

Network Working Group
Request for Comments: 4516
Obsoletes: 2255
Category: Standards Track

M. Smith, Ed.
Pearl Crescent, LLC
T. Howes
Opware, Inc.
June 2006

Lightweight Directory Access Protocol (LDAP): Uniform Resource Locator

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

This document describes a format for a Lightweight Directory Access Protocol (LDAP) Uniform Resource Locator (URL). An LDAP URL describes an LDAP search operation that is used to retrieve information from an LDAP directory, or, in the context of an LDAP referral or reference, an LDAP URL describes a service where an LDAP operation may be progressed.

Table of Contents

1. Introduction	2
2. URL Definition	2
2.1. Percent-Encoding	4
3. Defaults for Fields of the LDAP URL	5
4. Examples	6
5. Security Considerations	8
6. Normative References	9
7. Informative References	10
8. Acknowledgements	10
Appendix A: Changes Since RFC 2255	11
A.1. Technical Changes	11
A.2. Editorial Changes	11

1. Introduction

LDAP is the Lightweight Directory Access Protocol [RFC4510]. This document specifies the LDAP URL format for version 3 of LDAP and clarifies how LDAP URLs are resolved. This document also defines an extension mechanism for LDAP URLs. This mechanism may be used to provide access to new LDAP extensions.

Note that not all the parameters of the LDAP search operation described in [RFC4511] can be expressed using the format defined in this document. Note also that URLs may be used to represent reference knowledge, including that for non-search operations.

This document is an integral part of the LDAP technical specification [RFC4510], which obsoletes the previously defined LDAP technical specification, RFC 3377, in its entirety.

This document replaces RFC 2255. See Appendix A for a list of changes relative to RFC 2255.

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119].

2. URL Definition

An LDAP URL begins with the protocol prefix "ldap" and is defined by the following grammar, following the ABNF notation defined in [RFC4234].

```
ldapurl      = scheme COLON SLASH SLASH [host [COLON port]]
               [SLASH dn [QUESTION [attributes]
               [QUESTION [scope] [QUESTION [filter]
               [QUESTION extensions]]]]]
               ; <host> and <port> are defined
               ;   in Sections 3.2.2 and 3.2.3
               ;   of [RFC3986].
               ; <filter> is from Section 3 of
               ;   [RFC4515], subject to the
               ;   provisions of the
               ;   "Percent-Encoding" section
               ;   below.

scheme       = "ldap"
```

```

dn          = distinguishedName ; From Section 3 of [RFC4514],
              ; subject to the provisions of
              ; the "Percent-Encoding"
              ; section below.

attributes  = attrdesc *(COMMA attrdesc)
attrdesc    = selector *(COMMA selector)
selector     = attributeSelector ; From Section 4.5.1 of
              ; [RFC4511], subject to the
              ; provisions of the
              ; "Percent-Encoding" section
              ; below.

scope       = "base" / "one" / "sub"
extensions  = extension *(COMMA extension)
extension    = [EXCLAMATION] extype [EQUALS exvalue]
extype       = oid                ; From section 1.4 of [RFC4512].

exvalue     = LDAPString          ; From section 4.1.2 of
              ; [RFC4511], subject to the
              ; provisions of the
              ; "Percent-Encoding" section
              ; below.

EXCLAMATION = %x21                ; exclamation mark ("!")
SLASH       = %x2F                ; forward slash ("/")
COLON       = %x3A                ; colon (":")
QUESTION    = %x3F                ; question mark ("?")

```

The "ldap" prefix indicates an entry or entries accessible from the LDAP server running on the given hostname at the given portnumber. Note that the <host> may contain literal IPv6 addresses as specified in Section 3.2.2 of [RFC3986].

The <dn> is an LDAP Distinguished Name using the string format described in [RFC4514]. It identifies the base object of the LDAP search or the target of a non-search operation.

The <attributes> construct is used to indicate which attributes should be returned from the entry or entries.

The <scope> construct is used to specify the scope of the search to perform in the given LDAP server. The allowable scopes are "base" for a base object search, "one" for a one-level search, or "sub" for a subtree search.

