

Network Working Group  
Internet-Draft  
Updates: [1738](#) (if approved)  
Obsoletes: [2732](#), [2396](#), [1808](#) (if approved)  
Expires: August 16, 2004

T. Berners-Lee  
MIT/LCS  
R. Fielding  
Day Software  
L. Masinter  
Adobe  
February 16, 2004

Uniform Resource Identifier (URI): Generic Syntax  
draft-fielding-uri-rfc2396bis-04

## Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at  
<<http://www.ietf.org/ietf/lid-abstracts.txt>>.

The list of Internet-Draft Shadow Directories can be accessed at  
<<http://www.ietf.org/shadow.html>>.

This Internet-Draft will expire on August 16, 2004.

## Copyright Notice

Copyright (C) The Internet Society (2004). All Rights Reserved.

## Abstract

A Uniform Resource Identifier (URI) is a compact string of characters for identifying an abstract or physical resource. This specification defines the generic URI syntax and a process for resolving URI references that might be in relative form, along with guidelines and security considerations for the use of URIs on the Internet.

The URI syntax defines a grammar that is a superset of all valid URIs, such that an implementation can parse the common components of a URI reference without knowing the scheme-specific requirements of every possible identifier. This specification does not define a

generative grammar for URIs; that task is performed by the individual specifications of each URI scheme.

## Editorial Note

Discussion of this draft and comments to the editors should be sent to the [uri@w3.org](mailto:uri@w3.org) mailing list. An issues list and version history is available at <http://gbiv.com/protocols/uri/rev-2002/issues.html>.

## Table of Contents

|                       |  |                    |
|-----------------------|--|--------------------|
| <a href="#">1.</a>    | Introduction . . . . .                               | <a href="#">4</a>  |
| <a href="#">1.1</a>   | Overview of URIs . . . . .                           | <a href="#">4</a>  |
| <a href="#">1.1.1</a> | Generic Syntax . . . . .                             | <a href="#">5</a>  |
| <a href="#">1.1.2</a> | Examples . . . . .                                   | <a href="#">6</a>  |
| <a href="#">1.1.3</a> | URI, URL, and URN . . . . .                          | <a href="#">6</a>  |
| <a href="#">1.2</a>   | Design Considerations . . . . .                      | <a href="#">6</a>  |
| <a href="#">1.2.1</a> | Transcription . . . . .                              | <a href="#">6</a>  |
| <a href="#">1.2.2</a> | Separating Identification from Interaction . . . . . | <a href="#">7</a>  |
| <a href="#">1.2.3</a> | Hierarchical Identifiers . . . . .                   | <a href="#">9</a>  |
| <a href="#">1.3</a>   | Syntax Notation . . . . .                            | <a href="#">10</a> |
| <a href="#">2.</a>    | Characters . . . . .                                 | <a href="#">11</a> |
| <a href="#">2.1</a>   | Percent Encoding . . . . .                           | <a href="#">11</a> |
| <a href="#">2.2</a>   | Reserved Characters . . . . .                        | <a href="#">12</a> |
| <a href="#">2.3</a>   | Unreserved Characters . . . . .                      | <a href="#">12</a> |
| <a href="#">2.4</a>   | When to Encode or Decode . . . . .                   | <a href="#">13</a> |
| <a href="#">3.</a>    | Syntax Components . . . . .                          | <a href="#">15</a> |
| <a href="#">3.1</a>   | Scheme . . . . .                                     | <a href="#">15</a> |
| <a href="#">3.2</a>   | Authority . . . . .                                  | <a href="#">16</a> |
| <a href="#">3.2.1</a> | User Information . . . . .                           | <a href="#">16</a> |
| <a href="#">3.2.2</a> | Host . . . . .                                       | <a href="#">17</a> |
| <a href="#">3.2.3</a> | Port . . . . .                                       | <a href="#">20</a> |
| <a href="#">3.3</a>   | Path . . . . .                                       | <a href="#">20</a> |
| <a href="#">3.4</a>   | Query . . . . .                                      | <a href="#">21</a> |
| <a href="#">3.5</a>   | Fragment . . . . .                                   | <a href="#">22</a> |
| <a href="#">4.</a>    | Usage . . . . .                                      | <a href="#">24</a> |
| <a href="#">4.1</a>   | URI Reference . . . . .                              | <a href="#">24</a> |
| <a href="#">4.2</a>   | Relative URI . . . . .                               | <a href="#">24</a> |

|                       |  |                    |
|-----------------------|--|--------------------|
| <a href="#">4.3</a>   | Absolute URI . . . . .                           | <a href="#">25</a> |
| <a href="#">4.4</a>   | Same-document Reference . . . . .                | <a href="#">25</a> |
| <a href="#">4.5</a>   | Suffix Reference . . . . .                       | <a href="#">25</a> |
| <a href="#">5.</a>    | Reference Resolution . . . . .                   | <a href="#">27</a> |
| <a href="#">5.1</a>   | Establishing a Base URI . . . . .                | <a href="#">27</a> |
| <a href="#">5.1.1</a> | Base URI within Document Content . . . . .       | <a href="#">27</a> |
| <a href="#">5.1.2</a> | Base URI from the Encapsulating Entity . . . . . | <a href="#">28</a> |
| <a href="#">5.1.3</a> | Base URI from the Retrieval URI . . . . .        | <a href="#">28</a> |
| <a href="#">5.1.4</a> | Default Base URI . . . . .                       | <a href="#">28</a> |

|                       |   |                    |
|-----------------------|---|--------------------|
| <a href="#">5.2</a>   | Relative Resolution . . . . .                               | <a href="#">28</a> |
| <a href="#">5.2.1</a> | Pre-parse the Base URI . . . . .                            | <a href="#">29</a> |
| <a href="#">5.2.2</a> | Transform References . . . . .                              | <a href="#">29</a> |
| <a href="#">5.2.3</a> | Merge Paths . . . . .                                       | <a href="#">30</a> |
| <a href="#">5.2.4</a> | Remove Dot Segments . . . . .                               | <a href="#">30</a> |
| <a href="#">5.3</a>   | Component Recomposition . . . . .                           | <a href="#">32</a> |
| <a href="#">5.4</a>   | Reference Resolution Examples . . . . .                     | <a href="#">33</a> |
| <a href="#">5.4.1</a> | Normal Examples . . . . .                                   | <a href="#">33</a> |
| <a href="#">5.4.2</a> | Abnormal Examples . . . . .                                 | <a href="#">33</a> |
| <a href="#">6.</a>    | Normalization and Comparison . . . . .                      | <a href="#">35</a> |
| <a href="#">6.1</a>   | Equivalence . . . . .                                       | <a href="#">35</a> |
| <a href="#">6.2</a>   | Comparison Ladder . . . . .                                 | <a href="#">36</a> |
| <a href="#">6.2.1</a> | Simple String Comparison . . . . .                          | <a href="#">36</a> |
| <a href="#">6.2.2</a> | Syntax-based Normalization . . . . .                        | <a href="#">37</a> |
| <a href="#">6.2.3</a> | Scheme-based Normalization . . . . .                        | <a href="#">38</a> |
| <a href="#">6.2.4</a> | Protocol-based Normalization . . . . .                      | <a href="#">39</a> |
| <a href="#">6.3</a>   | Canonical Form . . . . .                                    | <a href="#">39</a> |
| <a href="#">7.</a>    | Security Considerations . . . . .                           | <a href="#">41</a> |
| <a href="#">7.1</a>   | Reliability and Consistency . . . . .                       | <a href="#">41</a> |
| <a href="#">7.2</a>   | Malicious Construction . . . . .                            | <a href="#">41</a> |
| <a href="#">7.3</a>   | Back-end Transcoding . . . . .                              | <a href="#">42</a> |
| <a href="#">7.4</a>   | Rare IP Address Formats . . . . .                           | <a href="#">42</a> |
| <a href="#">7.5</a>   | Sensitive Information . . . . .                             | <a href="#">43</a> |
| <a href="#">7.6</a>   | Semantic Attacks . . . . .                                  | <a href="#">43</a> |
| <a href="#">8.</a>    | Acknowledgments . . . . .                                   | <a href="#">45</a> |
|                       | Normative References . . . . .                              | <a href="#">46</a> |
|                       | Informative References . . . . .                            | <a href="#">47</a> |
|                       | Authors' Addresses . . . . .                                | <a href="#">48</a> |
| <a href="#">A.</a>    | Collected ABNF for URI . . . . .                            | <a href="#">50</a> |
| <a href="#">B.</a>    | Parsing a URI Reference with a Regular Expression . . . . . | <a href="#">52</a> |
| <a href="#">C.</a>    | Delimiting a URI in Context . . . . .                       | <a href="#">53</a> |
| <a href="#">D.</a>    | Summary of Non-editorial Changes . . . . .                  | <a href="#">55</a> |

|                     |  |                    |
|---------------------|--|--------------------|
| <a href="#">D.1</a> | Additions . . . . .                                      | <a href="#">55</a> |
| <a href="#">D.2</a> | Modifications from <a href="#">RFC 2396</a> . . . . .    | <a href="#">55</a> |
|                     | Index . . . . .  | <a href="#">58</a> |
|                     | Intellectual Property and Copyright Statements . . . . . | <a href="#">62</a> |

## [1](#). Introduction

A Uniform Resource Identifier (URI) provides a simple and extensible means for identifying a resource. This specification of URI syntax and semantics is derived from concepts introduced by the World Wide Web global information initiative, whose use of such identifiers dates from 1990 and is described in "Universal Resource Identifiers in WWW" [[RFC1630](#)], and is designed to meet the recommendations laid out in "Functional Recommendations for Internet Resource Locators" [[RFC1736](#)] and "Functional Requirements for Uniform Resource Names" [[RFC1737](#)].

This document obsoletes [[RFC2396](#)], which merged "Uniform Resource Locators" [[RFC1738](#)] and "Relative Uniform Resource Locators" [[RFC1808](#)] in order to define a single, generic syntax for all URIs. It excludes those portions of [RFC 1738](#) that defined the specific syntax of individual URI schemes; those portions will be updated as separate documents. The process for registration of new URI schemes is defined separately by [[RFC2717](#)]. Advice for designers of new URI schemes can be found in [[RFC2718](#)].

All significant changes from [RFC 2396](#) are noted in [Appendix D](#).

This specification uses the terms "character" and "character encoding" in accordance with the definitions provided in [[RFC2978](#)].

## [1.1](#) Overview of URIs

URIs are characterized as follows:

### Uniform

Uniformity provides several benefits: it allows different types of resource identifiers to be used in the same context, even when the mechanisms used to access those resources may differ; it allows uniform semantic interpretation of common syntactic conventions across different types of resource identifiers; it allows introduction of new types of resource identifiers without interfering with the way that existing identifiers are used; and, it allows the identifiers to be reused in many different contexts, thus permitting new applications or protocols to leverage a pre-existing, large, and widely-used set of resource identifiers.

### Resource

Anything that can be named or described can be a resource. Familiar examples include an electronic document, an image, a service (e.g., "today's weather report for Los Angeles"), and a

collection of other resources. A resource is not necessarily accessible via the Internet; e.g., human beings, corporations, and bound books in a library can also be resources. Likewise, abstract concepts can be resources, such as the operators and operands of a mathematical equation or the types of a relationship (e.g., "parent" or "employee").

### Identifier

An identifier embodies the information required to distinguish what is being identified from all other things within its scope of identification.

A URI is an identifier that consists of a sequence of characters matching the syntax defined by the syntax rule named "URI" in [Section 3](#). A URI can be used to refer to a resource. This specification does not place any limits on the nature of a resource or the reasons why an application might wish to refer to a resource. URIs have a global

scope and should be interpreted consistently regardless of context, but that interpretation may be defined in relation to the user's context (e.g., "http://localhost/" refers to a resource that is relative to the user's network interface and yet not specific to any one user).

