



Extensible Resource Identifier (XRI) Generic Syntax and Resolution Specification

Release Candidate 2, 20 November 2003

Document identifier:

wd-xri-specification-rc2

Location:

<http://www.oasis-open.org/committees/xri>

Editors:

Gabe Wachob, Visa International <gwachob@visa.com>
Drummond Reed, OneName <drummond.reed@onename.com>
Dave McAlpin, Epok <dave.mcalpin@epok.net>
Mike Lindelsee, Visa International <mlindels@visa.com>
Peter Davis, Neustar <peter.davis@neustar.biz>
Nat Sakimura, NRI <n-sakimura@nri.co.jp>

Abstract:

This document is the normative technical specification for XRI generic syntax and resolution. For a non-normative introduction to the uses and features of XRIs, see the *XRI Primer*.

Status:

This document is a working draft updated periodically on no particular schedule. Send comments to the editors.

Committee members should send comments on this specification to the xri@lists.oasis-open.org list. Others should subscribe to and send comments to the xri-comment@lists.oasis-open.org list. To subscribe, send an email message to xri-comment-request@lists.oasis-open.org with the word "subscribe" as the body of the message.

For information on whether any patents have been disclosed that may be essential to implementing this specification, and any offers of patent licensing terms, please refer to the Intellectual Property Rights section of the XRI TC web page (<http://www.oasis-open.org/committees/xri/>).

The errata page for this specification is at <http://www.oasis-open.org/committees/xri/yyy>.



34 **Table of Contents**

35 Introduction 5

36 1.1 Overview of XRIs..... 5

37 1.1.1 Generic Syntax 5

38 1.1.2 Examples 6

39 1.1.3 URI, URL, URN, and XRI..... 7

40 1.2 Design Considerations 7

41 1.2.1 Abstraction and Independence 7

42 1.2.2 Persistence and Reassignability 8

43 1.2.3 Human-Friendliness and Machine-Friendliness 8

44 1.2.4 Internationalization 8

45 1.2.5 Cross-Context Identification..... 8

46 1.2.6 Authority, Delegation, and Federation 8

47 1.2.7 Security and Privacy 8

48 1.2.8 Extensibility 8

49 1.3 Terminology and Notation 8

50 1.3.1 Keywords 8

51 1.3.2 Syntax Notation..... 9

52 1.3.3 Glossary 9

53 2 Syntax 13

54 2.1 Syntax Components 13

55 2.1.1 Authority 13

56 2.1.1.1 URI Authority 13

57 2.1.1.2 XRI Authority..... 15

58 2.1.1.3 Global Context Symbols (GCS) 15

59 2.1.1.4 Cross-References 16

60 2.1.1.5 Self-References 16

61 2.1.2 Path..... 17

62 2.1.3 Query 17

63 2.1.4 Fragment..... 18

64 2.2 Characters..... 18

65 2.2.1 Character Encoding 18

66 2.2.2 Reserved Characters 18

67 2.2.3 Unreserved Characters..... 18

68 2.2.4 Escaped Characters 19

69 2.2.4.1 Escaped Encoding..... 19

70 2.2.4.2 Encoding XRI Metadata..... 19

71 2.2.4.3 Transforming XRIs into IRIs and URIs 20

72 2.2.4.4 Special Escaping Rules for XRI Syntax 21

73 2.2.4.5 Transforming URIs and IRIs Back into XRIs 22

74 2.2.5 Excluded Characters..... 23

75 2.3 Relative XRI References 24

76 2.3.1 Establishing a Base XRI 24

| | | |
|-----|---|----|
| 77 | 2.3.2 Obtaining the Referenced XRI..... | 24 |
| 78 | 2.3.3 Leading Segments Containing a Colon..... | 25 |
| 79 | 2.4 Normalization and Comparison..... | 25 |
| 80 | 2.4.1 Case..... | 26 |
| 81 | 2.4.2 Encoding, Escaping, and Transformations..... | 26 |
| 82 | 2.4.3 Optional Syntax..... | 26 |
| 83 | 2.4.4 Cross-References..... | 27 |
| 84 | 2.4.5 Canonicalization..... | 27 |
| 85 | 3 Resolution..... | 29 |
| 86 | 3.1 Introduction..... | 29 |
| 87 | 3.1.1 Assumptions..... | 29 |
| 88 | 3.1.2 Phases of Resolution..... | 29 |
| 89 | 3.1.3 XRI vs. URI Authorities..... | 30 |
| 90 | 3.1.4 XRI Metadata Reserved for XRI Resolution..... | 30 |
| 91 | 3.2 XRI Authority Resolution..... | 30 |
| 92 | 3.2.1 Overview..... | 30 |
| 93 | 3.2.2 Identifier Authority Descriptors..... | 31 |
| 94 | 3.2.3 Initiating Resolution..... | 33 |
| 95 | 3.2.4 Iterating Resolution..... | 33 |
| 96 | 3.2.5 Examples..... | 34 |
| 97 | 3.2.6 Resolving Cross-References in XRI Authorities..... | 36 |
| 98 | 3.2.7 User Relative XRIs..... | 36 |
| 99 | 3.3 URI Authority Resolution..... | 37 |
| 100 | 3.4 Local Access..... | 37 |
| 101 | 3.4.1 Local Access Service Types..... | 37 |
| 102 | 3.4.2 HTTP/HTTPS Local Access..... | 37 |
| 103 | 3.4.3 Constructing a Local Access HTTP/HTTPS URI..... | 38 |
| 104 | 3.4.4 Using a Cross-Reference to Specify a Representation Type..... | 38 |
| 105 | 3.5 HTTP Headers..... | 39 |
| 106 | 3.5.1 Caching..... | 39 |
| 107 | 3.5.2 Location..... | 39 |
| 108 | 3.5.3 Content-Location..... | 39 |
| 109 | 3.5.4 Content-Type..... | 40 |
| 110 | 3.5.5 X-XRI-Canonical..... | 40 |
| 111 | 3.6 Other HTTP Features..... | 40 |
| 112 | 3.7 Caching and Efficiency..... | 40 |
| 113 | 3.8 Points of Extensibility..... | 41 |
| 114 | 4 Security and Data Protection..... | 42 |
| 115 | 4.1 Secure Resolution..... | 42 |
| 116 | 4.2 XRI Metadata..... | 42 |
| 117 | 4.3 XRI Usage in Legacy Infrastructure..... | 42 |
| 118 | 4.4 XRI Usage in Evolving Infrastructure..... | 42 |
| 119 | 5 References..... | 43 |
| 120 | 5.1 Normative..... | 43 |