The <filter> is used to specify the search filter to apply to entries within the specified scope during the search. It has the format specified in [RFC4515].

The <extensions> construct provides the LDAP URL with an extensibility mechanism, allowing the capabilities of the URL to be extended in the future. Extensions are a simple comma-separated list of type=value pairs, where the =value portion MAY be omitted for options not requiring it. Each type=value pair is a separate extension. These LDAP URL extensions are not necessarily related to any of the LDAP extension mechanisms. Extensions may be supported or unsupported by the client resolving the URL. An extension prefixed with a '!' character (ASCII 0x21) is critical. An extension not prefixed with a '!' character is non-critical.

If an LDAP URL extension is implemented (that is, if the implementation understands it and is able to use it), the implementation MUST make use of it. If an extension is not implemented and is marked critical, the implementation MUST NOT process the URL. If an extension is not implemented and is not marked critical, the implementation MUST ignore the extension.

The extension type (<extype>) MAY be specified using the numeric OID <numericoid> form (e.g., 1.2.3.4) or the descriptor <descr> form (e.g., myLDAPURLExtension). Use of the <descr> form SHOULD be restricted to registered object identifier descriptive names. See [RFC4520] for registration details and usage guidelines for descriptive names.

No LDAP URL extensions are defined in this document. Other documents or a future version of this document MAY define one or more extensions.

2.1. Percent-Encoding

A generated LDAP URL MUST consist only of the restricted set of characters included in one of the following three productions defined in [RFC3986]:

- <reserved>
- <unreserved>
- <pct-encoded>

Implementations SHOULD accept other valid UTF-8 strings [RFC3629] as input. An octet MUST be encoded using the percent-encoding mechanism described in section 2.1 of [RFC3986] in any of these situations:

The octet is not in the reserved set defined in section 2.2 of [RFC3986] or in the unreserved set defined in section 2.3 of [RFC3986].

It is the single Reserved character '?' and occurs inside a <dn>, <filter>, or other element of an LDAP URL.

It is a comma character ',' that occurs inside an <exvalue>.

Note that before the percent-encoding mechanism is applied, the extensions component of the LDAP URL may contain one or more null (zero) bytes. No other component may.

3. Defaults for Fields of the LDAP URL

Some fields of the LDAP URL are optional, as described above. In the absence of any other specification, the following general defaults SHOULD be used when a field is absent. Note that other documents MAY specify different defaulting rules; for example, section 4.1.10 of [RFC4511] specifies a different rule for determining the correct DN to use when it is absent in an LDAP URL that is returned as a referral.

<host>

If no <host> is given, the client must have some a priori knowledge of an appropriate LDAP server to contact.

<port>

The default LDAP port is TCP port 389.

<dn>

If no <dn> is given, the default is the zero-length DN, "".

<attributes>

If the <attributes> part is omitted, all user attributes of the entry or entries should be requested (e.g., by setting the attributes field AttributeDescriptionList in the LDAP search request to a NULL list, or by using the special <alluserattrs> selector "*").

<scope>

If <scope> is omitted, a <scope> of "base" is assumed.

<filter>

If <filter> is omitted, a filter of "(objectClass=*)" is assumed.

<extensions>

If <extensions> is omitted, no extensions are assumed.

4. Examples

The following are some example LDAP URLs that use the format defined above. The first example is an LDAP URL referring to the University of Michigan entry, available from an LDAP server of the client's choosing:

```
ldap:///o=University%20of%20Michigan,c=US
```

The next example is an LDAP URL referring to the University of Michigan entry in a particular ldap server:

```
ldap://ldap1.example.net/o=University%20of%20Michigan,c=US
```

Both of these URLs correspond to a base object search of the "o=University of Michigan,c=US" entry using a filter of "(objectclass=*)", requesting all attributes.

The next example is an LDAP URL referring to only the postalAddress attribute of the University of Michigan entry:

```
ldap://ldap1.example.net/o=University%20of%20Michigan,  
c=US?postalAddress
```

The corresponding LDAP search operation is the same as in the previous example, except that only the postalAddress attribute is requested.

