

Definition of an incorporation and usage policy for a Directory Schema Registry

TERENA Project Directory Schema Registry, Deliverable D

Peter Gietz, DAASI International Ltd., peter.gietz@daasi.de

Version 0.9, 10.2.2003

1. Status of this document

This is deliverable D of the TERENA project Directory Schema Registry, which is co-funded by TERENA, JISC (Joint Information Systems Committee, UK), REDIRIS (Spanish National Research Network), CESNET (Czech National Research Network), POZMAN SUPERCOMPUTING (Poznan Supercomputing and Networking Center, Poland) and DAASI International and performed by DAASI International. Together with four other deliverables [RegIntro], [RegSchema], [RegArchitecture], [RegBusiness] and the bibliography [RegBib] it forms the documentation of this project.

This version 0.9 reflects the status in January 2003 and is the first public version. A final version 1 will be published after all other project documents have been completed. That final version will reflect the comments made by the funding board of the project.

In some parts of this text, some details not relevant for the currently planned registry are discussed which points to a possible future broadening of the organisational scope of the registry after the end of this project. A version 2 of this document is planned for the time after the end of the project. For this future version comments and additions to the current document are welcomed. Please send them to the email address of the author.

Table of Contents

1.	Status of this document.....	1
2.	Introduction.....	2
3.	Lessons learned from the registry survey.....	2
4.	Procedures specified by the IETF schema WG.....	3
5.	List of Requirements for a policy of the DSR.....	8
6.	Policy specifications	9
7.	Conclusion.....	14
8.	Pointer to references	14
A.	Email from the IETF Application Area Director about the closure of the IETF schema WG	15
B.	Listing Request Formats according to [SchemaProc].....	16
C.	Detailed Table contents.....	17

2. Introduction

[RegIntro] discusses technologies and organisations relevant in the schema registry problem space. It also gives a survey of existing LDAP schema and of prior approaches for solving the schema registry problem. Chapter 3 of this document discusses the policy aspects of the registries described in the latter survey. Chapter 4 describes in detail policies and procedures specified by the IETF schema WG which aimed at the establishment of a directory schema listing service. This discussion is followed by a list of policy requirements for the Directory Schema Registry (DSR) in chapter 5. Chapter 6, the main part of this document, specifies all policy related rules for the DSR, including rules for a policy board, a schema registry operator, a registration process and for update and comments. These proposals will be the basis for the initial pilot phase of registry operation which is strongly related to the funding board of this project. Possible future alternatives are also mentioned. A short summary concludes this document.

All references can be found in [RegBib]. A number of acronyms are being used throughout the project documentation. Please view chapter 3 of [RegIntro] for short explanations of these acronyms. Appendix A of [RegIntro] consists of an alphabetical list of these acronyms which provides the reader with pointers to the places of that document where they are discussed.

3. Lessons learned from the registry survey

This chapter summarises the policy relevant parts of [RegIntro].

3.1. Registries in the IETF

IETF is the most important although not the only organisation in which relevant LDAP schema is being defined as could be shown in [RegIntro]. The initial strategy of the IETF for LDAP schema standardisation was to set up working groups (e.g. asid, ldaptext) which provided a forum to among other things discuss schema standardisation. An IETF WG is conceived as a short term living entity. Since new schema and new LDAP extensions will be produced for a long time, another strategy had to be found.

As described in section 5.3 of [RegIntro], the IETF schema WG proposed a strategy for a schema listing service that makes schema publishing possible without the RFC track. The policy aspects of this proposal are discussed below in chapter 4.

As described in section 5.2 of [RegIntro], [RFC 3383] specifies a process for registering extensible elements of LDAP in accordance with the guidelines for writing IANA considerations in RFCs [RFC 2434]. Since schema-wise only OIDs and Object Identifier descriptors can be registered with this process, it does not fulfil the requirements for a schema registry.

[RegIntro] in sections also gave an overview of other IANA registries. The main characteristics of the IANA registries discussed are the following:

- Process is controlled by "internet officials", namely IESG, Area Directors and so called experts appointed by the former.
- Public discussion on a mailing list with a fixed short time slot (mostly 2 weeks)
- A well defined process for new registry entries
 - several steps from proposal to registration
 - rules for accepting or rejecting registration applications
 - classification of different statuses like standard, private, experimental, etc.
- A well defined update mechanism

- The retrieval aspect of a registry problem space is not treated by the IETF specifications of IANA registries inspected, or only defined in a very general way like "assigned values can be seen at <URL>".
- The process depends on active involvement of experts, no formal quality checks are recommended, e.g. an automated parsing of defined structures with a check, whether the syntax is in accordance to the LDAP standard.