| | | |
|-----|---|----|
| 121 | 5.2 Informative..... | 44 |
| 122 | Appendix A. Collected ABNF for XRI (Normative) | 45 |
| 123 | Appendix B. XML Schema for XRI Identifier Authority Descriptor (Normative)..... | 48 |
| 124 | Appendix C. Transforming HTTP URIs to XRIs (Non-Normative)..... | 50 |
| 125 | Appendix D. Acknowledgments..... | 51 |
| 126 | Appendix E. Revision History | 52 |
| 127 | Appendix F. Notices | 53 |
| 128 | | |

129 Introduction

130 1.1 Overview of XRIs

131 An Extensible Resource Identifier (XRI) provides a standard means of abstractly identifying a
132 resource independent of any particular concrete representation of that resource—or, in the case
133 of a completely abstract resource, of any representation at all.

134 XRIs are similar to URIs as defined in “*Uniform Resource Identifiers (URI): Generic Syntax*”
135 [RFC2396], but contain additional syntactic elements and extend the unreserved character set to
136 include characters beyond those allowed in generic URIs. To accommodate applications that
137 expect generic URIs, the XRI specification defines rules for transforming an XRI into a valid URI
138 as defined by [RFC2396]. Since a revision of RFC 2396 is currently in progress, the XRI scheme
139 also incorporates some simplifications and enhancements to generic URI syntax as proposed in
140 [RFC2396bis].

141 XRI syntax is internationalized following the recommendations in “*Guidelines for New URL*
142 *Schemes*” [RFC2718] and “*Extensible Markup Language (XML) 1.0 (Second Edition)*” [XML], and
143 specifically the requirements of the “anyURI” datatype as specified in “XML Schema Part 2:
144 Datatypes” [XMLSchema2]. To do this, the XRI scheme incorporates the syntax recommended
145 in another work-in-progress, “*Internationalized Resource Identifiers (IRIs)*” [IRI].

146 Although an XRI is not a Uniform Resource Name (URN) as defined in “*URN Syntax*” [RFC2141],
147 XRIs consisting entirely of persistent segments are designed to meet the requirements set out in
148 “*Functional Requirements for Uniform Resource Names*” [RFC1737].

149 This document specifies the ABNF for the XRI scheme. In addition it specifies an HTTP-based
150 resolution protocol for XRIs. Use of this protocol is not required; XRIs may also be resolved using
151 other protocols or resolution mechanisms.

152 While [RFC2396bis] and [IRI] are cited in this document, they are both works in progress and are
153 consequently non-normative. All relevant information from these proposals is reproduced here, so
154 access to these documents, while very informative, is not required.

155 1.1.1 Generic Syntax

156 XRI syntax is designed to be as simple and extensible as URI syntax. A fully-qualified XRI
157 consists of the scheme name “xri:” followed by the same four optional components as a generic
158 URI.

159

```
160 xri: authority / path ? query # fragment
```

161

162 The definitions of these components are, for the most part, supersets of the equivalent
163 components in the generic URI syntax. One advantage of this approach is that the vast majority
164 of HTTP URIs, which inherit directly from generic URI syntax, can be transformed to valid XRIs
165 simply by changing the scheme from “http” to “xri”. The rules for this transformation are
166 summarized in Appendix C, “Transforming HTTP URIs to XRIs”.

167 XRI syntax extends generic URI syntax in six ways by providing support for:

- 168 1. *Persistent and reassignable segments*. Generic URI syntax does not distinguish between
169 persistent and reassignable identifiers. XRI syntax enables the top-level authority
170 segment as well as any subsequent path segment to be explicitly designated as either
171 persistent or reassignable.

- 172 2. *Unlimited delegation.* Generic URI syntax supports delegated identifiers (i.e., DNS names
173 or IP addresses) only within the top-level authority segment. XRI syntax supports
174 delegation of both persistent and reassignable identifiers at any level of the path.
- 175 3. *Global context symbols.* While XRI syntax supports the same generic URI syntax for DNS
176 and IP authorities, it also provides shorthand symbols for establishing the abstract
177 context of an identifier.
- 178 4. *Cross-references.* Generic URI syntax does not provide a way to share identifiers across
179 contexts. This capability is particularly useful with abstract identifiers (e.g., to establish
180 the generic type of a resource, or to share standardized identifier metadata such as
181 versioning). For this reason, XRI syntax allows XRIs (and URIs) to be shared across
182 contexts by means of parenthetical nesting.
- 183 5. *Self-references.* Generic URI syntax does not provide a way to indicate whether or not a
184 URI is intended for resolution. Since an XRI may itself be the full representation of a
185 abstract non-network resource (e.g., “love,” “Paris,” or “the planet Jupiter”), XRI syntax
186 provides a way to express self-reference.
- 187 6. *Internationalized character set.* Generic URI syntax limits legal characters to a subset of
188 the US-ASCII character set. XRI syntax, following the lead of Internationalized Resource
189 Identifiers [IRI], employs the broader Unicode character set, making the use of XRIs in
190 languages other than English much more straightforward.