#### [1.1.1](#) Generic Syntax

Each URI begins with a scheme name, as defined in [Section 3.1](#), that refers to a specification for assigning identifiers within that scheme. As such, the URI syntax is a federated and extensible naming system wherein each scheme's specification may further restrict the syntax and semantics of identifiers using that scheme.

This specification defines those elements of the URI syntax that are required of all URI schemes or are common to many URI schemes. It thus defines the syntax and semantics that are needed to implement a scheme-independent parsing mechanism for URI references, such that the scheme-dependent handling of a URI can be postponed until the scheme-dependent semantics are needed. Likewise, protocols and data formats that make use of URI references can refer to this specification as defining the range of syntax allowed for all URIs, including those schemes that have yet to be defined.

A parser of the generic URI syntax is capable of parsing any URI reference into its major components; once the scheme is determined, further scheme-specific parsing can be performed on the components. In other words, the URI generic syntax is a superset of the syntax of all URI schemes.

#### [1.1.2](#) Examples

The following examples illustrate URIs that are in common use.

<ftp://ftp.is.co.za/rfc/rfc1808.txt>

<http://www.ietf.org/rfc/rfc2396.txt>

mailto:John.Doe@example.com

news:comp.infosystems.www.servers.unix

telnet://melvyl.ucop.edu/

### [1.1.3](#) URI, URL, and URN

A URI can be further classified as a locator, a name, or both. The term "Uniform Resource Locator" (URL) refers to the subset of URIs that, in addition to identifying a resource, provide a means of locating the resource by describing its primary access mechanism (e.g., its network "location"). The term "Uniform Resource Name" (URN) has been used historically to refer to both URIs under the "urn" scheme [[RFC2141](#)], which are required to remain globally unique and persistent even when the resource ceases to exist or becomes unavailable, and to any other URI with the properties of a name.

An individual scheme does not need to be classified as being just one of "name" or "locator". Instances of URIs from any given scheme may have the characteristics of names or locators or both, often depending on the persistence and care in the assignment of identifiers by the naming authority, rather than any quality of the scheme. Future specifications and related documentation should use the general term "URI", rather than the more restrictive terms URL and URN [[RFC3305](#)].

## [1.2](#) Design Considerations

### [1.2.1](#) Transcription

The URI syntax has been designed with global transcription as one of its main considerations. A URI is a sequence of characters from a very limited set: the letters of the basic Latin alphabet, digits, and a few special characters. A URI may be represented in a variety of ways: e.g., ink on paper, pixels on a screen, or a sequence of integers from a coded character set. The interpretation of a URI depends only on the characters used and not how those characters are represented in a network protocol.

The goal of transcription can be described by a simple scenario. Imagine two colleagues, Sam and Kim, sitting in a pub at an international conference and exchanging research ideas. Sam asks Kim for a location to get more information, so Kim writes the URI for the

research site on a napkin. Upon returning home, Sam takes out the napkin and types the URI into a computer, which then retrieves the information to which Kim referred.

There are several design considerations revealed by the scenario:

- o A URI is a sequence of characters that is not always represented as a sequence of octets.
- o A URI might be transcribed from a non-network source, and thus should consist of characters that are most likely to be able to be entered into a computer, within the constraints imposed by keyboards (and related input devices) across languages and locales.
- o A URI often needs to be remembered by people, and it is easier for people to remember a URI when it consists of meaningful or familiar components.

These design considerations are not always in alignment. For example, it is often the case that the most meaningful name for a URI component would require characters that cannot be typed into some systems. The ability to transcribe a resource identifier from one medium to another has been considered more important than having a URI consist of the most meaningful of components.

In local or regional contexts and with improving technology, users might benefit from being able to use a wider range of characters; such use is not defined in this specification. Percent-encoded octets ([Section 2.1](#)) may be used within a URI to represent characters outside the range of the US-ASCII coded character set if such representation is defined by the scheme or by the protocol element in which the URI is referenced; such a definition will specify the character encoding scheme used to map those characters to octets prior to being percent-encoded for the URI.

### [1.2.2](#) Separating Identification from Interaction

A common misunderstanding of URIs is that they are only used to refer to accessible resources. In fact, the URI alone only provides identification; access to the resource is neither guaranteed nor implied by the presence of a URI. Instead, an operation (if any) associated with a URI reference is defined by the protocol element, data format attribute, or natural language text in which it appears.



Given a URI, a system may attempt to perform a variety of operations on the resource, as might be characterized by such words as "access", "update", "replace", or "find attributes". Such operations are defined by the protocols that make use of URIs, not by this specification. However, we do use a few general terms for describing common operations on URIs. URI "resolution" is the process of determining an access mechanism and the appropriate parameters necessary to dereference a URI; such resolution may require several iterations. To use that access mechanism to perform an action on the URI's resource is to "dereference" the URI.

When URIs are used within information systems to identify sources of information, the most common form of URI dereference is "retrieval": making use of a URI in order to retrieve a representation of its associated resource. A "representation" is a sequence of octets, along with representation metadata describing those octets, that constitutes a record of the state of the resource at the time that the representation is generated. Retrieval is achieved by a process that might include using the URI as a cache key to check for a locally cached representation, resolution of the URI to determine an appropriate access mechanism (if any), and dereference of the URI for the sake of applying a retrieval operation. Depending on the protocols used to perform the retrieval, additional information might be supplied about the resource (resource metadata) and its relation to other resources.

URI references in information systems are designed to be late-binding: the result of an access is generally determined at the time it is accessed and may vary over time or due to other aspects of the interaction. When an author creates a reference to such a resource, they do so with the intention that the reference be used in the future; what is being identified is not some specific result that was obtained in the past, but rather some characteristic that is expected to be true for future results. In such cases, the resource referred to by the URI is actually a sameness of characteristics as observed over time, perhaps elucidated by additional comments or assertions made by the resource provider.

Although many URI schemes are named after protocols, this does not imply that use of such a URI will result in access to the resource via the named protocol. URIs are often used simply for the sake of identification. Even when a URI is used to retrieve a representation of a resource, that access might be through gateways, proxies, caches, and name resolution services that are independent of the protocol associated with the scheme name, and the resolution of some URIs may require the use of more than one protocol (e.g., both DNS and HTTP are typically used to access an "http" URI's origin server

when a representation isn't found in a local cache).

### [1.2.3](#) Hierarchical Identifiers

The URI syntax is organized hierarchically, with components listed in order of decreasing significance from left to right. For some URI schemes, the visible hierarchy is limited to the scheme itself: everything after the scheme component delimiter (":") is considered opaque to URI processing. Other URI schemes make the hierarchy explicit and visible to generic parsing algorithms.

The generic syntax uses the slash ("/"), question mark ("?"), and number sign("#") characters for the purpose of delimiting components that are significant to the generic parser's hierarchical interpretation of an identifier. In addition to aiding the readability of such identifiers through the consistent use of familiar syntax, this uniform representation of hierarchy across naming schemes allows scheme-independent references to be made relative to that hierarchy.

It is often the case that a group or "tree" of documents has been constructed to serve a common purpose, wherein the vast majority of URIs in these documents point to resources within the tree rather than outside of it. Similarly, documents located at a particular site are much more likely to refer to other resources at that site than to resources at remote sites. Relative referencing of URIs allows document trees to be partially independent of their location and access scheme. For instance, it is possible for a single set of hypertext documents to be simultaneously accessible and traversable via each of the "file", "http", and "ftp" schemes if the documents refer to each other using relative references. Furthermore, such document trees can be moved, as a whole, without changing any of the relative references.

A relative URI reference ([Section 4.2](#)) refers to a resource by describing the difference within a hierarchical name space between the reference context and the target URI. The reference resolution algorithm, presented in [Section 5](#), defines how such a reference is transformed to the target URI. Since relative references can only be used within the context of a hierarchical URI, designers of new URI schemes should use a syntax consistent with the generic syntax's hierarchical components unless there are compelling reasons to forbid

relative referencing within that scheme.

All URIs are parsed by generic syntax parsers when used. A URI scheme that wishes to remain opaque to hierarchical processing must disallow the use of slash and question mark characters. However, since a non-relative URI reference is only modified by the generic parser if it contains complete path segments of "." or ".." (see [Section 3.3](#)),

URIs may safely use "/" for other purposes if they do not allow dot-segments.

### [1.3](#) Syntax Notation

This specification uses the Augmented Backus-Naur Form (ABNF) notation of [\[RFC2234\]](#), including the following core ABNF syntax rules defined by that specification: ALPHA (letters), CR (carriage return), CTL (control characters), DIGIT (decimal digits), DQUOTE (double quote), HEXDIG (hexadecimal digits), LF (line feed), and SP (space). The complete URI syntax is collected in [Appendix A](#).

## [2.](#) Characters

Although ABNF notation defines its terminal values to be non-negative integers (codepoints) based on the US-ASCII coded character set [[ASCII](#)], we must invert that relation in order to understand the URI syntax, since URIs are defined as strings of characters independent of any particular encoding. Therefore, the integer values must be mapped back to their corresponding characters via US-ASCII in order to complete the syntax rules.

This specification does not mandate the use of any particular character encoding scheme for mapping between URI characters and the octets used to store or transmit those characters. When a URI appears in a protocol element, the character encoding is defined by that protocol; absent such a definition, a URI is assumed to use the same character encoding as the surrounding text.

A URI is composed from a limited set of characters consisting of digits, letters, and a few graphic symbols. A reserved ([Section 2.2](#)) subset of those characters may be used to delimit syntax components within a URI, while the remaining characters, including both the unreserved ([Section 2.3](#)) set and those reserved characters not acting as delimiters, define each component's data.

### [2.1](#) Percent Encoding

A percent-encoding mechanism is used to represent a data octet in a component when that octet's corresponding character is outside the allowed set or is being used as a delimiter of, or within, the component. A percent-encoded octet is encoded as a character triplet, consisting of the percent character "%" followed by the two hexadecimal digits representing that octet's numeric value. For example, "%20" is the percent-encoding for the binary octet "00100000" (ABNF: %x20), which in US-ASCII corresponds to the space character (SP).

pct-encoded = "%" HEXDIG HEXDIG

The uppercase hexadecimal digits 'A' through 'F' are equivalent to the lowercase digits 'a' through 'f', respectively. Two URIs that differ only in the case of hexadecimal digits used in percent-encoded octets are equivalent. For consistency, URI producers and normalizers should use uppercase hexadecimal digits for all percent-encodings.

## [2.2](#) Reserved Characters

URIs include components and sub-components that are delimited by characters in the "reserved" set. These characters are called "reserved" because they may (or may not) be defined as delimiters by the generic syntax, by each scheme-specific syntax, or by the implementation-specific syntax of a URI's dereferencing algorithm. If data for a URI component would conflict with a reserved character's purpose as a delimiter, then the conflicting data must be percent-encoded before forming the URI.

reserved = gen-delims / sub-delims

gen-delims = ":" / "/" / "?" / "#" / "[" / "]" / "@"

sub-delims = "!" / "\$" / "&" / "'" / "(" / ")"  
/ "\*" / "+" / "," / ";" / "="

A subset of the reserved characters (gen-delims) are used as

delimiters of the generic URI components described in [Section 3](#). A component's ABNF syntax rule will not use the reserved or gen-delims rule names directly; instead, each syntax rule lists those reserved characters that are allowed within that component (i.e., not delimiting it). The allowed reserved characters, including those in the sub-delims set and any of the gen-delims that are not a delimiter of that component, are reserved for use as sub-component delimiters within the component. Only the most common sub-components are defined by this specification; other sub-components may be defined by a URI scheme's specification, or by the implementation-specific syntax of a URI's dereferencing algorithm, provided that such sub-components are delimited by characters in that component's reserved set. If no such delimiting role has been assigned, then a reserved character appearing in a component represents the data octet corresponding to its encoding in US-ASCII.

URIs that differ in the replacement of a reserved character with its corresponding percent-encoded octet are not equivalent. Percent-encoding a reserved character, or decoding a percent-encoded octet that corresponds to a reserved character, will change how the URI is interpreted by most applications.