The next example is an LDAP URL referring to the set of entries found by querying the given LDAP server on port 6666 and doing a subtree search of the University of Michigan for any entry with a common name of "Babs Jensen", retrieving all attributes:

```
ldap://ldap1.example.net:6666/o=University%20of%20Michigan,  
c=US??sub?(cn=Babs%20Jensen)
```

The next example is an LDAP URL referring to all children of the c=GB entry:

```
LDAP://ldap1.example.com/c=GB?objectClass?ONE
```

The objectClass attribute is requested to be returned along with the entries, and the default filter of "(objectclass=*)" is used.

The next example is an LDAP URL to retrieve the mail attribute for the LDAP entry named "o=Question?,c=US", illustrating the use of the percent-encoding mechanism on the reserved character '?':

```
ldap://ldap2.example.com/o=Question%3f,c=US?mail
```

The next example (which is broken into two lines for readability) illustrates the interaction between the LDAP string representation of the filters-quoting mechanism and the URL-quoting mechanisms.

```
ldap://ldap3.example.com/o=Babsco,c=US
    ???(four-octet=%5c00%5c00%5c00%5c04)
```

The filter in this example uses the LDAP escaping mechanism of \ to encode three zero or null bytes in the value. In LDAP, the filter would be written as (four-octet=\00\00\00\04). Because the \ character must be escaped in a URL, the \s are percent-encoded as %5c (or %5C) in the URL encoding.

The next example illustrates the interaction between the LDAP string representation of the DNS-quoting mechanism and URL-quoting mechanisms.

```
ldap://ldap.example.com/o=An%20Example%5C2C%20Inc.,c=US
```

The DN encoded in the above URL is:

```
o=An Example\2C Inc.,c=US
```

That is, the left-most RDN value is:

```
An Example, Inc.
```

The following three URLs are equivalent, assuming that the defaulting rules specified in Section 3 of this document are used:

```
ldap://ldap.example.net
ldap://ldap.example.net/
ldap://ldap.example.net/?
```

These three URLs point to the root DSE on the ldap.example.net server.

The final two examples show use of a hypothetical, experimental bind name extension (the value associated with the extension is an LDAP DN).

```
ldap:///??sub??e-bindname=cn=Manager%2cdc=example%2cdc=com
ldap:///??sub??!e-bindname=cn=Manager%2cdc=example%2cdc=com
```

The two URLs are the same, except that the second one marks the e-bindname extension as critical. Notice the use of the percent-encoding mechanism to encode the commas within the distinguished name value in the e-bindname extension.

5. Security Considerations

The general URL security considerations discussed in [RFC3986] are relevant for LDAP URLs.

The use of security mechanisms when processing LDAP URLs requires particular care, since clients may encounter many different servers via URLs, and since URLs are likely to be processed automatically, without user intervention. A client SHOULD have a user-configurable policy that controls which servers the client will establish LDAP sessions with and with which security mechanisms, and SHOULD NOT establish LDAP sessions that are inconsistent with this policy. If a client chooses to reuse an existing LDAP session when resolving one or more LDAP URLs, it MUST ensure that the session is compatible with the URL and that no security policies are violated.

Sending authentication information, no matter the mechanism, may violate a user's privacy requirements. In the absence of specific policy permitting authentication information to be sent to a server, a client should use an anonymous LDAP session. (Note that clients conforming to previous LDAP URL specifications, where all LDAP sessions are anonymous and unprotected, are consistent with this specification; they simply have the default security policy.) Simply opening a transport connection to another server may violate some users' privacy requirements, so clients should provide the user with a way to control URL processing.

Some authentication methods, in particular, reusable passwords sent to the server, may reveal easily-abused information to the remote server or to eavesdroppers in transit and should not be used in URL processing unless they are explicitly permitted by policy. Confirmation by the human user of the use of authentication information is appropriate in many circumstances. Use of strong authentication methods that do not reveal sensitive information is much preferred. If the URL represents a referral for an update operation, strong authentication methods SHOULD be used. Please refer to the Security Considerations section of [RFC4513] for more information.