3.2. Other registries

The short survey of other registries in [RegIntrod] showed following policy relevant aspects:

- most registries require a more or less elaborated set of metadata.
- In the case of the XML registry named XML.org a part of these data consist of key words and a classification of industrial areas
- In XML.org only registered participants are allowed to register schema
- The Novell OID registry requires a fee for registration of schema
- Both discussed OID registries provide valuable additional information, especially about the owner of the respective OID subtree. Thus pointers to these registries should be included in the DSR.
- The LDAP schema viewer discussed in [RegIntrod] provides a useful interface and may contain valuable additional information, e.g. in a comment given there. Thus pointers to the respective schema element should be included in the DSR.

4. Procedures specified by the IETF schema WG

As described in [RegIntrod], the specifications of the IETF schema WG, although mostly expired, specify a useful framework for a directory schema listing service. Even if they have not been processed to an RFC, these specifications have reached a certain degree of Internet consensus. The webpage and mailing list archive of the schema WG are still online at: <http://www.imc.org/ietf-schema-reg/index.html>. The last thread on this mailing list is a statement from the IETF Application Area Director on the closure of the schema WG dated February 2000 which is quoted below in Appendix A.

The term directory for the schema WG subsumes technologies like LDAP, Whois, Whois++ and Rwhois. The parts relevant to LDAP schema are quoted and discussed here as input for policy requirements of the DSR.

One of the specifications of the schema WG [SchemaProc] deals solely with policy aspects of the schema listing service. This will be discussed in detail in this chapter. All quotations in this chapter are from [SchemaProc] if not stated otherwise.

4.1. Requirements specified

[SchemReq] also specifies general requirements for the schema listing procedures:

"The schema listing procedures SHOULD be designed to enable:

- a) verification that all properly formed schema listing requests are submitted by the contact person claiming to originate them
- b) methods of ensuring that only properly-formed, high-quality directory schema are published in the schema listing repository
- c) verification that requests to change the identity of the contact person for a listing originate from the listing authority contact or the contact person

- d) coping with the situation in which the contact person and/or listing authority contact for a schema is no longer available or is unable to submit updates to the listing for which they hold update authority".

In addition to these general requirements, [SchemaProc] specifies following more specific requirements:

1. Functionality Requirements

Schema unit is a technical term defined by [SchemaProc] as

"a related or grouped set of object attributes that form a discrete unit within the context of a schema for a particular protocol; examples include an LDAP object class."

Schema unit content is defined as

"a formal specification of a schema unit using a profile of [MIMEDIR]"

Following functionality requirements are specified in [SchemaProc]:

"Schema unit listings MUST include two different types of information:

- (1) metadata
- (2) schema unit content

Metadata will be used to catalog repository files by characteristics that differentiate listings.

Schema unit content MUST be limited to the specification of a single schema unit."

It is important that the schema can be retrieved. Thus the policy has to define a minimum of metadata that have to be submitted together with the schema.

[SchemaMeta] specifies the following text/directory MIME types for the metadata:

- listingName, listingTitle, listingUse, specFile, relatedTo,
- contactLanguage, contactName, contactEmail, contactPhone, contactAddress
- authLanguage, authName, authEmail, authPhone, authAddress
- specURL, security, created, moreInfo, caveat, listingComments, schemaPak, pakMember

For further discussion of these metadata see [RegSchema].

The reference for "[MIMEDIR]" is a draft of [RFC 2425], the specification of a MIME Content-Type for Directory Information which is discussed in section 5.3.3. of [RegIntrod]. The LDAP profile defined according to [RFC 2425] is specified in [RFC 2927] which is discussed in section 5.3.2. of [RegIntrod].

Thus the functionality requirements say that each listed LDAP object class has to be accompanied by metadata. The format required for sending schema unit is [RFC 2425], in the case of LDAP in the profile specified in [RFC 2927]. The format required for sending the metadata is specified in [SchemaMeta] which is discussed in [RegSchema].

Schema must be defined for a special purpose. A group of schema elements that fulfil such a purpose can be grouped in one schema. The schema WG used the term Schema Pak which they defined as:

"a related or grouped set of schema units that collectively specify a schema associated with a particular protocol"

The schema WG discerns between a primary repository,

"the repository that masters the schema listings database"

and shadow repositories,

"a repository that mirrors the primary repository".

Only the operator of the primary repository is relevant in the communication between the schema writer and the listing service.