191 1.1.2 Examples

192 The following examples illustrate XRI syntax. These examples have minimal annotation and are
193 only intended to give a sense of the scope and flavor of XRI syntax. For more information on the
194 normative syntax, see section 2. For a complete description of the uses and features of XRIs, see
195 the non-normative XRI Primer [tktk need reference].
196

```
197 xri://www.example.com/pages/index.html  
198     --standard HTTP URI converted to an XRI  
199  
200 xri://[2010:836B:4179::836B:4179]/pages/index.html  
201     --using an IPv6 authority per RFC 2732  
202  
203 xri://www.example.com/inventory.parts/widget.subwidget.foobarator  
204     --delegation of reassignable identifiers  
205  
206 xri://www.example.com/:inventory:parts/:12:7:234  
207     --delegation of persistent identifiers  
208  
209 xri:@ExampleCorp  
210 xri:@ExampleCorp.www  
211 xri:@ExampleCorp.website  
212 xri:=JohnDoe  
213 xri:=JohnDoe.home  
214 xri:=JohnDoe.work  
215 xri:+flowers  
216 xri:+flowers.rose  
217 xri:+flowers.daisy  
218     --global context symbols  
219  
220 xri://www.example.com/(+management)/(+CEO)  
221 xri:(urn:oasis:spec:2040)/(+index)  
222 xri:(mailto:john.doe@example.com)/(+phone)  
223 xri:=JohnDoe.home/(+email)  
224 xri:=JohnDoe.home/(+email).($v/3)  
225     --cross-references  
226  
227 xri:(+flowers.rose)
```

228
229
230
231

```
xri://www.example.com/dictionary/flowers/rose)
xri:(http://www.example.com/dictionary/flowers/rose)
--self-references
```

232
233
234

Table 1 also illustrates several examples of internationalized XRIs.

| | |
|--------|-------------------------------|
| French | xri:@ALaFrançaise/areté |
| Hebrew | xri://אג.גה.הפ/גח/ט' /גז.html |
| Kanji | xri:=崎村夏彦/ (+本籍地) |

235
236

Table 1: Internationalized XRIs.

237 1.1.3 URI, URL, URN, and XRI

238 The evolution and interrelationships of the terms “URI”, “URL”, and “URN” are explained in a
239 report from the Joint W3C/IETF URI Planning Interest Group, “Uniform Resource Identifiers
240 (URIs), URLs, and Uniform Resource Names (URNs): Clarifications and Recommendations”
241 [RFC3305]. According to section 2.1:

242 “During the early years of discussion of web identifiers (early to mid 90s), people assumed
243 that an identifier type would be cast into one of two (or possibly more) classes. An identifier
244 might specify the location of a resource (a URL) or its name (a URN), independent of
245 location. Thus a URI was either a URL or a URN.”

246 This view has since changed, as the report goes on to state in section 2.2:

247 “Over time, the importance of this additional level of hierarchy seemed to lessen; the view
248 became that an individual scheme did not need to be cast into one of a discrete set of URI
249 types, such as ‘URL’, ‘URN’, ‘URC’, etc. Web-identifier schemes are, in general, URI
250 schemes, as a given URI scheme may define subspaces.”

251 This conclusion is shared by [RFC2396bis], which states in section 1.1.3:

252 “An individual [URI] scheme does not need to be classified as being just one of ‘name’ or
253 ‘locator’. Instances of URIs from any given scheme may have the characteristics of names or
254 locators or both, often depending on the persistence and care in the assignment of identifiers
255 by the identifier authority, rather than any quality of the scheme.”

256 The XRI scheme follows this philosophy. XRIs can be used either as persistent “names” for
257 resources or as concrete “locators” for resources, including other XRIs. The XRI scheme also
258 includes syntax for distinguishing whether an XRI is intended only for identification or also for
259 resolution. For more information, see section 2.1.1.4, *Self-References*.

260 1.2 Design Considerations

261 The full set of requirements for XRI syntax and resolution is documented in “XRI Requirements
262 and Glossary v1.0” [XRIReqs]. A synopsis of the major design considerations is included here.

263 1.2.1 Abstraction and Independence

264 The overarching requirement of the XRI design is that XRI syntax be fully abstract (i.e.,
265 independent of resource location, network, application, transport protocol, type, or security
266 method). Although XRI syntax may be extended for specific uses, the generic XRI syntax is

267 designed to represent logical associations between resources and therefore to be portable across
268 all networks, directories, domains, and applications.

269 **1.2.2 Persistence and Reassignability**

270 XRI syntax and resolution is designed to express and resolve fully persistent identifiers, fully
271 reassignable identifiers, or any combination of persistent and reassignable identifier segments.

272 **1.2.3 Human-Friendliness and Machine-Friendliness**

273 XRI syntax and resolution is designed to support both human-friendly identifiers (HFIs—those
274 optimized for human readability, memorability, and usability) and machine-friendly identifiers
275 (MFIs—those optimized for machine processing and network efficiency). XRI syntax allows any
276 combination of HFI and MFI components within a single XRI.

277 **1.2.4 Internationalization**

278 XRIs are designed to be rendered in the natural language of the intended user. They therefore
279 employ the Unicode character set [**Unicode**] and provide syntactical support for expressing
280 optional language-dependent context metadata. As a result, XRIs extend the virtues of human
281 readability, memorability, and usability to non-English speaking audiences.

282 **1.2.5 Cross-Context Identification**

283 XRI syntax and resolution is designed to allow the use of an identifier in the context of another
284 identifier (i.e., for an XRI or a URI to be contained within another XRI). Such embedded identifiers
285 are called *cross-references*, and they are vital to XRI extensibility.