The LDAP URL format allows the specification of an arbitrary LDAP search operation to be performed when evaluating the LDAP URL. Following an LDAP URL may cause unexpected results, for example, the retrieval of large amounts of data or the initiation of a long-lived

search. The security implications of resolving an LDAP URL are the same as those of resolving an LDAP search query.

6. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3629] Yergeau, F., "UTF-8, a transformation format of ISO 10646", STD 63, RFC 3629, November 2003.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, January 2005.
- [RFC4234] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", RFC 4234, October 2005.
- [RFC4510] Zeilenga, K., Ed., "Lightweight Directory Access Protocol (LDAP): Technical Specification Road Map", RFC 4510, June 2006.
- [RFC4511] Sermersheim, J., Ed., "Lightweight Directory Access Protocol (LDAP): The Protocol", RFC 4511, June 2006.
- [RFC4512] Zeilenga, K., "Lightweight Directory Access Protocol (LDAP): Directory Information Models", RFC 4512, June 2006.
- [RFC4513] Harrison, R., Ed., "Lightweight Directory Access Protocol (LDAP): Authentication Methods and Security Mechanisms", RFC 4513, June 2006.
- [RFC4514] Zeilenga, K., Ed., "Lightweight Directory Access Protocol (LDAP): String Representation of Distinguished Names", RFC 4514, June 2006.
- [RFC4515] Smith, M. Ed. and T. Howes, "Lightweight Directory Access Protocol (LDAP): String Representation of Search Filters", RFC 4515, June 2006.

7. Informative References

- [RFC2396] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifiers (URI): Generic Syntax", RFC 2396, August 1998.
- [RFC4520] Zeilenga, K., "Internet Assigned Numbers Authority (IANA) Considerations for the Lightweight Directory Access Protocol (LDAP)", BCP 64, RFC 4520, June 2006.

8. Acknowledgements

The LDAP URL format was originally defined at the University of Michigan. This material is based upon work supported by the National Science Foundation under Grant No. NCR-9416667. The support of both the University of Michigan and the National Science Foundation is gratefully acknowledged.

This document obsoletes RFC 2255 by Tim Howes and Mark Smith. Changes included in this revised specification are based upon discussions among the authors, discussions within the LDAP (v3) Revision Working Group (ldapbis), and discussions within other IETF Working Groups. The contributions of individuals in these working groups is gratefully acknowledged. Several people in particular have made valuable comments on this document: RL "Bob" Morgan, Mark Wahl, Kurt Zeilenga, Jim Sermersheim, and Hallvard Furuseth deserve special thanks for their contributions.

Appendix A: Changes Since RFC 2255

A.1. Technical Changes

The following technical changes were made to the contents of the "URL Definition" section:

Revised all of the ABNF to use common productions from [RFC4512].

Replaced references to [RFC2396] with a reference to [RFC3986] (this allows literal IPv6 addresses to be used inside the <host> portion of the URL, and a note was added to remind the reader of this enhancement). Referencing [RFC3986] required changes to the ABNF and text so that productions that are no longer defined by [RFC3986] are not used. For example, <hostport> is not defined by [RFC3986] so it has been replaced with host [COLON port]. Note that [RFC3986] includes new definitions for the "Reserved" and "Unreserved" sets of characters, and the net result is that the following two additional characters should be percent-encoded when they appear anywhere in the data used to construct an LDAP URL: "[" and "]" (these two characters were first added to the Reserved set by RFC 2732).

Changed the definition of <attrdesc> to refer to <attributeSelector> from [RFC4511]. This allows the use of "*" in the <attrdesc> part of the URL. It is believed that existing implementations of RFC 2255 already support this.

Avoided use of <prose-val> (bracketed-string) productions in the <dn>, <host>, <attrdesc>, and <exvalue> rules.

Changed the ABNF for <ldapurl> to group the <dn> component with the preceding <SLASH>.

Changed the <extype> rule to be an <oid> from [RFC4512].

Changed the text about extension types so it references [RFC4520]. Reordered rules to more closely follow the order in which the elements appear in the URL.

"Bindname Extension": removed due to lack of known implementations.