2. Naming Requirements

"All listings MUST have a valid OID as their name.

The primary listing repository operator MUST provide schema writers with the following components of a listing request name:

- + a sequence number assigned to each listing by the primary repository operator
- + a version number assigned to each listing version by the primary repository operator."

The operator of the primary repository is responsible for assigning unique listing and version numbers.

3. Content Formatting and Transfer Encoding Requirements

"All listings MUST employ a common data format. Metadata and schema unit content format and transfer encoding MUST utilize appropriate [...] profiles.

Appropriate profiles are profiles specified in accordance with [RFC 2425]. Thus the requirement is that format and encoding for the schema unit has to be as specified in [RFC 2927], and for the metadata as specified in [SchemaMeta].

4. Security requirements

Security considerations have to be included in the schema documentation. These may be extended by the "comments on schema listing" mechanism (see below). As issues that should be considered following two are listed:

"(1) A listing might include specifications mandating exposure of certain attributes which would result in compromising the privacy of an organization or individual.

(2) A listing might be intended for use by applications requiring some sort of security assurances not provided by the schema specified in the schema unit content or in the schema unit content files referenced in a schema pak listing."

5. Publication Requirements

"Requests for schema listed in the IETF schema listing service MAY be published as RFCs"

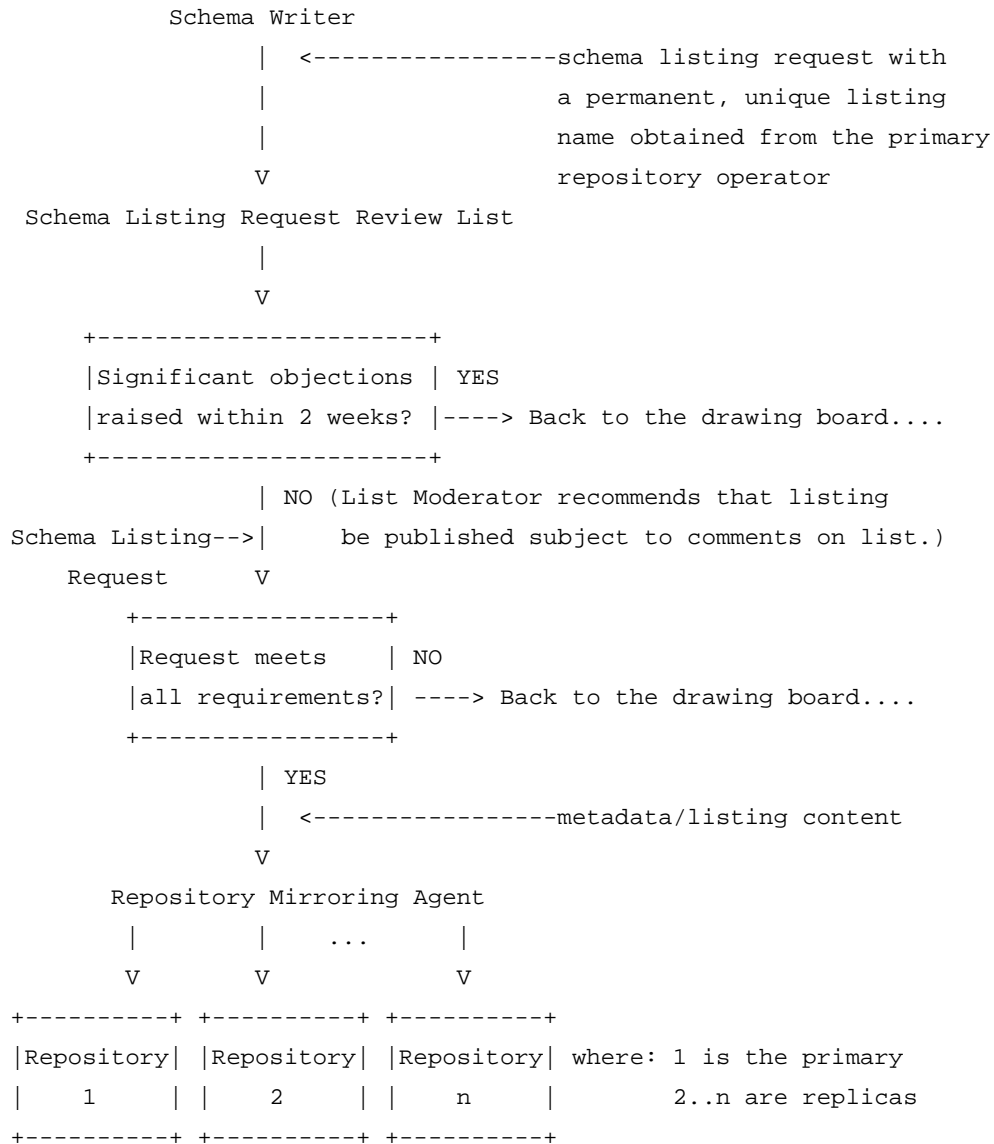
"The listing of a schema does not imply endorsement, approval, or recommendation by the IETF or even certification that the specification is adequate for the intended use of the schema."