### [2.3](#) Unreserved Characters

Characters that are allowed in a URI but do not have a reserved purpose are called unreserved. These include uppercase and lowercase letters, decimal digits, hyphen, period, underscore, and tilde.

unreserved = ALPHA / DIGIT / "-" / "." / "\_" / "~"

URIs that differ in the replacement of an unreserved character with its corresponding percent-encoded octet are equivalent: they identify the same resource. However, percent-encoded unreserved characters may change the result of some URI comparisons ([Section 6](#)), potentially leading to incorrect or inefficient behavior. For consistency, percent-encoded octets in the ranges of ALPHA (%41-%5A and %61-%7A), DIGIT (%30-%39), hyphen (%2D), period (%2E), underscore (%5F), or tilde (%7E) should not be created by URI producers and, when found in a URI, should be decoded to their corresponding unreserved character by URI normalizers.

### [2.4](#) When to Encode or Decode

Under normal circumstances, the only time that octets within a URI are percent-encoded is during the process of producing the URI from its component parts. It is during that process that an implementation determines which of the reserved characters are to be used as sub-component delimiters and which can be safely used as data. Once produced, a URI is always in its percent-encoded form.

When a URI is dereferenced, the components and sub-components significant to the scheme-specific dereferencing process (if any) must be parsed and separated before the percent-encoded octets within those components can be safely decoded, since otherwise the data may be mistaken for component delimiters. The only exception is for percent-encoded octets corresponding to characters in the unreserved set, which can be decoded at any time. For example, the octet corresponding to the tilde ("~") character is often encoded as "%7E" by older URI processing software; the "%7E" can be replaced by "~" without changing its interpretation.

Because the percent ("%") character serves as the indicator for percent-encoded octets, it must be percent-encoded as "%25" in order for that octet to be used as data within a URI. Implementations must not percent-encode or decode the same string more than once, since decoding an already decoded string might lead to misinterpreting a percent data octet as the beginning of a percent-encoding, or vice versa in the case of percent-encoding an already percent-encoded string.

URI characters serve as an external interface for identification between systems. A system that internally provides identifiers in the form of a different character encoding, such as EBCDIC, will generally perform character translation of textual identifiers to UTF-8 [[RFC3629](#)] (or some other superset of the US-ASCII character encoding) at an internal interface, since that results in more meaningful identifiers than simply percent-encoding the original octets. When interpreting an incoming URI on such an interface,

percent-encoded octets must be decoded before the reverse transcoding can be applied.

In some cases, the interface between a URI component and the identifying data it has been crafted to represent is much less direct

than a character encoding translation. For example, portions of a URI might reflect a query on non-ASCII data, numeric coordinates on a map, etc. Likewise, a URI scheme may define components with additional encoding requirements, such as base64, that are applied prior to forming the component and producing the URI.

When a URI scheme defines a component that represents textual data consisting of characters from the Unicode (ISO/IEC 10646-1) character set, the data should be encoded first as octets according to the UTF-8 character encoding [[RFC3629](#)], and then only those octets that do not correspond to characters in the unreserved set should be percent-encoded. For example, the character A would be represented as "A", the character LATIN CAPITAL LETTER A WITH GRAVE would be represented as "%C3%80", and the character KATAKANA LETTER A would be represented as "%E3%82%A2".



### 3. Syntax Components

The generic URI syntax consists of a hierarchical sequence of components referred to as the scheme, authority, path, query, and fragment.

URI = scheme ":" ["/" authority] path ["?" query] ["#" fragment]

The scheme and path components are required, though path may be empty (no characters). An ABNF-driven parser will find that the border between authority and path is ambiguous; they are disambiguated by the "first-match-wins" (a.k.a. "greedy") algorithm. In other words, if authority is present then the first segment of the path must be empty.

The following are two example URIs and their component parts:

```
foo://example.com:8042/over/there?name=ferret#nose
  \_/  \-----/ \-----/ \-----/ \_/
   |      |           |           |           |
scheme authority   path       query   fragment
   |
 / \ / ----- | \
urn:example:animal:ferret:nose
```

#### 3.1 Scheme

Each URI begins with a scheme name that refers to a specification for assigning identifiers within that scheme. As such, the URI syntax is a federated and extensible naming system wherein each scheme's specification may further restrict the syntax and semantics of identifiers using that scheme.

Scheme names consist of a sequence of characters beginning with a letter and followed by any combination of letters, digits, plus ("+"), period ("."), or hyphen ("-"). Although scheme is case-insensitive, the canonical form is lowercase and documents that specify schemes must do so using lowercase letters. An implementation should accept uppercase letters as equivalent to lowercase in scheme names (e.g., allow "HTTP" as well as "http"), for the sake of robustness, but should only produce lowercase scheme names, for consistency.

scheme = ALPHA \*( ALPHA / DIGIT / "+" / "-" / "." )

Individual schemes are not specified by this document. The process for registration of new URI schemes is defined separately by

[[RFC2717](#)]. The scheme registry maintains the mapping between scheme names and their specifications. Advice for designers of new URI schemes can be found in [[RFC2718](#)].

When presented with a URI that violates one or more scheme-specific restrictions, the scheme-specific resolution process should flag the reference as an error rather than ignore the unused parts; doing so reduces the number of equivalent URIs and helps detect abuses of the generic syntax that might indicate the URI has been constructed to mislead the user ([Section 7.6](#)).

### [3.2](#) Authority

Many URI schemes include a hierarchical element for a naming authority, such that governance of the name space defined by the remainder of the URI is delegated to that authority (which may, in turn, delegate it further). The generic syntax provides a common means for distinguishing an authority based on a registered name or server address, along with optional port and user information.

The authority component is preceded by a double slash ("//") and is terminated by the next slash ("/"), question mark ("?"), or number sign("#") character, or by the end of the URI.

authority = [ userinfo "@" ] host [ ":" port ]

URI producers and normalizers should omit the "@" delimiter that separates userinfo from host if the userinfo component is empty (zero length) and should omit the ":" delimiter that separates host from port if the port component is empty. Some schemes do not allow the userinfo and/or port sub-components.

#### [3.2.1](#) User Information

The userinfo sub-component may consist of a user name and, optionally, scheme-specific information about how to gain authorization to access the resource. The user information, if present, is followed by a commercial at-sign("@") that delimits it from the host.

userinfo = \*( unreserved / pct-encoded / sub-delims / ":" )

Use of the format "user:password" in the userinfo field is deprecated. Applications should not render as clear text any data after the first colon (":") character found within a userinfo sub-component unless such data is the empty string (indicating no password) or "anonymous". Applications may choose to ignore or reject such data when received as part of a reference, and should reject the

storage of such data in unencrypted form. The passing of authentication information in clear text has proven to be a security risk in almost every case where it has been used.

Applications that render a URI for the sake of user feedback, such as in graphical hypertext browsing, should render userinfo in a way that is distinguished from the rest of a URI, when feasible. Such rendering will assist the user in cases where the userinfo has been misleadingly crafted to look like a trusted domain name ([Section 7.6](#)).

### [3.2.2](#) Host

The host sub-component of authority is identified by an IP literal encapsulated within square brackets, an IPv4 address in dotted-decimal form, or a host name.

host = IP-literal / IPv4address / reg-name

The syntax rule for host is ambiguous because it does not completely distinguish between an IPv4address and a reg-name. Again, the "first-match-wins" algorithm applies: If host matches the rule for IPv4address, then it should be considered an IPv4 address literal and not a reg-name. Although host is case-insensitive, producers and normalizers should use lowercase for host names and hexadecimal addresses for the sake of uniformity, while only using uppercase letters for percent-encodings.

A host identified by an Internet Protocol literal address, version 6 [[RFC3513](#)] or later, is distinguished by enclosing the IP literal within square brackets ("[" and "]"). This is the only place where square bracket characters are allowed in the URI syntax. In anticipation of future, as-yet-undefined IP literal address formats, an optional version flag may be used to indicate such a format explicitly rather than relying on heuristic determination.

IP-literal = "[" ( IPv6address / IPvFuture ) "]"

IPvFuture = "v" HEXDIG "." 1\*( unreserved / sub-delims / ":" )

The version flag does not indicate the IP version; rather, it indicates future versions of the literal format. As such, implementations must not provide the version flag for existing IPv4 and IPv6 literal addresses. If a URI containing an IP-literal that starts with "v" (case-insensitive), indicating that the version flag is present, is dereferenced by an application that does not know the meaning of that version flag, then the application should return an appropriate error for "address mechanism not supported".

A host identified by an IPv6 literal address is represented inside the square brackets without a preceding version flag. The ABNF provided here is a translation of the text definition of an IPv6 literal address provided in [\[RFC3513\]](#). A 128-bit IPv6 address is divided into eight 16-bit pieces. Each piece is represented numerically in case-insensitive hexadecimal, using one to four hexadecimal digits (leading zeroes are permitted). The eight encoded pieces are given most-significant first, separated by colon characters. Optionally, the least-significant two pieces may instead be represented in IPv4 address textual format. A sequence of one or more consecutive zero-valued 16-bit pieces within the address may be elided, omitting all their digits and leaving exactly two consecutive colons in their place to mark the elision.

```
IPv6address =
    / 6( h16 ":" ) ls32
    / "::" 5( h16 ":" ) ls32
    / [ h16 ] "::" 4( h16 ":" ) ls32
    / [ *1( h16 ":" ) h16 ] "::" 3( h16 ":" ) ls32
    / [ *2( h16 ":" ) h16 ] "::" 2( h16 ":" ) ls32
    / [ *3( h16 ":" ) h16 ] "::" h16 ":" ls32
    / [ *4( h16 ":" ) h16 ] "::" ls32
    / [ *5( h16 ":" ) h16 ] "::" h16
    / [ *6( h16 ":" ) h16 ] "::"
```

```
ls32 = ( h16 ":" h16 ) / IPv4address
      ; least-significant 32 bits of address
```

```
h16 = 1*4HEXDIG
```

; 16 bits of address represented in hexadecimal

A host identified by an IPv4 literal address is represented in dotted-decimal notation (a sequence of four decimal numbers in the range 0 to 255, separated by "."), as described in [\[RFC1123\]](#) by reference to [\[RFC0952\]](#). Note that other forms of dotted notation may be interpreted on some platforms, as described in [Section 7.4](#), but only the dotted-decimal form of four octets is allowed by this grammar.

IPv4address = dec-octet "." dec-octet "." dec-octet "." dec-octet

|           |                     |           |
|-----------|---------------------|-----------|
| dec-octet | = DIGIT             | ; 0-9     |
|           | / %x31-39 DIGIT     | ; 10-99   |
|           | / "1" 2DIGIT        | ; 100-199 |
|           | / "2" %x30-34 DIGIT | ; 200-249 |
|           | / "25" %x30-35      | ; 250-255 |

A host identified by a registered name is a string of characters that is intended for lookup within a locally-defined host or service name registry. The most common of such registry mechanisms is the Domain Name System (DNS), as defined by [Section 3 of \[RFC1034\]](#) and [Section 2.1 of \[RFC1123\]](#). A DNS name consists of a sequence of domain labels separated by ".", each domain label starting and ending with an alphanumeric character and possibly also containing "-" characters. The rightmost domain label of a fully qualified domain name in DNS may be followed by a single "." and should be followed by one if it is necessary to distinguish between the complete domain name and some local domain.

reg-name = 0\*255( unreserved / pct-encoded / sub-delims )

If the host component is defined and the registered name is empty (zero length), then the name defaults to "localhost" ([Section 6.2.3](#) discusses how this should be normalized). If "localhost" is not determined by a host name lookup, then it should be interpreted to mean the machine on which the URI is being resolved.

This specification does not mandate a particular registered name

lookup technology and therefore does not restrict the syntax of reg-name beyond that necessary for interoperability. Instead, it delegates the issue of host name syntax conformance to the operating system of each application performing URI resolution, and that operating system decides what it will allow for the purpose of host identification. A URI resolution implementation might use DNS, host tables, yellow pages, NetInfo, WINS, or any other system for lookup of host and service names. However, a globally-scoped naming system, such as DNS fully-qualified domain names, is necessary for URIs that are intended to have global scope. URI producers should use host names that conform to the DNS syntax, even when use of DNS is not immediately apparent.

