht
The Netherlands
Phone: +31-30 310290
Email: Erik.Huizer@SURFnet.nl

Vinton Cerf
Corporation for National Research Initiatives
1895 Preston White Drive, Suite 100
Reston, VA 22091
Phone: (703) 620-8990
EMail: vcerf@cnri.reston.va.us

Russ Hobby
University of California, Davis
Computing Services
Surge II Room 1400
Davis, CA 95616
Phone: (916) 752-0236
EMail: rdhobby@ucdavis.edu

Steve Kent
Bolt, Beranek, and Newman
50 Moulton Street
Cambridge, MA 02138
Phone: (617) 873-3988
EMail: skent@bbn.com