These statements make clear that a listing in the repository is not equivalent to an IETF standardisation process. On the contrary, the "convenient" schema listing service is also meant for schema that is not intended to go through the "too difficult and too lengthy a process" of IETF standardisation.

4.2. Overview of the listing request process

Besides these requirements, [SchemaProc] specifies the listing procedure by first showing an ASCII architecture diagram, which is reproduced here:

Schema Listing Request Architecture Diagram [SchemaProc]:



4.3. Process for inclusion of new schema

The following process is defined for inclusion of new schema into the schema listing service:

1. Listing request can function as alternative to RFC publication. Schema unit content, schema pak listings, and/or other supporting material may additionally be published as an Informational RFC.

2. Schema MUST be specified according to [RFC 2927], Metadata specified according to [SchemaMeta] MUST also be included.
3. Listing request must be posted to the mailing list (the original mailing list thought of was ietf-schema-review@pilot.schema.dir.reg.int) where moderated discussion takes place.
4. During a comment period of at least 2 weeks consensus is gauged "as to whether or not the schema being proposed is bogus".
5. If there is no reasonable objection against the schema, the mailing list moderator may submit the listing request to the primary listing repository operator.
6. The primary listing repository operator checks listing requests for validity and makes the listings available to the community. Following requirements apply:
 - a) All listings MUST have a valid OID as their name.
 - b) All schema unit listing requests MUST include both metadata and schema unit content.
 - c) All schema pak listing requests MUST be limited to metadata.
 - d) Metadata MUST comply with [SchemaMeta].
 - e) "Any security considerations given must not be obviously bogus. [...]the listing request review committee (the members of the listing request review mailing list) has the authority and expertise to identify obviously incompetent material and exclude it."
 - f) Schema listing requests MAY be signed using PGP/MIME as described in [RFC2015]. The primary listing repository operator MUST be able to accept PGP/MIME, but is NOT REQUIRED to validate or retain the signature.
 - g) Listing request MUST be formatted according to the specification of [SchemaProc] reproduced in Appendix B.
 - h) "If a listing request includes one or more URI-based references to information that would not be included in a resulting listing, but is associated with the schema or schema unit specified by the request, a fingerprint of this information MUST be included with each such reference".
7. Listings are supposed to be posted in an anonymous FTP directory. The FTP URL specified (<ftp://www.pilot.schema.dir.reg.int/schema/>) does not exist, as expected.
8. All listings will be summarised in a periodically issued RFC.

4.4. Change control and comments mechanism

Following specifications are made concerning change control and comments on already included schema.

1. The metadata for the schema listing defined by [SchemaMeta] include a contact person who holds the authority to update a listing and who should be contacted if questions or concerns arise related to a listing or listing request .
2. The contact person for a listing can pass responsibility for a listing to another person or agency by informing the primary listing repository operator and the mailing list.
3. The IESG may reassign responsibility for a listing, if e.g. the contact person has died, moved out of contact, etc.
4. As to comments:
 - a) Comments on listings may be submitted by members of the IETF community for consideration by the community and the primary listing repository operator.

- b) Such comments will be passed on to the contact person for the listing by the primary listing repository operator if possible.
 - c) Submitters of comments may request that their comment be attached to the listing itself especially if the contact person is not reacting.
 - d) If a consensus can be found on the mailing list, such a comment will be made accessible in conjunction with the listing itself.
5. As to change control:
- a) The contact person may request a change to the schema definition and thus starts the same procedure as the listing request (see above).
 - b) Reasons for change requests may be
 - yet undetected serious omissions or errors in the published listing
 - an emerging need of the user community for a listed schema which cannot be addressed by that listed schema in its present form
 - c) a change request may be denied if it renders entities that were valid under the previous definition invalid under the new definition
 - d) Before implementing the changes, the primary listing repository provider SHOULD attempt to verify the authority of a person claiming to be the contact person.
 - e) Listings will not be deleted. They can instead be declared OBSOLETE.