The reg-name syntax allows percent-encoded octets in order to represent non-ASCII host or service names in a uniform way that is independent of the underlying name resolution technology; such octets must represent characters encoded in the UTF-8 character encoding [[RFC3629](#)] prior to being percent-encoded. When a non-ASCII host name represents an internationalized domain name intended for resolution via DNS, the name must be transformed to the IDNA encoding [[RFC3490](#)] prior to name lookup. URI producers should provide such host names in the IDNA encoding, rather than a percent-encoding, if they wish to maximize interoperability with legacy URI resolvers.

The presence of host within a URI does not imply that the scheme requires access to the given host on the Internet. In many cases, the host syntax is used only for the sake of reusing the existing

registration process created and deployed for DNS, thus obtaining a globally unique name without the cost of deploying another registry. However, such use comes with its own costs: domain name ownership may change over time for reasons not anticipated by the URI producer.

### [3.2.3](#) Port

The port sub-component of authority is designated by an optional port number in decimal following the host and delimited from it by a single colon (":") character.

port = \*DIGIT

A scheme may define a default port. For example, the "http" scheme

defines a default port of "80", corresponding to its reserved TCP port number. The type of port designated by the port number (e.g., TCP, UDP, SCTP, etc.) is defined by the URI scheme. URI producers and normalizers should omit the port component and its ":" delimiter if port is empty or its value would be the same as the scheme's default.

### [3.3](#) Path

The path component contains data, usually organized in hierarchical form, that, along with data in the non-hierarchical query component ([Section 3.4](#)), serves to identify a resource within the scope of the URI's scheme and naming authority (if any). If a URI contains an authority component, then the initial path segment must be empty (i.e., the path must begin with a slash ("/") character or be entirely empty). The path is terminated by the first question mark ("?",) or number sign ("#") character, or by the end of the URI.

```
path          = segment *( "/" segment )
segment       = *pchar
```

```
pchar         = unreserved / pct-encoded / sub-delims / ":" / "@"
```

A path consists of a sequence of path segments separated by a slash ("/") character. A path is always defined for a URI, though the defined path may be empty (zero length). Use of the slash character to indicate hierarchy is only required when a URI will be used as the context for relative references. For example, the URI `<mailto:fred@example.com>` has a path of "fred@example.com", whereas the URI `<foo://info.example.com?fred>` has an empty path.

The path segments "." and ".." are defined for relative reference within the path name hierarchy. They are intended for use at the beginning of a relative path reference ([Section 4.2](#)) for indicating

relative position within the hierarchical tree of names. This is similar to their role within some operating systems' file directory structure to indicate the current directory and parent directory, respectively. However, unlike a file system, these dot-segments are only interpreted within the URI path hierarchy and are removed as part of the resolution process ([Section 5.2](#)).

Aside from dot-segments in hierarchical paths, a path segment is considered opaque by the generic syntax. URI-producing applications often use the reserved characters allowed in a segment for the purpose of delimiting scheme-specific or dereference-handler-specific sub-components. For example, the semicolon (";") and equals ("=") reserved characters are often used for delimiting parameters and parameter values applicable to that segment. The comma (",") reserved character is often used for similar purposes. For example, one URI producer might use a segment like "name;v=1.1" to indicate a reference to version 1.1 of "name", whereas another might use a segment like "name,1.1" to indicate the same. Parameter types may be defined by scheme-specific semantics, but in most cases the syntax of a parameter is specific to the implementation of the URI's dereferencing algorithm.

### [3.4](#) Query

The query component contains non-hierarchical data that, along with data in the path component ([Section 3.3](#)), serves to identify a resource within the scope of the URI's scheme and naming authority (if any). The query component is indicated by the first question mark ("?") character and terminated by a number sign ("#") character or by the end of the URI.

query            = \*( pchar / "/" / "?" )

The characters slash ("/") and question mark ("?") may represent data within the query component, but should not be used as such within a URI that is expected to be the base for relative references ([Section 5.1](#)). Incorrect implementations of reference resolution often fail to distinguish query data from path data when looking for hierarchical separators, thus resulting in non-interoperable results. However, since query components are often used to carry identifying information in the form of "key=value" pairs, and one frequently used value is a reference to another URI, it is sometimes better for usability to avoid percent-encoding those characters.



### [3.5](#) Fragment

The fragment identifier component of a URI allows indirect identification of a secondary resource by reference to a primary resource and additional identifying information. The identified secondary resource may be some portion or subset of the primary resource, some view on representations of the primary resource, or some other resource defined or described by those representations. A fragment identifier component is indicated by the presence of a number sign ("#") character and terminated by the end of the URI.

fragment = \*( pchar / "/" / "?" )

The semantics of a fragment identifier are defined by the set of representations that might result from a retrieval action on the primary resource. The fragment's format and resolution is therefore dependent on the media type [[RFC2046](#)] of a potentially retrieved representation, even though such a retrieval is only performed if the URI is dereferenced. Individual media types may define their own restrictions on, or structure within, the fragment identifier syntax for specifying different types of subsets, views, or external references that are identifiable as secondary resources by that media type. If the primary resource has multiple representations, as is often the case for resources whose representation is selected based on attributes of the retrieval request (a.k.a., content negotiation), then whatever is identified by the fragment should be consistent across all of those representations: each representation should either define the fragment such that it corresponds to the same secondary resource, regardless of how it is represented, or the fragment should be left undefined by the representation (i.e., not found).

As with any URI, use of a fragment identifier component does not imply that a retrieval action will take place. A URI with a fragment identifier may be used to refer to the secondary resource without any implication that the primary resource is accessible or will ever be accessed.

Fragment identifiers have a special role in information systems as the primary form of client-side indirect referencing, allowing an author to specifically identify those aspects of an existing resource that are only indirectly provided by the resource owner. As such, interpretation of the fragment identifier during a retrieval action is performed solely by the user agent; the fragment identifier is not passed to other systems during the process of retrieval. Although this is often perceived to be a loss of information, particularly in regards to accurate redirection of references as content moves over time, it also serves to prevent information providers from denying

reference authors the right to selectively refer to information within a resource.

The characters slash ("/") and question mark ("?") are allowed to represent data within the fragment identifier, but should not be used as such within a URI that is expected to be the base for relative references ([Section 5.1](#)) for the same reasons as described above for query.

## [4. Usage](#)

When applications make reference to a URI, they do not always use the full form of reference defined by the "URI" syntax rule. In order to save space and take advantage of hierarchical locality, many Internet protocol elements and media type formats allow an abbreviation of a URI, while others restrict the syntax to a particular form of URI. We define the most common forms of reference syntax in this specification because they impact and depend upon the design of the generic syntax, requiring a uniform parsing algorithm in order to be interpreted consistently.

### [4.1 URI Reference](#)

URI-reference is used to denote the most common usage of a resource identifier.

$$\text{URI-reference} = \text{URI} / \text{relative-URI}$$

A URI-reference may be relative: if the reference's prefix matches the syntax of a scheme followed by its colon separator, then the reference is a URI rather than a relative-URI.

A URI-reference is typically parsed first into the five URI components, in order to determine what components are present and whether or not the reference is relative, and then each component is parsed for its subparts and their validation. The ABNF of URI-reference, along with the "first-match-wins" disambiguation rule, is sufficient to define a validating parser for the generic syntax. Readers familiar with regular expressions should see [Appendix B](#) for an example of a non-validating URI-reference parser that will take any given string and extract the URI components.

### [4.2 Relative URI](#)

A relative URI reference takes advantage of the hierarchical syntax

([Section 1.2.3](#)) in order to express a reference that is relative to the name space of another hierarchical URI.

relative-URI = ["//"<sup>authority</sup>] path ["?"<sup>query</sup>] ["#"<sup>fragment</sup>]

The URI referred to by a relative reference, also known as the target URI, is obtained by applying the reference resolution algorithm of [Section 5](#).

A relative reference that begins with two slash characters is termed a network-path reference; such references are rarely used. A relative reference that begins with a single slash character is termed an

absolute-path reference. A relative reference that does not begin with a slash character is termed a relative-path reference.

A path segment that contains a colon character (e.g., "this:that") cannot be used as the first segment of a relative-path reference because it would be mistaken for a scheme name. Such a segment must be preceded by a dot-segment (e.g., "./this:that") to make a relative-path reference.

### [4.3](#) Absolute URI

Some protocol elements allow only the absolute form of a URI without a fragment identifier. For example, defining a base URI for later use by relative references calls for an absolute-URI syntax rule that does not allow a fragment.

absolute-URI = scheme ":" ["//"<sup>authority</sup>] path ["?"<sup>query</sup>]

### [4.4](#) Same-document Reference

When a URI reference refers to a URI that is, aside from its fragment component (if any), identical to the base URI ([Section 5.1](#)), that reference is called a "same-document" reference. The most frequent examples of same-document references are relative references that are empty or include only the number sign ("#") separator followed by a fragment identifier.

When a same-document reference is dereferenced for the purpose of a

retrieval action, the target of that reference is defined to be within the same entity (representation, document, or message) as the reference; therefore, a dereference should not result in a new retrieval action.

Normalization of the base and target URIs prior to their comparison, as described in [Section 6.2.2](#) and [Section 6.2.3](#), is allowed but rarely performed in practice. Normalization may increase the set of same-document references, which may be of benefit to some caching applications. As such, reference authors should not assume that a slightly different, though equivalent, reference URI will (or will not) be interpreted as a same-document reference by any given application.

#### [4.5](#) Suffix Reference

The URI syntax is designed for unambiguous reference to resources and extensibility via the URI scheme. However, as URI identification and usage have become commonplace, traditional media (television, radio,

newspapers, billboards, etc.) have increasingly used a suffix of the URI as a reference, consisting of only the authority and path portions of the URI, such as

`www.w3.org/Addressing/`

or simply a DNS registered name on its own. Such references are primarily intended for human interpretation, rather than for machines, with the assumption that context-based heuristics are sufficient to complete the URI (e.g., most host names beginning with "www" are likely to have a URI prefix of "http://"). Although there is no standard set of heuristics for disambiguating a URI suffix, many client implementations allow them to be entered by the user and heuristically resolved.

While this practice of using suffix references is common, it should be avoided whenever possible and never used in situations where long-term references are expected. The heuristics noted above will change over time, particularly when a new URI scheme becomes popular, and are often incorrect when used out of context. Furthermore, they can lead to security issues along the lines of those described in [\[RFC1535\]](#).

Since a URI suffix has the same syntax as a relative path reference, a suffix reference cannot be used in contexts where a relative reference is expected. As a result, suffix references are limited to those places where there is no defined base URI, such as dialog boxes and off-line advertisements.

## [5.](#) Reference Resolution

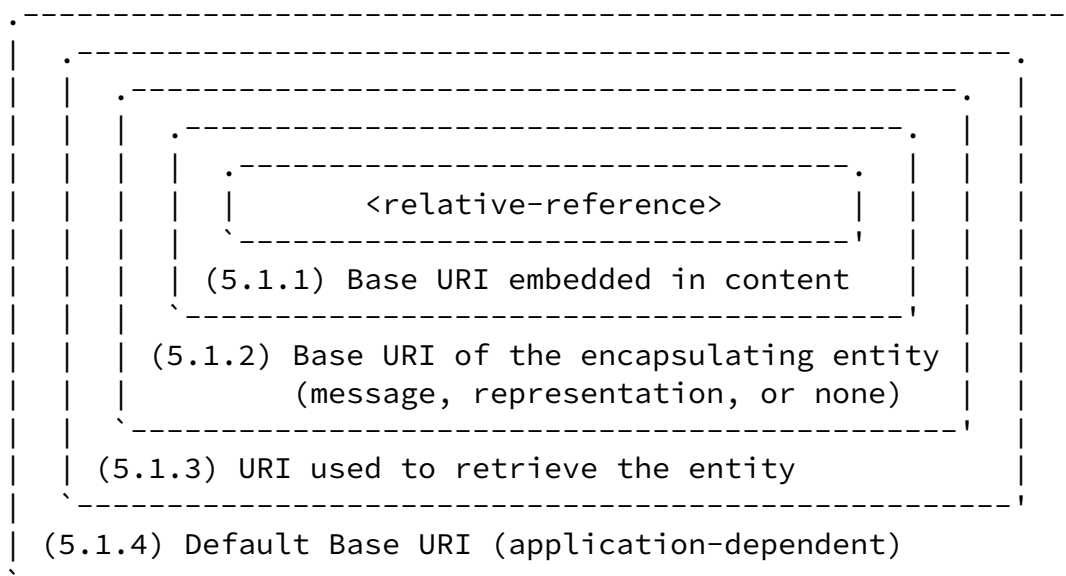
This section defines the process of resolving a URI reference within a context that allows relative references, such that the result is a string matching the "URI" syntax rule of [Section 3](#).