A.2. Editorial Changes

Changed document title to include "LDAP:" prefix.

IESG Note: removed note about lack of satisfactory mandatory authentication mechanisms.

"Status of this Memo" section: updated boilerplate to match current I-D guidelines.

"Abstract" section: separated from introductory material.

"Table of Contents" and "Intellectual Property" sections: added.

"Introduction" section: new section; separated from the Abstract. Changed the text indicate that RFC 2255 is replaced by this document (instead of RFC 1959). Added text to indicate that LDAP URLs are used for references and referrals. Fixed typo (replaced the nonsense phrase "to perform to retrieve" with "used to retrieve"). Added a note to let the reader know that not all of the parameters of the LDAP search operation described in [RFC4511] can be expressed using this format.

"URL Definition" section: removed second copy of <ldapurl> grammar and following two paragraphs (editorial error in RFC 2255). Fixed line break within '!' sequence. Reformatted the ABNF to improve readability by aligning comments and adding some blank lines. Replaced "residing in the LDAP server" with "accessible from the LDAP server" in the sentence immediately following the ABNF. Removed the sentence "Individual attrdesc names are as defined for AttributeDescription in [RFC4511]." because [RFC4511]'s <attributeSelector> is now used directly in the ABNF. Reworded last paragraph to clarify which characters must be percent-encoded. Added text to indicate that LDAP URLs are used for references and referrals. Added text that refers to the ABNF from RFC 4234. Clarified and strengthened the requirements with respect to processing of URLs that contain implemented and not implemented extensions (the approach now closely matches that specified in [RFC4511] for LDAP controls).

"Defaults for Fields of the LDAP URL" section: added; formed by moving text about defaults out of the "URL Definition" section. Replaced direct reference to the attribute name "*" with a reference to the special <alluserattrs> selector "*" defined in [RFC4511].

"URL Processing" section: removed.

"Examples" section: Modified examples to use example.com and example.net hostnames. Added missing '?' to the LDAP URL example whose filter contains three null bytes. Removed space after one comma within a DN. Revised the bindname example to use e-bindname. Changed the name of an attribute used in one example from "int" to "four-octet" to avoid potential confusion. Added an example that demonstrates the interaction between DN escaping and URL percent-encoding. Added some examples to show URL equivalence with respect

to the <dn> portion of the URL. Used uppercase in some examples to remind the reader that some tokens are case-insensitive.

"Security Considerations" section: Added a note about connection reuse. Added a note about using strong authentication methods for updates. Added a reference to [RFC4513]. Added note that simply opening a connection may violate some users' privacy requirements. Adopted the working group's revised LDAP terminology specification by replacing the word "connection" with "LDAP session" or "LDAP connection" as appropriate.

"Acknowledgements" section: added statement that this document obsoletes RFC 2255. Added Kurt Zeilenga, Jim Sermersheim, and Hallvard Furuseth.

"Normative References" section: renamed from "References" per new RFC guidelines. Changed from [1] style to [RFC4511] style throughout the document. Added references to RFC 4234 and RFC 3629. Updated all RFC 1738 references to point to the appropriate sections within [RFC3986]. Updated the LDAP references to refer to LDAPBis WG documents. Removed the reference to the LDAP Attribute Syntaxes document and added references to the [RFC4513], [RFC4520], and [RFC4510] documents.

"Informative References" section: added.

Header and "Authors' Addresses" sections: added "editor" next to Mark Smith's name. Updated affiliation and contact information.

Copyright: updated the year.

Throughout the document: surrounded the names of all ABNF productions with "<" and ">" where they are used in descriptive text.

Authors' Addresses

Mark Smith, Editor
Pearl Crescent, LLC
447 Marlpool Dr.
Saline, MI 48176
USA

Phone: +1 734 944-2856
EMail: mcs@pearlcrscent.com

Tim Howes
Opsware, Inc.
599 N. Mathilda Ave.
Sunnyvale, CA 94085
USA

Phone: +1 408 744-7509
EMail: howes@opsware.com

Full Copyright Statement

Copyright (C) The Internet Society (2006).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).