4.5. Responsibilities of the Primary Repository Operator

Following responsibilities and constraints of the Primary Repository Operator (PRO) are specified:

- 6. The data residing in the repository is not the property of the PRO.
- 7. The intellectual property remains at the schema writer or his organisation.
- 8. All metadata surrounding the system is considered to be either in the public domain or is owned by the IANA.
- 9. The PRO can also make no determinations of appropriateness or suitability of a schema to be listed which is solely the task of the listing request review committee, the members of the mailing list.
- 10. If the list coordinator requests the PRO to publish a schema listing, the PRO MUST include the schema listing or be relieved of the responsibility of running the repository.
- 11. The PRO will take reasonable steps to ensure that information associated with the service is as accurate and authentic as possible.

5. List of Requirements for a policy of the DSR

Since the above described procedures as specified in [SchemaProc] represent a certain amount of IETF WG consensus and since such a new consensus process is not defined to be a part of this project, the processes of the DSR should be modelled as close to these specifications as possible. Additional features, especially for schema retrieval also have to be discussed in the DSR policy. For instance although of course careful reading by an expert as prescribed by the above discussed IETF procedures is always better than automated checks, the latter should be additionally included into registration policy.

In Summary following policy requirements apply:

- 1. Establishment of an open list for discussion about schema inclusion
- 2. Specification of a moderator who interacts with the DSR operator

3. Specification of the DSR moderator
4. Specification of the syntactical requirements for schema submission:
 - a) formats
 - b) encoding
 - c) naming
 - d) process for automatically checking syntax and OID.
5. Specification of semantic requirements for schema submission defining a mandatory minimal set of metadata for:
 - a) single schema elements
 - b) a whole schema, including the specification of different statuses, like "standard of <standardisation board>", recommended, experimental, individual
 - c) bibliographical data
 - d) additional information on author and contact person
6. Specification of a version control
7. Specification of the registration process
8. Specification of the comment mechanism
9. Specification of the update process
10. Specification of the actions and responsibilities of the DSR operator.

6. Policy specifications

This section contains all policy specifications for the DSR. A copy of the finalised version of this chapter will be published on the Website of the DSR.

6.1. Specification of a policy board

In the first phase of the DSR, a pilot service will be run. The policy board for this pilot should consist of members of the organisations funding this project which are TERENA, JISC, REDIRIS, CESNET, POZMAN Supercomputing, and DAASI International. Since DAASI International will be the operator of the schema registry it should not vote in decisions of this board, but could function as advisor, that can make proposals to the board. In fact, especially this document can be seen as such a proposal. DAASI can also be appointed as expert for schema review by the policy board.

This policy board as a whole or as a dedicated secretary appointed by the board is the formal review body. Its function is very close to the function of the IESG in the discussed IETF procedures. Thus the board will:

- finalise the decisions about the processes
- control the whole process
- appoint experts for review
- reassign responsibility for a schema
- moderate the discussion list
- decide about schema inclusion and classification of its status

In future phases of the operation of the DSR it will have to be decided who will take over the role of the policy board. Besides, with representatives of Internet2 who already showed interest in the future operation of the DSR, this will have to be discussed with IETF, especially IESG and IANA, as well as with other standardisation boards that may be interested in the DSR, like The Open Group, OASIS, DMTF and GGF.

6.2. Specification of the registry operator

The operator of the DSR is DAASI International.

Following functions are assigned to the operator:

1. Provide and run the technical infrastructure for operating the DSR:
 - a) adequate Hardware
 - b) good internet connectivity
 - c) directory server and additional software needed to provide the specified interfaces and check routines. In post pilot phases this includes appropriate replication mechanisms
 - d) Web server and portal to the DSR, including an LDAP2Web gateway
 - e) Mailing list for public review
2. Provide OIDs and additional numbers for uniquely identifying schema submissions, including versioning.
3. Perform the specified schema checks.
4. Forward schema registration requests to the policy board.
5. Include schema according to the instructions of the policy board.
6. Provide technical advice to the policy board.
7. Contribute to the dissemination of project results and to Public Relations of the DSR.
8. Act as a communication mediator between different interest groups.
9. As to the responsibilities and constraints of the registry operator, all rules for the PRO in [SchemaProc] as described above in section 4.5. apply.