### [5.1](#) Establishing a Base URI

The term "relative" implies that there exists a "base URI" against which the relative reference is applied. Aside from fragment-only references ([Section 4.4](#)), relative references are only usable when a base URI is known. A base URI must be established by the parser prior to parsing URI references that might be relative.

The base URI of a reference can be established in one of four ways,

discussed below in order of precedence. The order of precedence can be thought of in terms of layers, where the innermost defined base URI has the highest precedence. This can be visualized graphically as:



#### [5.1.1](#) Base URI within Document Content

Within certain media types, a base URI for relative references can be embedded within the content itself such that it can be readily obtained by a parser. This can be useful for descriptive documents, such as tables of content, which may be transmitted to others through protocols other than their usual retrieval context (e.g., E-Mail or USENET news).

It is beyond the scope of this specification to specify how, for each

media type, a base URI can be embedded. The appropriate syntax, when available, is described by each media type's specification.

#### [5.1.2](#) Base URI from the Encapsulating Entity

If no base URI is embedded, the base URI is defined by the representation's retrieval context. For a document that is enclosed within another entity, such as a message or archive, the retrieval

context is that entity; thus, the default base URI of a representation is the base URI of the entity in which the representation is encapsulated.

A mechanism for embedding a base URI within MIME container types (e.g., the message and multipart types) is defined by MHTML [RFC2110]. Protocols that do not use the MIME message header syntax, but do allow some form of tagged metadata to be included within messages, may define their own syntax for defining a base URI as part of a message.

### [5.1.3](#) Base URI from the Retrieval URI

If no base URI is embedded and the representation is not encapsulated within some other entity, then, if a URI was used to retrieve the representation, that URI shall be considered the base URI. Note that if the retrieval was the result of a redirected request, the last URI used (i.e., the URI that resulted in the actual retrieval of the representation) is the base URI.

### [5.1.4](#) Default Base URI

If none of the conditions described above apply, then the base URI is defined by the context of the application. Since this definition is necessarily application-dependent, failing to define a base URI using one of the other methods may result in the same content being interpreted differently by different types of application.

A sender of a representation containing relative references is responsible for ensuring that a base URI for those references can be established. Aside from fragment-only references, relative references can only be used reliably in situations where the base URI is well-defined.

## [5.2](#) Relative Resolution

This section describes an algorithm for converting a URI reference that might be relative to a given base URI into the parsed components of the reference's target. The components can then be recomposed, as described in [Section 5.3](#), to form the target URI. This algorithm

provides definitive results that can be used to test the output of



other implementations. Applications may implement relative reference resolution using some other algorithm, provided that the results match what would be given by this algorithm.

#### [5.2.1](#) Pre-parse the Base URI

The base URI (Base) is established according to the procedure of [Section 5.1](#) and parsed into the five main components described in [Section 3](#). Note that only the scheme component is required to be present in a base URI; the other components may be empty or undefined. A component is undefined if its associated delimiter does not appear in the URI reference; the path component is never undefined, though it may be empty.

Normalization of the base URI, as described in [Section 6.2.2](#) and [Section 6.2.3](#), is optional. A URI reference must be transformed to its target URI before it can be normalized.

#### [5.2.2](#) Transform References

For each URI reference (R), the following pseudocode describes an algorithm for transforming R into its target URI (T):

```
-- The URI reference is parsed into the five URI components
--
(R.scheme, R.authority, R.path, R.query, R.fragment) = parse(R);

-- A non-strict parser may ignore a scheme in the reference
-- if it is identical to the base URI's scheme.
--
if ((not strict) and (R.scheme == Base.scheme)) then
    undefine(R.scheme);
endif;

if defined(R.scheme) then
    T.scheme      = R.scheme;
    T.authority   = R.authority;
    T.path        = remove_dot_segments(R.path);
    T.query       = R.query;
else
    if defined(R.authority) then
        T.authority = R.authority;
        T.path      = remove_dot_segments(R.path);
        T.query     = R.query;
    else
        if (R.path == "") then
            T.path = Base.path;
```

```
    if defined(R.query) then
        T.query = R.query;
    else
        T.query = Base.query;
    endif;
else
    if (R.path starts-with "/" ) then
        T.path = remove_dot_segments(R.path);
    else
        T.path = merge(Base.path, R.path);
        T.path = remove_dot_segments(T.path);
    endif;
    T.query = R.query;
endif;
T.authority = Base.authority;
endif;
T.scheme = Base.scheme;
endif;

T.fragment = R.fragment;
```

### [5.2.3 Merge Paths](#)

The pseudocode above refers to a "merge" routine for merging a relative-path reference with the path of the base URI. This is accomplished as follows:

- o If the base URI has a defined authority component and an empty path, then return a string consisting of "/" concatenated with the reference's path; otherwise,
- o Return a string consisting of the reference's path component appended to all but the last segment of the base URI's path (i.e., excluding any characters after the right-most "/" in the base URI path, or excluding the entire base URI path if it does not contain any "/" characters).

### [5.2.4 Remove Dot Segments](#)

The pseudocode also refers to a "remove\_dot\_segments" routine for interpreting and removing the special "." and ".." complete path segments from a referenced path. This is done after the path is extracted from a reference, whether or not the path was relative, in order to remove any invalid or extraneous dot-segments prior to

forming the target URI. Although there are many ways to accomplish this removal process, we describe a simple method using a two string

buffers.

1. The input buffer is initialized with the now-appended path components and the output buffer is initialized to the empty string.
2. Replace any prefix of "./" or "../" at the beginning of the input buffer with "/".
3. While the input buffer is not empty, loop:
  1. If the input buffer begins with a prefix of "/./" or "/.", where "." is a complete path segment, then replace that prefix with "/"; otherwise
  2. If the input buffer begins with a prefix of "../" or "../.", where ".." is a complete path segment, then replace that prefix with "/" and remove the last segment and its preceding "/" (if any) from the output buffer; otherwise
  3. Remove the first segment and its preceding "/" (if any) from the input buffer and append them to the output buffer.
4. Finally, the output buffer is returned as the result of `remove_dot_segments`.

The following illustrates how the above steps are applied for two example merged paths, showing the state of the two buffers after each step.

| STEP | OUTPUT BUFFER | INPUT BUFFER  |
|------|---------------|---------------|
| 1 :  |               | /a/b/c/./../g |
| 3c:  | /a            | /b/c/./../g   |
| 3c:  | /a/b          | /c/./../g     |
| 3c:  | /a/b/c        | /./../g       |
| 3a:  | /a/b/c        | /../g         |
| 3b:  | /a/b          | /g            |
| 3b:  | /a            |               |

3c:    /a/g

| STEP       | OUTPUT BUFFER | INPUT BUFFER          |
|------------|---------------|-----------------------|
| <u>1</u> : |               | mid/content=5/../../6 |
| 3c:        | mid           | /content=5/../../6    |
| 3c:        | mid/content=5 | /../../6              |
| 3b:        | mid           | /6                    |
| 3c:        | mid/6         |                       |

Some applications may find it more efficient to implement the `remove_dot_segments` algorithm using two segment stacks rather than strings.

Note: Some client applications will fail to separate a reference's query component from its path component before merging the base and reference paths. This may result in loss of information if the query component contains the strings `"../"` or `"./"`.

### [5.3](#) Component Recomposition

Parsed URI components can be recomposed to obtain the corresponding URI reference string. Using pseudocode, this would be:

```
result = ""

if defined(scheme) then
    append scheme to result;
    append ":" to result;
endif;

if defined(authority) then
    append "//" to result;
    append authority to result;
endif;

append path to result;

if defined(query) then
    append "?" to result;
    append query to result;
```

```

endif;

if defined(fragment) then
    append "#" to result;
    append fragment to result;
endif;

return result;

```

Note that we are careful to preserve the distinction between a component that is undefined, meaning that its separator was not present in the reference, and a component that is empty, meaning that the separator was present and was immediately followed by the next component separator or the end of the reference.

## [5.4](#) Reference Resolution Examples

Within a representation with a well-defined base URI of

```
http://a/b/c/d;p?q
```

a relative URI reference is transformed to its target URI as follows.

### [5.4.1](#) Normal Examples

|           |   |                        |
|-----------|---|------------------------|
| "g:h"     | = | "g:h"                  |
| "g"       | = | "http://a/b/c/g"       |
| "/g"      | = | "http://a/b/c/g"       |
| "g/"      | = | "http://a/b/c/g/"      |
| "/g"      | = | "http://a/g"           |
| //g"      | = | "http://g"             |
| "?y"      | = | "http://a/b/c/d;p?y"   |
| "g?y"     | = | "http://a/b/c/g?y"     |
| "#s"      | = | "http://a/b/c/d;p?q#s" |
| "g#s"     | = | "http://a/b/c/g#s"     |
| "g?y#s"   | = | "http://a/b/c/g?y#s"   |
| ";x"      | = | "http://a/b/c/;x"      |
| "g;x"     | = | "http://a/b/c/g;x"     |
| "g;x?y#s" | = | "http://a/b/c/g;x?y#s" |
| "         | = | "http://a/b/c/d;p?q"   |

|           |   |                 |
|-----------|---|-----------------|
| "."       | = | "http://a/b/c/" |
| "./"      | = | "http://a/b/c/" |
| ".."      | = | "http://a/b/"   |
| "../"     | = | "http://a/b/"   |
| "../g"    | = | "http://a/b/g"  |
| "../.."   | = | "http://a/"     |
| "../../"  | = | "http://a/"     |
| "../../g" | = | "http://a/g"    |

#### [5.4.2](#) Abnormal Examples

Although the following abnormal examples are unlikely to occur in normal practice, all URI parsers should be capable of resolving them consistently. Each example uses the same base as above.