286 **1.2.6 Authority, Delegation, and Federation**

287 XRI syntax and resolution are designed to allow any resource to serve as an identifier authority,
288 and for any authority to delegate to any other authority at any level of the path. Thus XRI design
289 imposes no specific delegation model, network topology, or federation structure.

290 **1.2.7 Security and Privacy**

291 XRI syntax and resolution is designed to be adapted to any security model, method, or
292 infrastructure, as well as to any privacy policy or framework. XRIs never require sensitive data,
293 such as passwords or account numbers, to be included in an identifier. If a particular application
294 ever needs to include such data in an XRI, the syntax permits encryption and obfuscation of
295 identifier segments for enhanced security and privacy.

296 **1.2.8 Extensibility**

297 The XRI scheme is designed to provide the same interoperable extensibility for identifiers that
298 XML provides for markup languages. In other words, by design, the XRI scheme should be able
299 to be extended and specialized by various identifier authorities, and these extensions and
300 specializations should be interoperable.

301 **1.3 Terminology and Notation**

302 **1.3.1 Keywords**

303 The key words “MUST”, “MUST NOT”, “REQUIRED”, “SHALL”, “SHALL NOT”, “SHOULD”, “SHOULD
304 NOT”, “RECOMMENDED”, “MAY”, and “OPTIONAL” in this document are to be interpreted as

305 described in [RFC2119]. When these words are not capitalized in this document, they are meant
306 in their natural language sense.

307 1.3.2 Syntax Notation

308 This specification uses the syntax notation employed in [RFC2396]: Augmented Backus-Naur
309 Form (ABNF), defined in [RFC2234]. Although the ABNF defines syntax in terms of the US-ASCII
310 character encoding, XRI syntax should be interpreted in terms of the character that the ASCII-
311 encoded octet represents, rather than the octet encoding itself, as explained in [RFC2396]. As
312 with URIs, the precise bit-and-byte representation of an XRI on the wire or in a document is
313 dependent upon the character encoding of the protocol used to transport it, or the character set of
314 the document that contains it.

315 The following core ABNF productions are used by this specification as defined by section 6.1 of
316 [RFC2234]: ALPHA, CR, CTL, DIGIT, DQUOTE, HEXDIG, LF, OCTET, and SP. The complete
317 XRI ABNF syntax is collected in Appendix A.

318 To simplify comparison between generic XRI syntax and generic URI syntax, the ABNF
319 productions that are unique to XRIs are shown with light green shading, while those inherited
320 from [RFC2396] or [RFC2396bis] are shown with light yellow shading.

321

322 This is an example of ABNF specific to XRI.

323

324 This is an example of generic URI ABNF from RFC 2396 or 2396bis.

325

326 In addition, productions inherited from the IRI proposal [IRI] are prefixed with the letter “i” just as
327 they are in that document.

328 1.3.3 Glossary

329 The following definitions are central to this specification.

330

331 **Absolute Identifier**

332 An identifier that refers to a resource independent of the current context, i.e., using a
333 global context. Mutually exclusive with “Relative Identifier.”

334 **Abstract Identifier**

335 An identifier that is not directly resolvable to a resource, but is either:

336 a) a self-reference because it completely represents a non-network resource and is not
337 further resolvable (see “Self-Reference”), or

338 b) an indirect reference to a resource because it must first be resolved to another
339 identifier (either another abstract identifier or a concrete identifier.)

340 A URN as described in [RFC2141] is an example of an abstract identifier. Abstract
341 identifiers provide additional levels of indirection in referencing resources, which can be
342 useful for a variety of purposes, including persistence, equivalence, human-friendliness,
343 and data protection.

344 **Authority (or Identifier Authority)**

345 A resource that assigns identifiers to other resources. Note that in URI syntax as defined
346 in [RFC2396] and [RFC2396bis], the “authority” production refers explicitly to the top-
347 level authority identified by a DNS name or an IP address. Since XRI syntax supports
348 unlimited delegation, the term “authority” can technically refer to an identifier authority at

349 any level. However, in the XRI “authority-path” production (section 2.1.1), it explicitly
350 refers to the top-level identifier authority.

351 **Base Identifier**

352 An absolute identifier that identifies the current context for a relative identifier. See
353 “Relative Identifier.”

354 **Canonical Form**

355 The state of an identifier after applying transformation rules for the purpose of
356 determining equivalence. See also “Normal Form”.

357 **Community (or Identifier Community)**

358 The set of resources that share a common identifier authority, often (but not always) a
359 common root authority. Technically, the set of resources whose identifiers form a directed
360 graph or tree.



361 **Concrete Identifier**

362 An identifier that can be directly resolved to a resource or resource representation, rather
363 than to another identifier. Examples include the MAC address of a networked computer, a
364 phone number that rings directly to a specific device, and a postal address that is not a
365 forwarding address. All concrete identifiers are intended to be resolvable identifiers.
366 Contrast with “Abstract Identifier.”



367 **Context (or Identifier Context)**

368 The resource of which an identifier is an attribute. For example, in the string of delegated
369 identifiers “a/b/c”, the context of the identifier “b” is “a/”, and the context of the identifier “c”
370 is “a/b/”. Since multiple resources may assign an identifier for a target resource, the
371 resource can be said to be identified in multiple contexts. For absolute identifiers, the
372 context is global, i.e., there is a known starting point. For relative identifiers, the context is
373 implicit.

374 **Cross-reference**

375 An identifier assigned in one context that is reused in another context. Cross-references
376 are used primarily to identify logically equivalent resources in different domains or
377 physical locations. For example, a cross-reference may be used to identify the same
378 logical invoice stored in two accounting systems (the originating system and the receiving
379 system), the same logical Web page stored on multiple proxy servers, the same datatype
380 used in multiple databases or XML schemas, or the same abstract concept used in
381 multiple taxonomies or ontologies.