6.3. Specification of the syntactical requirements for schema inclusion

Schema submitted to the DSR has to fulfil the syntactical requirements specified here:

1. The specification of the schema elements has to be submitted in one of the following formats:
 - a) ABNF definition in accordance with [RFC 2252] and in future with the final versions of [LDAPModels] and [LDAPSyntax].
 - b) Format for OpenLDAP schema files which is very close to a).
 - c) MIME format defined in [RFC 2927].
2. The encoding has to be in accordance with the specifications of the format, thus either to [RFC 2252] and its follow ups, or to [RFC 2927].
3. Every schema submission must have a unique name. This will be provided by using OIDs as specified in [SchemaFiles] which is discussed in [RegIntrod].

Two automated checks will be made before inclusion into the registry: One process checks the conformance to the format used for submission. One checks whether valid OIDs have been used for identifying the single schema elements.

6.4. Specification of the semantic requirements for schema inclusion

This section describes the metadata that have to be included in a schema submission. The exact format for specifying these metadata for both types of schema submission (see above 6.3. number 1 a)/b) for the first type and c) for the second) is specified in [RegSchema].

6.4.1. Semantic requirements for single schema elements

Every single schema element included into a schema submission has to be accompanied by following metadata:

1. Pointer to a specification document which describes the schema element.
2. In addition to the value of the DESC field which contains a short description of one sentence, an description of the functionality, use case, etc. This description will mostly be a short paragraph, as you can normally find in RFCs, describing the schema element.
3. Keywords. In the first phase of the DSR there are no restrictions to keywords. Future phases will at least maintain a controlled vocabulary, if not connect the keywords to an ontology.

6.4.2. Semantic requirements for a schema submission

Every schema submission has to be accompanied by the following metadata:

1. Metadata on the specification document, which describes the schema elements:
 - title
 - file name
 - version number
 - publication date
 - format (e.g.: ASCII, PS, PDF, HTML, ...)
 - name of the standardisation board (if applicable)
 - following metadata about the schema author and if not the same about the contact person:
 - name
 - email address
 - postal address
 - telephone number
 - fax number
 - affiliation
 - language (which the person understands)
 - An abstract about the motivation and use cases of the schema, as you can normally find in the abstract section of RFCs, describing the schema.
 - An abstract about security, as you can normally find in the security considerations section of RFCs, describing the schema. As a minimum form of such a statement, a sentence like " A security analysis was not performed" has to be included.
 - Keywords. In the first phase of the DSR there are no restrictions to keywords. Future phases will at least maintain a controlled vocabulary, if not connect the keywords to an ontology.

- RelatedTo statements that indicate a relationship of the schema to another published schema specification. The syntax of such statements is specified in [SchemaMeta] as described in [RegSchema]. It includes a classification of the relation which consists of the following values: "obsoletes" / "obsoleted-by" / "updates" / "inherits". New relation types may be added.
- The submitter must make a proposal for classification (the final decision about classification will be done by the policy board). Following statuses are defined:
 - "standard of <standardisation body>". In IETF terms this would include all RFCs in the Standard Track. Substatuses like "proposed standard", "recommendation" etc. may be used.
 - "publication of <standardisation body>". In IETF terms this would include all RFCs in the categories Experimental, Best current practice, Informational. Substatuses like "experimental" etc. may be used.
 - "specific for <vendor name>". This includes schema published by a single vendor outside of an standardisation body which in most cases will be schema for proprietary functionalities.
 - "private". All other new schema submissions by individuals or groups will be classified "private".
 - "draft". Any schema specification that is yet a draft document, like IETF Internet Drafts are classified as such. As subtypes of draft all above mentioned statuses may be used.
 - "obsoleted". If a schema is found to be either obsoleted by a new version or by the fact that its use in a changing IT world is no more relevant it is classified as obsoleted. Substatuses "by new version" and "by new technology" may be used. As example all IETF RFCs in the historical track will be classified as obsoleted.

6.5. Specification of the registration process

The process for registering schema at the DSR leans against the IETF processes, especially the process specified in [SchemaProc] as described above in section 4.3. The very few differences are explicitly mentioned here.

1. Schema registration is unrelated to the publication procedure of a standardisation board, but a specification document is required that describes the semantics and usage of the schema. Such a specification document has to describe every single element of the schema or refer to another specification document which contains such a description. Such specification documents must be freely available. The requirement for a specification document is not specified in [SchemaProc].
2. Schema must be specified as described above in sections 6.3 and 6.4.
3. Registration request must be posted to the mailing list schema-registry-review@daasi.de.
 - a. This is a moderated public list to which anyone can subscribe.
 - b. Postings on this list can only be made by list members.
 - c. The list may be filtered by spam and virus detection programs.
 - d. The list is moderated by a list moderator appointed by the policy board. In addition, a moderator for each posted schema may be appointed by the policy board who takes up the responsibilities of the moderator for one particular schema.
4. In addition to the posting to the mailing list, the schema may be sent to the registry operator, who will publish it in a separate part of the registry's database marked as pending.