Parsers must be careful in handling cases where there are more relative path ".." segments than there are hierarchical levels in the base URI's path. Note that the ".." syntax cannot be used to change the authority component of a URI.

|                |   |              |
|----------------|---|--------------|
| "../.../g"     | = | "http://a/g" |
| "../.../.../g" | = | "http://a/g" |

Similarly, parsers must remove the dot-segments "." and ".." when they are complete components of a path, but not when they are only part of a segment.

|        |   |                    |
|--------|---|--------------------|
| "./g"  | = | "http://a/g"       |
| "../g" | = | "http://a/g"       |
| "g."   | = | "http://a/b/c/g."  |
| ".g"   | = | "http://a/b/c/.g"  |
| "g.."  | = | "http://a/b/c/g.." |
| "..g"  | = | "http://a/b/c/..g" |

Less likely are cases where the relative URI reference uses unnecessary or nonsensical forms of the "." and ".." complete path segments.

|            |   |                    |
|------------|---|--------------------|
| "../.../g" | = | "http://a/b/g"     |
| "./g/."    | = | "http://a/b/c/g/"  |
| "g/./h"    | = | "http://a/b/c/g/h" |

```

"g/../h"      = "http://a/b/c/h"
"g;x=1/../y"  = "http://a/b/c/g;x=1/y"
"g;x=1/../y"  = "http://a/b/c/y"

```

Some applications fail to separate the reference's query and/or fragment components from a relative path before merging it with the base path and removing dot-segments. This error is rarely noticed, since typical usage of a fragment never includes the hierarchy ("/") character, and the query component is not normally used within relative references.

```

"g?y/../x"    = "http://a/b/c/g?y/../x"
"g?y/../x"    = "http://a/b/c/g?y/./x"
"g#s/../x"    = "http://a/b/c/g#s/../x"
"g#s/../x"    = "http://a/b/c/g#s/./x"

```

Some parsers allow the scheme name to be present in a relative URI reference if it is the same as the base URI scheme. This is considered to be a loophole in prior specifications of partial URI [[RFC1630](#)]. Its use should be avoided, but is allowed for backward compatibility.

```

"http:g"      = "http:g"          ; for strict parsers
               / "http://a/b/c/g" ; for backward compatibility

```

## [6](#). Normalization and Comparison

One of the most common operations on URIs is simple comparison: determining if two URIs are equivalent without using the URIs to access their respective resource(s). A comparison is performed every time a response cache is accessed, a browser checks its history to color a link, or an XML parser processes tags within a namespace. Extensive normalization prior to comparison of URIs is often used by spiders and indexing engines to prune a search space or reduce duplication of request actions and response storage.

URI comparison is performed in respect to some particular purpose, and software with differing purposes will often be subject to differing design trade-offs in regards to how much effort should be spent in reducing duplicate identifiers. This section describes a variety of methods that may be used to compare URIs, the trade-offs between them, and the types of applications that might use them.

## [6.1](#) Equivalence

Since URIs exist to identify resources, presumably they should be considered equivalent when they identify the same resource. However, such a definition of equivalence is not of much practical use, since there is no way for software to compare two resources without knowledge of the implementation-specific syntax of each URI's dereferencing algorithm. For this reason, determination of equivalence or difference of URIs is based on string comparison, perhaps augmented by reference to additional rules provided by URI scheme definitions. We use the terms "different" and "equivalent" to describe the possible outcomes of such comparisons, but there are many application-dependent versions of equivalence.

Even though it is possible to determine that two URIs are equivalent, it is never possible to be sure that two URIs identify different resources. For example, an owner of two different domain names could decide to serve the same resource from both, resulting in two different URIs. Therefore, comparison methods are designed to minimize false negatives while strictly avoiding false positives.

In testing for equivalence, applications should not directly compare relative URI references; the references should be converted to their target URI forms before comparison. When URIs are being compared for the purpose of selecting (or avoiding) a network action, such as retrieval of a representation, the fragment components (if any) should be excluded from the comparison.

## [6.2](#) Comparison Ladder

A variety of methods are used in practice to test URI equivalence. These methods fall into a range, distinguished by the amount of



processing required and the degree to which the probability of false negatives is reduced. As noted above, false negatives cannot in principle be eliminated. In practice, their probability can be reduced, but this reduction requires more processing and is not cost-effective for all applications.

If this range of comparison practices is considered as a ladder, the following discussion will climb the ladder, starting with those practices that are cheap but have a relatively higher chance of producing false negatives, and proceeding to those that have higher computational cost and lower risk of false negatives.

#### [6.2.1](#) Simple String Comparison

If two URIs, considered as character strings, are identical, then it is safe to conclude that they are equivalent. This type of equivalence test has very low computational cost and is in wide use in a variety of applications, particularly in the domain of parsing.

Testing strings for equivalence requires some basic precautions. This procedure is often referred to as "bit-for-bit" or "byte-for-byte" comparison, which is potentially misleading. Testing of strings for equality is normally based on pairwise comparison of the characters that make up the strings, starting from the first and proceeding until both strings are exhausted and all characters found to be equal, a pair of characters compares unequal, or one of the strings is exhausted before the other.

Such character comparisons require that each pair of characters be put in comparable form. For example, should one URI be stored in a byte array in EBCDIC encoding, and the second be in a Java String object (UTF-16), bit-for-bit comparisons applied naively will produce both false-positive and false-negative errors. It is better to speak of equality on a character-for-character rather than byte-for-byte or bit-for-bit basis. In practical terms, character-by-character comparisons should be done codepoint-by-codepoint after conversion to a common character encoding.

### [6.2.2](#) Syntax-based Normalization

Software may use logic based on the definitions provided by this specification to reduce the probability of false negatives. Such processing is moderately higher in cost than character-for-character string comparison. For example, an application using this approach could reasonably consider the following two URIs equivalent:

```
example://a/b/c/%7Bfoo%7D
eXAMPLE://a/./b/../b/%63/%7bfoo%7d
```

Web user agents, such as browsers, typically apply this type of URI normalization when determining whether a cached response is available. Syntax-based normalization includes such techniques as case normalization, encoding normalization, empty-component normalization, and removal of dot-segments.

#### [6.2.2.1](#) Case Normalization

When a URI scheme uses components of the generic syntax, it will also use the common syntax equivalence rules, namely that the scheme and host are case-insensitive and therefore should be normalized to lowercase. For example, the URI `<HTTP://www.EXAMPLE.com/>` is equivalent to `<http://www.example.com/>`. Applications should not assume anything about the case sensitivity of other URI components, since that is dependent on the implementation used to handle a dereference.

The hexadecimal digits within a percent-encoding triplet (e.g., `"%3a"` versus `"%3A"`) are case-insensitive and therefore should be normalized to use uppercase letters for the digits A-F.

#### [6.2.2.2](#) Encoding Normalization

The percent-encoding mechanism ([Section 2.1](#)) is a frequent source of variance among otherwise identical URIs. In addition to the case-insensitivity issue noted above, some URI producers percent-encode octets that do not require percent-encoding, resulting in URIs that are equivalent to their non-encoded counterparts. Such URIs should be normalized by decoding any percent-encoded octet that corresponds to an unreserved character, as described in [Section 2.3](#).

#### [6.2.2.3](#) Empty-component Normalization

Components of the generic URI syntax are delimited from other components by optional separators. For example, a query component is separated from the path by a question mark ("?") and a port sub-component is separated from host by a colon (":"). A URI in which a delimiter is present and the (sub-)component it delimits is empty is equivalent to the same URI without that delimiter. For example, the following are all equivalent:

```
ftp://example.com/  
ftp://example.com/  
ftp://@example.com/  
ftp://@example.com:/?  
ftp://@example.com:/?#
```

URI producers and normalizers should omit a delimiter if the component it delimits is empty, as exemplified by the first URI above, with one exception: a double-slash delimiter indicating an authority component should not be removed, even when the authority is empty, since doing so can lead to misinterpreting the path.

#### [6.2.2.4](#) Path Segment Normalization

The complete path segments "." and ".." have a special meaning within hierarchical URI schemes. As such, they should not appear in absolute paths; if they are found, they can be removed by applying the `remove_dot_segments` algorithm to the path, as described in [Section 5.2](#).

#### [6.2.3](#) Scheme-based Normalization

The syntax and semantics of URIs vary from scheme to scheme, as described by the defining specification for each scheme. Software may use scheme-specific rules, at further processing cost, to reduce the probability of false negatives. For example, since the "http" scheme makes use of an authority component, has a default port of "80", and defines an empty path to be equivalent to "/", the following four URIs are equivalent:

`http://example.com`  
`http://example.com/`  
<http://example.com:/>  
<http://example.com:80/>

In general, a URI that uses the generic syntax for authority with an empty path should be normalized to a path of `"/`"; likewise, an explicit `":port"`, where the port is empty or the default for the

scheme, is equivalent to one where the port and its `":"` delimiter are elided. In other words, the second of the above URI examples is the normal form for the "http" scheme.

Another case where normalization varies by scheme is in the handling of an empty authority component. For many scheme specifications, an empty authority is considered an error; for others, it is considered equivalent to "localhost". For the sake of uniformity, future scheme specifications should define an empty authority as being equivalent to "localhost", and URI producers and normalizers should use "localhost" instead of an empty authority.

#### [6.2.4](#) Protocol-based Normalization

Web spiders, for which substantial effort to reduce the incidence of false negatives is often cost-effective, are observed to implement even more aggressive techniques in URI comparison. For example, if they observe that a URI such as

`http://example.com/data`

redirects to a URI differing only in the trailing slash

`http://example.com/data/`

they will likely regard the two as equivalent in the future. This kind of technique is only appropriate when equivalence is clearly indicated by both the result of accessing the resources and the common conventions of their scheme's dereference algorithm (in this case, use of redirection by HTTP origin servers to avoid problems with relative references).

### [6.3](#) Canonical Form

It is in the best interests of everyone to avoid false-negatives in comparing URIs and to minimize the amount of software processing for such comparisons. Those who produce and make reference to URIs can reduce the cost of processing and the risk of false negatives by consistently providing them in a form that is reasonably canonical with respect to their scheme. Specifically:

- o Always provide the URI scheme in lowercase characters.
- o Always provide the host, if any, in lowercase characters.
- o Only perform percent-encoding where it is essential.
- o Always use uppercase A-through-F characters when percent-encoding.
- o Prevent `./` and `../` from appearing in non-relative URI paths.
- o Omit delimiters when their associated (sub-)component is empty.
- o For schemes that define an empty authority to be equivalent to "localhost", use "localhost".
- o For schemes that define an empty path to be equivalent to a path of `/`, use `/`.

## [7.](#) Security Considerations

A URI does not in itself pose a security threat. However, since URIs are often used to provide a compact set of instructions for access to network resources, care must be taken to properly interpret the data within a URI, to prevent that data from causing unintended access, and to avoid including data that should not be revealed in plain text.

### [7.1](#) Reliability and Consistency

There is no guarantee that, having once used a given URI to retrieve some information, the same information will be retrievable by that URI in the future. Nor is there any guarantee that the information retrievable via that URI in the future will be observably similar to that retrieved in the past. The URI syntax does not constrain how a given scheme or authority apportions its name space or maintains it over time. Such a guarantee can only be obtained from the person(s) controlling that name space and the resource in question. A specific

URI scheme may define additional semantics, such as name persistence, if those semantics are required of all naming authorities for that scheme.

## [7.2](#) Malicious Construction

It is sometimes possible to construct a URI such that an attempt to perform a seemingly harmless, idempotent operation, such as the retrieval of a representation, will in fact cause a possibly damaging remote operation to occur. The unsafe URI is typically constructed by specifying a port number other than that reserved for the network protocol in question. The client unwittingly contacts a site that is running a different protocol service and data within the URI contains instructions that, when interpreted according to this other protocol, cause an unexpected operation. A frequent example of such abuse has been the use of a protocol-based scheme with a port component of "25", thereby fooling user agent software into sending an unintended or impersonating message via an SMTP server.

Applications should prevent dereference of a URI that specifies a TCP port number within the "well-known port" range (0 - 1023) unless the protocol being used to dereference that URI is compatible with the protocol expected on that well-known port. Although IANA maintains a registry of well-known ports, applications should make such restrictions user-configurable to avoid preventing the deployment of new services.

When a URI contains percent-encoded octets that match the delimiters for a given resolution or dereference protocol (for example, CR and

LF characters for the TELNET protocol), such percent-encoded octets must not be decoded before transmission across that protocol. Transfer of the percent-encoding, which might violate the protocol, is less harmful than allowing decoded octets to be interpreted as additional operations or parameters, perhaps triggering an unexpected and possibly harmful remote operation.

## [7.3](#) Back-end Transcoding

When a URI is dereferenced, the data within it is often parsed by both the user agent and one or more servers. In HTTP, for example, a typical user agent will parse a URI into its five major components,

access the authority's server, and send it the data within the authority, path, and query components. A typical server will take that information, parse the path into segments and the query into key/value pairs, and then invoke implementation-specific handlers to respond to the request. As a result, a common security concern for server implementations that handle a URI, either as a whole or split into separate components, is proper interpretation of the octet data represented by the characters and percent-encodings within that URI.