382 **Delegated Identifier**

383 A multi-segment identifier in which some segments are assigned by different identifier
384 authorities. Mutually exclusive with “Local Identifier.”

385 **Federated Identifier**

386 A delegated identifier which spans independent identifier authorities. See also “Delegated
387 Identifier.”

388 **Human-Friendly Identifier (HFI)**

389 An identifier containing words or phrases intended to convey meaning in a specific
390 human language and thus be easy for people to remember and use. Compare with
391 “Machine-Friendly Identifier.”

392 **Identifier**

393 Per **[RFC2396bis]**, anything that “embodies the information required to distinguish what
394 is being identified from all other things within its scope of identification.” In UML terms, an
395 identifier is an attribute of a resource (the identifier context) that forms an association with



396 another resource (the identifier target). The general term “identifier” does not specify
397 whether the identifier is abstract or concrete, persistent or reassignable, human-friendly
398 or machine-friendly, absolute or relative, resolvable or self-referential, or delegated or
399 local.

400 **Local Identifier**

401 Any identifier, or any set of segments in a multi-segment identifier, that is assigned by the
402 same identifier authority. Each of these segments is “local” to that authority. Mutually
403 exclusive with “Delegated Identifier.”

404 **Machine-Friendly Identifier (MFI)**

405 An identifier containing digits, hex values, or other character sequences optimized for
406 efficient machine searching, routing, caching, and resolvability. MFIs generally do not
407 contain human semantics. Compare with “Human-Friendly Identifier.”

408 **Normal Form**

409 The character-by-character format of an identifier after encoding, escaping, or other
410 character transformation rules have been applied in order to satisfy syntactic
411 requirements. Four normal forms are defined for XRIs—escaped normal form, IRI normal
412 form, anyURI normal form, and URI normal form. See section 2.2.4 for details. See also
413 “Canonical Form”.



414 **Persistent Identifier**

415 An identifier that is permanently assigned to a resource and is intended never to be
416 reassigned to another resource, even if the original resource goes off the network, is
417 terminated, or no longer exists. A URN as described in **[RFC2141]** is a persistent
418 identifier. Persistent identifiers tend to be machine-friendly identifiers, since human-
419 friendly identifiers typically reflect human semantic relationships that may change over
420 time. Mutually exclusive with “Reassignable Identifier.”

421 **Reassignable Identifier**

422 An identifier that may be reassigned from one resource to another. Example: the domain
423 name “example.com” may be reassigned from ABC Company to XYZ Company, or the
424 email address “john@example.com” may be reassigned from John Smith to John Jones.
425 Reassignable identifiers tend to be human-friendly identifiers because they often
426 represent the potentially transitory mapping of human semantic relationships onto
427 network resources or resource representations. Mutually exclusive with “Persistent
428 Identifier.”



429 **Relative Identifier**

430 An identifier that refers to a resource only in relationship to the current context (for
431 example, the current community, the current document, or the current position in a
432 delegated identifier). A relative identifier can be converted into an absolute identifier by
433 combining it with a base identifier (an absolute identifier that identifies the current context
434 of the relative identifier.) See “Base Identifier”. Mutually exclusive with “Absolute
435 Identifier.”

436 **Resolvable Identifier**

437 An identifier that references a network resource or resource representation and that can
438 be resolved into a network endpoint for communicating with the target resource. Mutually
439 exclusive with “Self-Reference.”

440 **Resource**

441 Per **[RFC2396bis]**, “ thing that can be named or described.” Resources are of two
442 types: network resource  (those that are network addressable) and non-network
443 resources (those that exist entirely independent of a network). Network resources may be
444 either direct resources or resource representations (see “Resource Representation”).

445 **Resource Representation**

446 A network resource that represents the attributes of another resource. A resource
447 representation may represent either another network resource (such as a machine or an
448 application) or a non-network resource (such as a person, organization, or concept).

449 **Segment**

450 Any syntactically delimited portion of an identifier. In generic URI syntax, all segments
451 after the authority portion are delimited by forward slashes (“/segment1/segment2/...”). In
452 XRI syntax, slash segments can be further subdivided into sub-segments called *dot*
453 *segments* (for reassignable identifiers) and *colon segments* (for persistent identifiers).
454 See section 2.1.2.

455 **Self-Reference (or Self-Referential Identifier)**

456 An identifier which is itself the representation of the resource it references. Self-
457 references are typically used to represent abstract non-network resources (e.g., “love”,
458 “Paris”, “the planet Jupiter”) in contexts where they are not intended to be resolved to a
459 separate network representation of that resource. The primary purpose of self-references
460 is to establish equivalence across contexts (see “Cross-References”). Mutually exclusive
461 with “Resolvable Identifier.”



462 **Target (or Identifier Target)**

463 The resource referenced by an identifier. A target may be either a network resource
464 (including a resource representation) or a non-network resource.

465 **XRI Reference**

466 A term that includes both absolute and relative XRIs. Used the same way as “IRI
467 reference” and “IRI reference”. Note that to transform an XRI reference into an XRI, it
468 must be converted into its absolute form.



469

2 Syntax

470

2.1 Syntax Components

471

Generic XRI syntax builds on generic URI syntax. Since it includes syntactic elements and characters outside the range allowed by [RFC2396], however, this specification does not technically define a new URI scheme. Instead, it follows the example of [IRI] and defines a new identifier scheme, along with a specification for transforming XRIs into IRIs or generic URIs for applications that expect them (see section 2.2.4.3).

476

As with URIs, an XRI may be either absolute or relative.

477

478

```
XRI = absolute-xri / relative-xri
```

479

480

An absolute XRI consists of the scheme name “xri:” followed by the same set of hierarchical components as an absolute URI – authority, path, query, and fragment. For convenience in the ABNF, these components are broken into the *authority path*, the *local path*, and the *query-frag* (a query segment, a fragment segment, or both).