- a. This publication is not a registration but is meant only for convenient reference until the status of registration is reached.
 - b. Such a publication is only possible if a minimum set of formal requirements are fulfilled.
 - c. This concept of pending data is not specified in [SchemaProc].
5. During a comment period of at least 2 weeks the schema posted for registration is being discussed, until a consensus has been established which is defined to be reached, if no unsolved reasonable objection remains. The moderator determines such a consensus.
6. If consensus to register the schema has been reached, the moderator submits the registration request to the schema registry operator.
7. The schema registry operator checks the registration request for validity, as defined above in sections 6.3 and 6.4. and publishes it in the registry.
 - a. If a pending entry of the schema has been published, as described in number 4., it will be deleted by the registry operator.
8. The contact person specified in the schema submission who holds the authority on the schema publication, can pass responsibility for the schema to another person or agency by informing the schema registry operator and the mailing list.
9. The rules of [SchemaProc] about PGP/MIME submissions and about fingerprint requirement for URI-based references as described above in section 4.3. numbers 6. f) and 6. h) are not adopted in the initial phase of the registry. Future versions of this text may adopt these rules.
10. The rule of [SchemaProc] on a periodically issued RFC that summarises the listings as described above in section 4.3. number 8 is outside the scope of this project.

6.6. Specification of the comment mechanism

The DSR maintains a comment mechanism which is divided into two categories, namely formal comments and informal comments.

6.6.1. Formal comment mechanism

The formal comment mechanism is modelled after [SchemaProc], as described above in section 4.4. number 4.

1. Anyone can send a formal comment on a registered schema to the registry operator.
2. Reasons for such a comment will generally be an objection to the specification that hadn't been raised during the comment period specified above in section 6.5. number 5.
3. The registry operator forwards this comment to the contact person for the schema in question as well as to the mailing list schema-registry-review@daasi.de described above in section 6.5. number 3.
4. Similiar to the process specified above in section 6.5. numbers 3, 5, and 6, consensus has to be reached and if so the comment will be added to the registered schema.

6.6.2. Informal comment mechanism

The informal comment mechanism provides a way to add interesting information about the schema registered in the DSR.

1. Such a comment could be the information that this schema is implemented in a particular software, is used by a particular project, etc.
2. It can be made by filling in a web form available at the website of the DSR.

3. Besides the comment itself, the name and an operational email address of the commentator has to be provided. If the comment includes the name of an organisation, the connection of the commentator with this organisation must be given as well.
4. The comments will be reviewed and possibly edited by the schema registry operator and, if appropriate, added to the registered schema.
5. Before such an addition the schema registry operator has to send back the comment to the email address specified in the web form to ask for an authorisation for publishing. Such an authorisation can be done via an authorisation URL.

6.7. Specification of the update process and version control

The same process as specified in [SchemProc] and described in 4.4. numbers 5 a) to e) is followed by the DSR.

7. Conclusion

Chapter 6 of this document specifies the policy for the operation of the DSR. It is very likely that through operation experience details of this policy may change. The current specification is needed as input for the software and schema design of the DSR. Future versions of chapter 6 which so to speak represents a contract between policy board, the registry operator and the registry users (as schema posters or schema retriever) will be published on the website of the DSR.

Since the schema WG almost reached Internet consensus, the policy specified here was modelled as close to [SchemaProc] as was thought reasonable for a first pilot phase.

This document clarifies the process for including new schema. It does not handle the situation at the beginning of the project, in which a reasonable amount of existing schema has to be inputted to the registry. What kind of schema already exists is described in chapter 4 of [RegIntro] which can be seen as a proposal for the set of schema definitions which should be included in the first phase.

[RegBusiness] will deal about possible business plans for permanently running the DSR. The specifications made there might also have influence on the policy.

8. Pointer to references

[RegBib] Gietz, Peter, "Bibliography for the Directory Schema Registry Project", Version 1, Deliverable B of the TERENA project Directory Schema Registry, January 2003

[RegBib] contains all references of the project documents.