Percent-encoded octets must be decoded at some point during the dereference process. Applications must split the URI into its components and sub-components prior to decoding the octets, since otherwise the decoded octets might be mistaken for delimiters. Security checks of the data within a URI should be applied after decoding the octets. Note, however, that the "%00" percent-encoding (NUL) may require special handling and should be rejected if the application is not expecting to receive raw data within a component.

Special care should be taken when the URI path interpretation process involves the use of a back-end filesystem or related system functions. Filesystems typically assign an operational meaning to special characters, such as the "/", "\", ":", "[", and "]" characters, and special device names like ".", "..", "...", "aux", "lpt", etc. In some cases, merely testing for the existence of such a name will cause the operating system to pause or invoke unrelated system calls, leading to significant security concerns regarding denial of service and unintended data transfer. It would be impossible for this specification to list all such significant characters and device names; implementers should research the reserved names and characters for the types of storage device that may be attached to their application and restrict the use of data obtained from URI components accordingly.

#### [7.4](#) Rare IP Address Formats

Although the URI syntax for IPv4address only allows the common, dotted-decimal form of IPv4 address literal, many implementations that process URIs make use of platform-dependent system routines,



such as `gethostbyname()` and `inet_aton()`, to translate the string literal to an actual IP address. Unfortunately, such system routines often allow and process a much larger set of formats than those described in [Section 3.2.2](#).

For example, many implementations allow dotted forms of three numbers, wherein the last part is interpreted as a 16-bit quantity and placed in the right-most two bytes of the network address (e.g., a Class B network). Likewise, a dotted form of two numbers means the last part is interpreted as a 24-bit quantity and placed in the right most three bytes of the network address (Class A), and a single number (without dots) is interpreted as a 32-bit quantity and stored directly in the network address. Adding further to the confusion, some implementations allow each dotted part to be interpreted as decimal, octal, or hexadecimal, as specified in the C language (i.e., a leading 0x or 0X implies hexadecimal; otherwise, a leading 0 implies octal; otherwise, the number is interpreted as decimal).

These additional IP address formats are not allowed in the URI syntax due to differences between platform implementations. However, they can become a security concern if an application attempts to filter access to resources based on the IP address in string literal format. If such filtering is performed, literals should be converted to numeric form and filtered based on the numeric value, rather than a prefix or suffix of the string form.

## [7.5](#) Sensitive Information

URI producers should not provide a URI that contains a username or password which is intended to be secret: URIs are frequently displayed by browsers, stored in clear text bookmarks, and logged by user agent history and intermediary applications (proxies). A password appearing within the userinfo component is deprecated and should be considered an error (or simply ignored) except in those rare cases where the 'password' parameter is intended to be public.

## [7.6](#) Semantic Attacks

Because the userinfo sub-component is rarely used and appears before the host in the authority component, it can be used to construct a URI that is intended to mislead a human user by appearing to identify one (trusted) naming authority while actually identifying a different authority hidden behind the noise. For example

```
ftp://ftp.example.com&story=breaking_news@10.0.0.1/top_story.htm
```

might lead a human user to assume that the host is 'trusted.example.com', whereas it is actually '10.0.0.1'. Note that a misleading userinfo sub-component could be much longer than the example above.

A misleading URI, such as the one above, is an attack on the user's preconceived notions about the meaning of a URI, rather than an attack on the software itself. User agents may be able to reduce the impact of such attacks by distinguishing the various components of the URI when rendered, such as by using a different color or tone to render userinfo if any is present, though there is no general panacea. More information on URI-based semantic attacks can be found in [[Siedzik](#)].

## 8. Acknowledgments

This specification is derived from [RFC 2396](#) [[RFC2396](#)], [RFC 1808](#) [[RFC1808](#)], and [RFC 1738](#) [[RFC1738](#)]; the acknowledgments in those documents still apply. It also incorporates the update (with corrections) for IPv6 literals in the host syntax, as defined by Robert M. Hinden, Brian E. Carpenter, and Larry Masinter in [[RFC2732](#)]. In addition, contributions by Gisle Aas, Reese Anschultz, Daniel Barclay, Tim Bray, Mike Brown, Rob Cameron, Jeremy Carroll, Dan Connolly, Adam M. Costello, John Cowan, Jason Diamond, Martin Duerst, Stefan Eissing, Clive D.W. Feather, Tony Hammond, Pat Hayes, Henry Holtzman, Ian B. Jacobs, Michael Kay, John C. Klensin, Graham Klyne, Dan Kohn, Bruce Lilly, Andrew Main, Ira McDonald, Michael Mealling, Stephen Pollei, Julian Reschke, Tomas Rokicki, Miles Sabin, Mark Thomson, Ronald Tschalaer, Norm Walsh, Marc Warne, Stuart Williams, and Henry Zongaro are gratefully acknowledged.

Internet-Draft

URI Generic Syntax

February 2004

## Normative References

- [ASCII] American National Standards Institute, "Coded Character Set -- 7-bit American Standard Code for Information Interchange", ANSI X3.4, 1986.
- [RFC2234] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", [RFC 2234](#), November 1997.
- [RFC3629] Yergeau, F., "UTF-8, a transformation format of ISO 10646", STD 63, [RFC 3629](#), November 2003.

Internet-Draft

URI Generic Syntax

February 2004

#### Informative References

- [RFC0952] Harrenstien, K., Stahl, M. and E. Feinler, "DoD Internet host table specification", [RFC 952](#), October 1985.
- [RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, [RFC 1034](#), November 1987.
- [RFC1123] Braden, R., "Requirements for Internet Hosts - Application and Support", STD 3, [RFC 1123](#), October 1989.
- [RFC1535] Gavron, E., "A Security Problem and Proposed Correction With Widely Deployed DNS Software", [RFC 1535](#), October 1993.
- [RFC1630] Berners-Lee, T., "Universal Resource Identifiers in WWW: A Unifying Syntax for the Expression of Names and Addresses of Objects on the Network as used in the World-Wide Web", [RFC 1630](#), June 1994.
- [RFC1736] Kunze, J., "Functional Recommendations for Internet Resource Locators", [RFC 1736](#), February 1995.
- [RFC1737] Masinter, L. and K. Sollins, "Functional Requirements for Uniform Resource Names", [RFC 1737](#), December 1994.
- [RFC1738] Berners-Lee, T., Masinter, L. and M. McCahill, "Uniform Resource Locators (URL)", [RFC 1738](#), December 1994.

- [RFC1808] Fielding, R., "Relative Uniform Resource Locators", [RFC 1808](#), June 1995.
- [RFC2046] Freed, N. and N. Borenstein, "Multipurpose Internet Mail Extensions (MIME) Part Two: Media Types", [RFC 2046](#), November 1996.
- [RFC2110] Palme, J. and A. Hopmann, "MIME E-mail Encapsulation of Aggregate Documents, such as HTML (MHTML)", [RFC 2110](#), March 1997.
- [RFC2141] Moats, R., "URN Syntax", [RFC 2141](#), May 1997.
- [RFC2277] Alvestrand, H., "IETF Policy on Character Sets and Languages", [BCP 18](#), [RFC 2277](#), January 1998.
- [RFC2396] Berners-Lee, T., Fielding, R. and L. Masinter, "Uniform Resource Identifiers (URI): Generic Syntax", [RFC 2396](#), August 1998.

Berners-Lee, et al.

Expires August 16, 2004

[Page 47]

---

Internet-Draft

URI Generic Syntax

February 2004

- [RFC2518] Goland, Y., Whitehead, E., Faizi, A., Carter, S. and D. Jensen, "HTTP Extensions for Distributed Authoring -- WEBDAV", [RFC 2518](#), February 1999.
- [RFC2717] Petke, R. and I. King, "Registration Procedures for URL Scheme Names", [BCP 35](#), [RFC 2717](#), November 1999.
- [RFC2718] Masinter, L., Alvestrand, H., Zigmond, D. and R. Petke, "Guidelines for new URL Schemes", [RFC 2718](#), November 1999.
- [RFC2732] Hinden, R., Carpenter, B. and L. Masinter, "Format for Literal IPv6 Addresses in URL's", [RFC 2732](#), December 1999.
- [RFC2978] Freed, N. and J. Postel, "IANA Charset Registration Procedures", [BCP 19](#), [RFC 2978](#), October 2000.
- [RFC3305] Mealling, M. and R. Denenberg, "Report from the Joint W3C/IETF URI Planning Interest Group: Uniform Resource Identifiers (URIs), URLs, and Uniform Resource Names (URNs): Clarifications and Recommendations", [RFC 3305](#), August 2002.

- [RFC3490] Faltstrom, P., Hoffman, P. and A. Costello,  
"Internationalizing Domain Names in Applications (IDNA)",  
[RFC 3490](#), March 2003.
- [RFC3513] Hinden, R. and S. Deering, "Internet Protocol Version 6  
(IPv6) Addressing Architecture", [RFC 3513](#), April 2003.
- [Siedzik] Siedzik, R., "Semantic Attacks: What's in a URL?", April  
2001, <[http://www.giac.org/practical/gsec/  
Richard\\_Siedzik\\_GSEC.pdf](http://www.giac.org/practical/gsec/Richard_Siedzik_GSEC.pdf)>.

#### Authors' Addresses

Tim Berners-Lee  
World Wide Web Consortium  
MIT/LCS, Room NE43-356  
200 Technology Square  
Cambridge, MA 02139  
USA

Phone: +1-617-253-5702  
Fax: +1-617-258-5999  
EMail: [timbl@w3.org](mailto:timbl@w3.org)  
URI: <http://www.w3.org/People/Berners-Lee/>

Roy T. Fielding  
Day Software  
5251 California Ave., Suite 110  
Irvine, CA 92612-3074  
USA

Phone: +1-949-679-2960  
Fax: +1-949-679-2972  
EMail: [fielding@gbiv.com](mailto:fielding@gbiv.com)  
URI: <http://roy.gbiv.com/>

Larry Masinter  
Adobe Systems Incorporated  
345 Park Ave  
San Jose, CA 95110  
USA

Phone: +1-408-536-3024  
EMail: [LMM@acm.org](mailto:LMM@acm.org)  
URI: <http://larry.masinter.net/>

## [Appendix A](#). Collected ABNF for URI

URI = scheme ":" ["/" authority] path ["?" query] ["#" fragment]

URI-reference = URI / relative-URI

relative-URI = ["/" authority] path ["?" query] ["#" fragment]



```

absolute-URI = scheme ":" ["/" authority] path ["?" query]

scheme        = ALPHA *( ALPHA / DIGIT / "+" / "-" / "." )

authority     = [ userinfo "@" ] host [ ":" port ]
userinfo      = *( unreserved / pct-encoded / sub-delims / ":" )
host          = IP-literal / IPv4address / reg-name
port          = *DIGIT

IP-literal    = "[" ( IPv6address / IPvFuture  ) "]"

IPvFuture     = "v" HEXDIG "." 1*( unreserved / sub-delims / ":" )

IPv6address   =
/                               6( h16 ":" ) ls32
/                               "::" 5( h16 ":" ) ls32
/ [                               h16 ] "::" 4( h16 ":" ) ls32
/ [ *1( h16 ":" ) h16 ] "::" 3( h16 ":" ) ls32
/ [ *2( h16 ":" ) h16 ] "::" 2( h16 ":" ) ls32
/ [ *3( h16 ":" ) h16 ] "::"   h16 ":"   ls32
/ [ *4( h16 ":" ) h16 ] "::"
/ [ *5( h16 ":" ) h16 ] "::"
/ [ *6( h16 ":" ) h16 ] "::"

h16           = 1*4HEXDIG
ls32          = ( h16 ":" h16 ) / IPv4address

IPv4address   = dec-octet "." dec-octet "." dec-octet "." dec-octet

dec-octet     = DIGIT                     ; 0-9
/ %x31-39 DIGIT                     ; 10-99
/ "1" 2DIGIT                          ; 100-199
/ "2" %x30-34 DIGIT                   ; 200-249
/ "25" %x30-35                        ; 250-255

reg-name      = 0*255( unreserved / pct-encoded / sub-delims )

path          = segment *( "/" segment )
segment       = *pchar

query         = *( pchar / "/" / "?" )

```

```
fragment      = *( pchar / "/" / "?" )

pct-encoded    = "%" HEXDIG HEXDIG

pchar          = unreserved / pct-encoded / sub-delims / ":" / "@"

unreserved     = ALPHA / DIGIT / "-" / "." / "_" / "~"
reserved       = gen-delims / sub-delims
gen-delims     = ":" / "/" / "?" / "#" / "[" / "]" / "@"
sub-delims     = "!" / "$" / "&" / "'" / "(" / ")"
               / "*" / "+" / "," / ";" / "="
```

Appendix B. Parsing a URI Reference with a Regular Expression

Since the "first-match-wins" algorithm is identical to the "greedy" disambiguation method used by POSIX regular expressions, it is natural and commonplace to use a regular expression for parsing the potential five components of a URI reference.