484

485

```
absolute-xri = "xri:" global-path
global-path  = authority-path [ local-path ] [ query-frag ]
local-path   = "/" relative-path
relative-path = *( [ "." ] "/" ) xri-segments
query-frag   = [ "?" xri-query ] [ "#" xri-fragment ]
```

486

487

488

489

490

491

A relative XRI consists of the same set of components as a relative URI.

492

493

```
relative-xri = ( local-path / relative-path ) [ query-frag ]
```

494

495

Finally, in certain contexts such as cross-references (section 2.1.1.4), the “xri:” scheme name is redundant. These contexts can use the *xri-value* production, which includes all levels of XRI paths.

496

497

498

499

```
xri-value = [ global-path / local-path / relative-path ]
```

500

```
[ query-frag ]
```

501

502

2.1.1 Authority

503

XRI syntax supports the same types of authorities as generic URI syntax, called *URI authorities*.

504

In addition, it supports *XRI authorities* that provide two other mechanisms for specifying the global context of an identifier, as defined in section 2.1.1.2.

505

506

507

```
authority-path = URI-authority / XRI-authority
```

508

509

2.1.1.1 URI Authority

510

In the context of an XRI, a URI authority is distinguished by an initial double slash (“//”).

511

512 URI-authority = "//" [userinfo "@"] host [":" port]

513

514 The syntax following this starting delimiter is inherited directly from [RFC2396bis], which
515 simplifies the syntax in [RFC2396] and includes support for IPv6 addresses defined in
516 [RFC2732]. First, the "userinfo" sub-component permits identifying a user in the context of a host.

517

518 userinfo = *(unreserved / escaped / ";" /
519 ":" / "&" / "=" / "+" / "\$" / ",")

520

521 Next, the "host" sub-component has three options for identifying the host: a domain name, an
522 IPv4 address, or an IPv6 literal.

523

524 host = [hostname / IPv4address / IPv6reference]

525

526 Note that the host identifier may be omitted. This is because in generic URI syntax, a default may
527 be defined by the semantics of a particular URI scheme. No default is specified for the XRI
528 scheme; this allows a default to be inherited from the particular protocol used to resolve the XRI.

529 A hostname, after the transformation described in step 4 of section 2.2.4.3, MUST meet the rules
530 defined in section 3.2.2 of [RFC2396]. The productions for idomainlabel, qualified, and hostname,
531 therefore, have additional restrictions not reflected in the ABNF.

532

533 hostname = idomainlabel qualified
534 qualified = *("." idomainlabel) ["."]
535 idomainlabel = 1*ucschar
536 domainlabel = alphanum [0*61(alphanum / "-") alphanum]
537 alphanum = ALPHA / DIGIT



538

539 IPv4address = dec-octet "." dec-octet "." dec-octet "." dec-octet
540 dec-octet = DIGIT ; 0-9
541 / %x31-39 DIGIT ; 10-99
542 / "1" 2DIGIT ; 100-199
543 / "2" %x30-34 DIGIT ; 200-249
544 / "25" %x30-35 ; 250-255

545

546 Support for an IPv6 address literal was added by [RFC2396bis] following the syntax originally
547 specified in [RFC2732]. Because IPv6 literals use colons as delimiters, they must be
548 encapsulated within square brackets.

549

550 IPv6reference = "[" IPv6address "]"
551 IPv6address = 6(h4 ":") ls32
552 / "::" 5(h4 ":") ls32
553 / [h4] "::" 4(h4 ":") ls32
554 / [*1(h4 ":") h4] "::" 3(h4 ":") ls32
555 / [*2(h4 ":") h4] "::" 2(h4 ":") ls32
556 / [*3(h4 ":") h4] "::" h4 ":" ls32
557 / [*4(h4 ":") h4] "::" ls32
558 / [*5(h4 ":") h4] "::" h4
559 / [*6(h4 ":") h4] "::"
560 ls32 = (h4 ":" h4) / IPv4address
561 ; least-significant 32 bits of address
562 h4 = 1*4HEXDIG

563

564 Finally, a host identifier can be followed by an optional port number. Because XRIs are abstract
565 identifiers, the XRI syntax specification does not define a default port. It is expected that the
566 default port will be inherited from the resolution protocol, such as the HTTP/HTTPS protocol
567 specified in section 3. Therefore, if the port is omitted in an XRI, it is undefined.

568

569 `port = *DIGIT`

570

571 2.1.1.2 XRI Authority

572 In addition to the authorities supported in generic URI syntax, XRIs support two other
573 mechanisms for specifying the global context of an identifier. The first is the global context symbol
574 (GCS), and the second is the cross-reference (abbreviated in the ABNF as *xref*).

575

576 `XRI-authority = (gcs-char xri-segment) / xref-authority`

577

578 2.1.1.3 Global Context Symbols (GCS)

579 To support the abstraction and human-friendly identifier (HFI) requirements, XRIs offer a simple,
580 compact syntax for indicating the logical global context of an identifier: a single prefix character.

581

582 `gcs-char = "+" / "=" / "@" / "$" / "*" / "!"`

583

584 The global context symbol characters were selected from the set of symbol characters that are
585 valid in a URI under [RFC2396] to represent the global contexts shown in Table 2:
586

| Symbol Character | Authority Type | Establishes global context for |
|------------------|----------------------------------|---|
| + | General public | Identifiers for generic concepts for which there is no specific authority, i.e., that are established by public convention. (In the English language, for example, these would be the generic nouns.) |
| = | Person | Identifiers that represent an individual person. |
| @ | Organization | Identifiers that represent an organization of any kind. |
| \$ | OASIS XRI Metadata Specification | Special identifiers established by the XRI Metadata Specification for interoperable identifier metadata (e.g., language, version, type, query syntax, etc.). See [tktk – reference needed.] |
| * | User-relative | Identifiers for which the authority is relative to the current user (“user-shortcut XRIs”). See section 3.2.6. |
| ! | XRI author | Identifiers used only for human-readable annotations of XRIs (ignored by machine processing.) |

587



Table 2: XRI global context symbols.