A. Email from the IETF Application Area Director about the closure of the IETF schema WG

The following email describes the reasons for the closure of the IETF schema WG, stating that the documents themselves have reached a basically mature status and the reason for not processing to RFC status is rather the fact that the listing service has not been established (although initial plans were to establish it January 1998).

```
To: ietf-schema-reg@xxxxxxxx
Subject: Closure of the SCHEMA wg
From: Patrik Fältström <paf@xxxxxxxx>
Date: Thu, 27 Jan 2000 20:15:37 +0100
Cc: IESG <iesg@xxxxxxxx>, IAB <iab@xxxxxxxx>
List-archive: <http://www.imc.org/ietf-schema-reg/mail-archive/>
List-id: <ietf-schema-reg.imc.org>
List-unsubscribe: <mailto:ietf-schema-reg-
request@imc.org?body=unsubscribe>
Sender: owner-ietf-schema-reg@xxxxxxxx
```

To the SCHEMA wg

Because of the unfortunate problems of actually creating the listing service specified in the documents of this working group, I don't see that the working group will succeed with the milestones given in the charter. The documents are from my point of view basically done apart from the text which is supposed to describe the actual listing service itself, and that text can not be created until the listing service exists.

I have therefore reported to the IESG on jan 27 2000 that I have the intention to close this working group if nothing happens within the next two weeks (before feb 10 2000).

If the schema registry is created some time in the future, I see no problem restarting the work with these issues where I see the wg has stalled, but not in this wg.

Patrik Faltstrom

Area Director, Applications Area

B. Listing Request Formats according to [SchemaProc]

For convenience the specification of the listing request formats for schema unit listing request and schema pak listing request as specified in section 2.8 of [SchemaProc] are reproduced here:

B.1. Schema Unit Listing Request Format

```
To: ietf-schema-review@pilot.schema.dir.reg.int
Subject: schema unit listing request
MIME-Version: 1.0
Message-Id: <ids1@wherever.com>
Content-Type: multipart/related; start=3@foo.com; boundary="boundary"
```

```
Content-ID: top@foo.com
```

```
--boundary
Content-Type: text/directory;
           profile="schema-metadata-0";
           charset="utf-8"
Content-Transfer-Encoding: Quoted-Printable
Content-ID: 3@foo.com
```

(metadata elements and values as specified in [SchemaMeta] and embedded in a text/directory MIME component of profile "schema-meta-0")

```
--boundary
Content-Type: text/directory;
           profile="<mimedir-profile-type>";
           charset="utf-8"
Content-Transfer-Encoding: Quoted-Printable
Content-ID: 3@foo.com
```

(a schema specification compliant with a profile of [RFC 2425] corresponding to the value of <mimedir-profile-type>)

```
--boundary
```

B.2. Schema Pak Listing Request Format

```
To: ietf-schema-review@pilot.schema.dir.reg.int
Subject: schema pak listing request
MIME-Version: 1.0
Message-Id: <ids1@wherever.com>
Content-Type: text/directory;
           profile="schema-metadata-0";
           charset="utf-8"
Content-Transfer-Encoding: Quoted-Printable
```

(metadata elements and values as specified in [SchemaMeta] and embedded in a text/directory MIME component of profile "schema-meta-0")

C. Detailed Table contents

- 1. Status of this document..... 1
- 2. Introduction..... 2
- 3. Lessons learned from the registry survey..... 2
 - 3.1. Registries in the IETF 2
 - 3.2. Other registries 3
- 4. Procedures specified by the IETF schema WG..... 3
 - 4.1. Requirements specified 3
 - 4.2. Overview of the listing request process 6
 - 4.3. Process for inclusion of new schema..... 6
 - 4.4. Change control and comments mechanism..... 7
 - 4.5. Responsibilities of the Primary Repository Operator 8
- 5. List of Requirements for a policy of the DSR..... 8
- 6. Policy specifications 9
 - 6.1. Specification of a policy board 9
 - 6.2. Specification of the registry operator 10
 - 6.3. Specification of the syntactical requirements for schema inclusion 10
 - 6.4. Specification of the semantic requirements for schema inclusion..... 11
 - 6.4.1. Semantic requirements for single schema elements..... 11
 - 6.4.2. Semantic requirements for a schema submission..... 11
 - 6.5. Specification of the registration process 12
 - 6.6. Specification of the comment mechanism 13
 - 6.6.1. Formal comment mechanism..... 13
 - 6.6.2. Informal comment mechanism 13
 - 6.7. Specification of the update process and version control 14
- 7. Conclusion..... 14
- 8. Pointer to references 14
- A. Email from the IETF Application Area Director about the closure of the IETF schema WG 15
- B. Listing Request Formats according to [SchemaProc]..... 16
 - B.1. Schema Unit Listing Request Format 16
 - B.2. Schema Pak Listing Request Format 16
- C. Detailed Table contents..... 17