The following line is the regular expression for breaking-down a well-formed URI reference into its components.

```
^((([^\?#]+):)?(//([^\?#]*)?)?([^\?#]*)?(\?([^\#]*)?)?(#(.*)?)?)
```

12                      3   4                      5                      6   7                      8   9

The numbers in the second line above are only to assist readability; they indicate the reference points for each subexpression (i.e., each paired parenthesis). We refer to the value matched for subexpression <n> as \$<n>. For example, matching the above expression to

<http://www.ics.uci.edu/pub/ietf/uri/#Related>

results in the following subexpression matches:

```
$1 = http:
$2 = http
$3 = //www.ics.uci.edu
$4 = www.ics.uci.edu
$5 = /pub/ietf/uri/
$6 = <undefined>
$7 = <undefined>
$8 = #Related
$9 = Related
```

where <undefined> indicates that the component is not present, as is the case for the query component in the above example. Therefore, we can determine the value of the four components and fragment as

```
scheme      = $2
authority   = $4
path        = $5
query       = $7
fragment    = $9
```

and, going in the opposite direction, we can recreate a URI reference

from its components using the algorithm of [Section 5.3](#).

## [Appendix C](#). Delimiting a URI in Context

URIs are often transmitted through formats that do not provide a clear context for their interpretation. For example, there are many occasions when a URI is included in plain text; examples include text sent in electronic mail, USENET news messages, and, most importantly, printed on paper. In such cases, it is important to be able to delimit the URI from the rest of the text, and in particular from punctuation marks that might be mistaken for part of the URI.

In practice, URIs are delimited in a variety of ways, but usually within double-quotes "http://example.com/", angle brackets <http://example.com/>, or just using whitespace

```
http://example.com/
```

These wrappers do not form part of the URI.

In some cases, extra whitespace (spaces, line-breaks, tabs, etc.) may need to be added to break a long URI across lines. The whitespace should be ignored when extracting the URI.

No whitespace should be introduced after a hyphen ("-") character. Because some typesetters and printers may (erroneously) introduce a hyphen at the end of line when breaking a line, the interpreter of a URI containing a line break immediately after a hyphen should ignore all whitespace around the line break, and should be aware that the hyphen may or may not actually be part of the URI.

Using <> angle brackets around each URI is especially recommended as a delimiting style for a reference that contains embedded whitespace.

The prefix "URL:" (with or without a trailing space) was formerly recommended as a way to help distinguish a URI from other bracketed designators, though it is not commonly used in practice and is no longer recommended.

For robustness, software that accepts user-typed URI should attempt to recognize and strip both delimiters and embedded whitespace.

For example, the text:

Yes, Jim, I found it under "http://www.w3.org/Addressing/", but you can probably pick it up from <ftp://foo.example.com/rfc/>. Note the warning in <<http://www.ics.uci.edu/pub/ietf/uri/historical.html#WARNING>>.

contains the URI references

<http://www.w3.org/Addressing/>  
<ftp://foo.example.com/rfc/>  
<http://www.ics.uci.edu/pub/ietf/uri/historical.html#WARNING>

## [Appendix D](#). Summary of Non-editorial Changes

### [D.1](#) Additions

IPv6 (and later) literals have been added to the list of possible identifiers for the host portion of a authority component, as described by [\[RFC2732\]](#), with the addition of "[" and "]" to the reserved set and a version flag to anticipate future versions of IP literals. Square brackets are now specified as reserved within the authority component and not allowed outside their use as delimiters for an IP literal within host. In order to make this change without changing the technical definition of the path, query, and fragment components, those rules were redefined to directly specify the characters allowed rather than be defined in terms of uric.

Since [\[RFC2732\]](#) defers to [\[RFC3513\]](#) for definition of an IPv6 literal address, which unfortunately lacks an ABNF description of IPv6address, we created a new ABNF rule for IPv6address that matches the text representations defined by [Section 2.2 of \[RFC3513\]](#). Likewise, the definition of IPv4address has been improved in order to limit each decimal octet to the range 0-255.

[Section 6](#) ([Section 6](#)) on URI normalization and comparison has been completely rewritten and extended using input from Tim Bray and discussion within the W3C Technical Architecture Group.

An ABNF rule for URI has been introduced to correspond to the common usage of the term: an absolute URI with optional fragment.

## [D.2](#) Modifications from [RFC 2396](#)

The ad-hoc BNF syntax has been replaced with the ABNF of [[RFC2234](#)]. This change required all rule names that formerly included underscore characters to be renamed with a dash instead.

[Section 2](#) on characters has been rewritten to explain what characters are reserved, when they are reserved, and why they are reserved even when not used as delimiters by the generic syntax. The mark characters that are typically unsafe to decode, including the exclamation mark ("!"), asterisk ("\*"), single-quote ("'"), and open and close parentheses ("(" and ")"), have been moved to the reserved set in order to clarify the distinction between reserved and unreserved and hopefully answer the most common question of scheme designers. Likewise, the section on percent-encoded characters has been rewritten, and URI normalizers are now given license to decode any percent-encoded octets corresponding to unreserved characters. In general, the terms "escaped" and "unesaped" have been replaced with "percent-encoded" and "decoded", respectively, to reduce

confusion with other forms of escape mechanisms.

The ABNF for URI and URI-reference has been redesigned to make them more friendly to LALR parsers and significantly reduce complexity. As a result, the layout form of syntax description has been removed, along with the `uric`, `uric_no_slash`, `hier_part`, `opaque_part`, `net_path`, `abs_path`, `rel_path`, `path_segments`, `rel_segment`, and `mark` rules. All references to "opaque" URIs have been replaced with a better description of how the path component may be opaque to hierarchy. The ambiguity regarding the parsing of URI-reference as a URI or a relative-URI with a colon in the first segment is now explained and disambiguated in the section defining relative-URI.

The fragment identifier has been moved back into the section on generic syntax components and within the URI and relative-URI rules,

though it remains excluded from absolute-URI. The number sign ("#") character has been moved back to the reserved set as a result of reintegrating the fragment syntax.

The ABNF has been corrected to allow a relative path to be empty. This also allows an absolute-URI to consist of nothing after the "scheme:", as is present in practice with the "dav:" namespace [[RFC2518](#)] and the "about:" scheme used internally by many WWW browser implementations. The ambiguity regarding the boundary between authority and path is now explained and disambiguated in the same section.

Registry-based naming authorities that use the generic syntax are now defined within the host rule and limited to 255 path characters. This change allows current implementations, where whatever name provided is simply fed to the local name resolution mechanism, to be consistent with the specification and removes the need to re-specify DNS name formats here. It also allows the host component to contain percent-encoded octets, which is necessary to enable internationalized domain names to be provided in URIs, processed in their native character encodings at the application layers above URI processing, and passed to an IDNA library as a registered name in the UTF-8 character encoding. The server, hostport, hostname, domainlabel, toplabel, and alphanum rules have been removed.

The resolving relative references algorithm of [[RFC2396](#)] has been rewritten using pseudocode for this revision to improve clarity and fix the following issues:

- o [[RFC2396](#)] [section 5.2](#), step 6a, failed to account for a base URI with no path.

- o Restored the behavior of [[RFC1808](#)] where, if the reference contains an empty path and a defined query component, then the target URI inherits the base URI's path component.
- o Removed the special-case treatment of same-document references within the URI parser in favor of a section that explains when a reference should be interpreted by a dereferencing engine as a same-document reference: when the target URI and base URI,



excluding fragments, match. This change does not modify the behavior of existing same-document references as defined by [RFC 2396](#) (fragment-only references); it merely adds the same-document distinction to other references that refer to the base URI and simplifies the interface between applications and their URI parsers, as is consistent with the internal architecture of deployed URI processing implementations.

- o Separated the path merge routine into two routines: merge, for describing combination of the base URI path with a relative-path reference, and remove\_dot\_segments, for describing how to remove the special "." and ".." segments from a composed path. The remove\_dot\_segments algorithm is now applied to all URI reference paths in order to match common implementations and improve the normalization of URIs in practice. This change only impacts the parsing of abnormal references and same-scheme references wherein the base URI has a non-hierarchical path.

A  
    ABNF 10  
    absolute 25  
    absolute-path 24  
    absolute-URI 25  
    access 7  
    authority 15, 16

B  
    base URI 27

C  
    characters 11

D  
    dec-octet 18  
    dereference 7  
    dot-segments 20

F  
    fragment 22

G  
    gen-delims 12  
    generic syntax 5

H  
    h16 17  
    hierarchical 9  
    host 17

I  
    identifier 5  
    IP-literal 17  
    IPv4 18  
    IPv4address 18  
    IPv6 17  
    IPv6address 17  
    IPvFuture 17

L  
    locator 6  
    ls32 17

M  
    merge 30

## N

name 6  
network-path 24

## P

path 15, 20  
pchar 20  
pct-encoded 11  
percent-encoding 11  
port 20

## Q

query 21

## R

reg-name 19  
registered name 19  
relative 9, 27  
relative-path 24  
relative-URI 24  
remove\_dot\_segments 30  
representation 8  
reserved 12  
resolution 7, 27  
resource 4  
retrieval 8

## S

same-document 25  
sameness 8  
scheme 15  
segment 20  
sub-delims 12  
suffix 25

## T

transcription 6

Internet-Draft

URI Generic Syntax

February 2004

## U

- uniform 4
- unreserved 12
- URI grammar
  - absolute-URI 25
  - ALPHA 10
  - authority 15, 16
  - CR 10
  - CTL 10
  - dec-octet 18
  - DIGIT 10
  - DQUOTE 10
  - fragment 15, 22, 24
  - gen-delims 12
  - h16 18
  - HEXDIG 10
  - host 16, 17
  - IP-literal 17
  - IPv4address 18
  - IPv6address 17, 18
  - IPvFuture 17
  - LF 10
  - ls32 18
  - mark 12
  - OCTET 10
  - path 15
  - path-segments 20
  - pchar 20, 21, 22
  - pct-encoded 11
  - port 16, 20
  - query 15, 21, 24, 25
  - reg-name 19
  - relative-URI 24, 24
  - reserved 12
  - scheme 15, 15, 25
  - segment 20
  - SP 10
  - sub-delims 12
  - unreserved 12

URI 15, 24  
URI-reference 24  
userinfo 16, 16  
URI 15  
URI-reference 24  
URL 6  
URN 6  
userinfo 16

Berners-Lee, et al. Expires August 16, 2004

[Page 60]

---

Internet-Draft

URI Generic Syntax

February 2004

#### Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any intellectual property 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; neither does it represent that it has made any effort to identify any such rights. Information on the IETF's procedures with respect to rights in standards-track and standards-related documentation can be found in [BCP-11](#). Copies of claims of rights made available for publication 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 implementors or users of this specification can be obtained from the IETF Secretariat.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this standard. Please address the information to the IETF Executive Director.

#### Full Copyright Statement

Copyright (C) The Internet Society (2004). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are

included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assignees.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION

Berners-Lee, et al. Expires August 16, 2004

[Page 61]

---

Internet-Draft

URI Generic Syntax

February 2004

HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

#### Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.