588 Note that because a global context symbol may precede an xri-segment, and an xri-segment may
589 start with a cross-reference (below), a global context symbol can be used to express the abstract
590 logical context of a conventional URI authority. For example:
591

```
592 xri:=(http://www.my-website.com)/favorites.html  
593 --expresses that this resource represents an individual
```

594

595 2.1.1.4 Cross-References

596 Cross-references are the primary extensibility mechanism in XRI. A cross-reference may be
597 either an XRI value or an absolute URI. In either case, it is enclosed in parentheses the same
598 way an IPv6 literal is encapsulated in square brackets as specified in [RFC2732] (see section
599 2.1.1.1).

600

```
601 xref-authority = xref ( "." sub-segment / ":" sub-segment ) * ( "."  
602 sub-segment / ":" sub-segment )  
603 xref = "(" ( xri-value / URI ) ")"
```

604

605 It is important that the value of a cross-reference be syntactically unambiguous, whether it is an
606 absolute URI or one of the various forms of an XRI value. Since an absolute URI must start with a
607 legal URI scheme name character (i.e., an ALPHA), an XRI value used as a cross-reference
608 MUST start with a symbol character. Since the only XRI value that is not required to start with a
609 symbol is a dot segment (see section 2.1.2), the effect of this rule is that if a relative XRI begins
610 with a dot segment, the leading dot is not optional. For example, if the relative XRI “foo/bar” is
611 used as a cross-reference, it must include the optional leading dot, i.e., “(.foo/bar)”.

612 A cross-reference may appear at any node of any XRI except within a URI authority segment.
613 The use of cross-references as the very first segment in an XRI enables any globally-unique
614 identifier in any URI scheme (e.g., an HTTP URI, mailto URI, URN, etc.) to specify a global
615 authority.
616

```
617 xri:(mailto:john.doe@example.com)/favorites/home  
618 --example of using a URI as an XRI global authority
```

619



620 2.1.1.5 Self-References

621 Cross-reference syntax is also the means by which an XRI can express that it is not intended for
622 resolution, but only for the purpose of establishing equivalence across contexts. Such an XRI is
623 called a *self-reference*. To express a self-reference, the entire XRI value is enclosed in
624 parentheses—in essence, it becomes a global cross-reference. This is the XRI equivalent of the
625 English language convention of putting a word or phrase in quotes to express that the author is
626 referring to the word or phrase itself and not to its normal meaning. (In linguistics and philosophy,
627 this is called the “use-mention distinction.”) For example:

628

```
629 The term "user-friendly" is used frequently in computing.  
630 --English-language usage of a quoted term  
631  
632 xri:(+user-friendly)  
633 --XRI syntax for expressing a self-reference
```

634

635 2.1.2 Path

636 As with URIs, the XRI path component is a hierarchal sequence of path segments separated by
637 slash ("/) characters and terminated by the first question-mark ("?) or number sign ("#")
638 character, or by the end of the XRI. The key difference is that while a URI path segment is
639 considered opaque by a generic URI processor, an XRI path segment can be parsed by an XRI
640 processor into two types of sub-segments: *dot segments* and *colon segments* after their leading
641 characters ("." and ":").

642

```
643 xri-segments = xri-segment * ( "/" xri-segment )  
644 xri-segment  = ( [ "." ] sub-segment / ":" sub-segment )  
645              *( "." sub-segment / ":" sub-segment )  
646 sub-segment  = *xri-pchar / xref
```

647

648 Dot segments are used to specify *reassignable identifiers*—identifiers that may be reassigned by
649 an identifier authority to represent a different resource at some future date. Colon segments
650 (following the lead of URN syntax in **[RFC2141]**) are used to specify *persistent identifiers*—
651 identifiers that are permanently assigned to a resource and will not be reassigned at a future
652 date. The default is a dot segment, so no leading dot is required if this is the first (or only) sub-
653 segment.

654 Other than these special uses of the dot (".") and the colon (":") characters, an XRI path segment
655 can contain the same characters as a URI path segment. If a dot or colon is used, it will be
656 interpreted as a delimiter. If this interpretation is not desired for these characters, or for any other
657 special XRI delimiters, these characters **MUST** be escaped when they appear in the path
658 segment. See section 2.2.4, "Escaped Characters".

659

```
660 xri-pchar = xri-unreserved / escaped / ";" / "!" / "*" /  
661           "@" / "&" / "=" / "+" / "$" / ","
```

662

663 With the exception of dot and colon sub-segments, an XRI path segment is considered opaque
664 by generic XRI syntax. As with URIs in general, XRI extensions or generating applications may
665 define special meanings for other URI reserved characters for the purpose of delimiting
666 extension-specific or generator-specific sub-components. For example, section 3.4 of **[RFC2396]**
667 specifies the set of URI reserved characters that can be used within a query segment.

668 2.1.3 Query

669 The XRI query component is identical to the URI query component as described in section 3.4 of
670 **[RFC2396]**, except that it may begin with a cross-reference. This permits the incorporation of XRI
671 metadata describing the query string syntax. See the XRI Metadata Specification [tktk need
672 reference] for more about query syntax metadata.

673

```
674 xri-query = [ xref ] * ( pchar / "/" / "?" )
```

675

676 The characters permitted in a query segment are the same ones allowed in a URI query segment.

677

```
678 pchar = unreserved / escaped / ";" /  
679        ":" / "@" / "&" / "=" / "+" / "$" / ","
```

680 2.1.4 Fragment

681 XRI syntax also supports fragments as described in section 4.1 of [RFC2396], except that an XRI
682 fragment may begin with a cross-reference.

683

```
684 xri-fragment = [ xref ] * ( pchar / "/" / "?" )
```

685

686 Since XRI syntax can directly address attributes or secondary representations of a primary
687 resource to any depth, fragments are supported primarily for compatibility with generic URI
688 syntax. XRIs can also employ cross-references to identify media types or other alternative
689 representations of a resource.

690 2.2 Characters

691 The character set and encoding of an XRI is primarily inherited from generic URI syntax as
692 defined in [RFC2396] and clarified in [RFC2396bis]. However, it also includes the expanded
693 character set defined in [IRI].

694 All XRI characters fall into the same three subsets as URI characters.

695

```
696 xri-characters = xri-reserved / xri-unreserved / escaped
```

697 2.2.1 Character Encoding

698 The basic character encoding of XRI is UTF-8, as recommended by [RFC2718]. When an XRI is
699 presented as a human readable identifier, the representation of the XRI in the underlying
700 document should use the character encoding of the underlying document. However, this string
701 must be converted to UTF-8 before any processing external to the underlying document.

702 Note that not all ASCII sequences can be derived from UTF-8 sequences. A valid XRI character
703 sequence MUST be derivable by unescaping an equivalent UTF-8 sequence. For example, the
704 ASCII sequence '%FC', which would represent U+00FC LATIN SMALL LETTER U WITH
705 DIAERESIS in an iso-8859-1 encoding, when unescaped will not result in a valid UTF-8
706 sequence.

707 2.2.2 Reserved Characters

708 Because additional characters are used to delimit XRI syntax components not present in URIs,
709 the XRI reserved character set is a superset of the URI reserved character set. Specifically, five
710 characters have been added: opening parenthesis ("("), closing parenthesis (")"), dot ("."), asterisk
711 ("*"), and exclamation point ("!").

712

```
713 xri-reserved = "/" / "?" / "#" / "[" / "]" / "(" / ")" / ";" / ":" /  
714 ", " / "." / "&" / "@" / "=" / "+" / "*" / "$" / "!"
```

715

716 If the use of an unescaped XRI reserved character as a data character would cause the
717 interpretation of the XRI to be ambiguous, the character MUST be escaped as per the rules in
718 section 2.2.4, "Escaped Characters", and particularly section 2.2.4.4.

719 2.2.3 Unreserved Characters

720 Aside from the expanded UCS character set for internationalization, the unreserved character set
721 for XRIs is the same as that of URIs after the subtraction of the five characters noted above (all of
722 which are in of the "mark" production of [RFC2396] and [RFC2396bis]).

723

724 xri-unreserved = ALPHA / DIGIT / ucschar / xri-mark
725 xri-mark = "-" / "_" / "~" / "'"

726

727 The principal difference between XRI and URI reserved character sets is the inclusion of the UCS
728 character set.

729

730 ucschar = %xA0-D7FF / %xF900-FDCF / %xFDF0-FFEF /
731 %x10000-1FFFFD / %x20000-2FFFFD / %x30000-3FFFFD /
732 %x40000-4FFFFD / %x50000-5FFFFD / %x60000-6FFFFD /
733 %x70000-7FFFFD / %x80000-8FFFFD / %x90000-9FFFFD /
734 %xA0000-AFFFFD / %xB0000-BFFFFD / %xC0000-CFFFFD /
735 %xD0000-DFFFFD / %xE1000-EFFFFD

736

737 Escaping unreserved characters in an XRI does not impact what resource is identified by that
738 XRI. However, it may change the result of an XRI comparison (see section 2.4, “Normalization
739 and Comparison”), so unreserved characters should not be escaped unless necessary.

740 2.2.4 Escaped Characters

741 XRIs follow the same rules for escaping characters as URIs. That is, any data in an XRI MUST be
742 escaped if: a) it does not have a representation using an unreserved character, and b) using a
743 reserved character could cause the XRI to be misinterpreted.

744 An XRI thus escaped is said to be in *escaped normal form*. This does not imply that it is
745 necessarily a valid IRI or URI. Rules for converting an XRI into a valid IRI or URI are discussed in
746 section 2.2.4.3. An XRI is in escaped normal form if it is unambiguous per the ABNF provided in
747 this document, but it is a valid IRI or URI only after it is escaped according to the transformation
748 described in section 2.2.4.3.

749 2.2.4.1 Escaped Encoding

750 XRIs use the same percent-encoding as URIs, described in section 2.4.1 of [RFC2396]. An
751 escaped octet is encoded as a character triplet consisting of the percent character “%” followed
752 by the two hexadecimal digits representing that octet’s numerical value.

753

754 escaped = "%" HEXDIG HEXDIG

755

756 The uppercase hexadecimal digits “A” through “F” are equivalent to the lowercase digits “a”
757 through “f”, respectively. XRIs that differ only in the case of hexadecimal digits used in escaped
758 octets are equivalent. For consistency, uppercase digits SHOULD be used by XRI generators and
759 normalizers.

760 Note that the % symbol used by itself in an XRI must be escaped as described in section 2.2.5.

761 2.2.4.2 Encoding XRI Metadata

762 In some cases, the transformation from an identifier in its native language and display format into
763 an XRI in escaped normal form may lose information that cannot be retained through character
764 escaping. For example, in certain languages displaying the glyph of a UTF-8 encoded character
765 requires additional language and font information not available in UTF-8. The loss of this
766 information during UTF-8 encoding can cause the resulting XRI to be ambiguous.

767 Another case is when the normalization or canonicalization rules of a particular identifier authority
768 do not permit the inclusion of whitespace, mixed case letters, or certain punctuation in an XRI
769 segment even when escaped, yet the authority would like to retain this metadata for purposes of
770 presentation. XRI syntax offers an option for encoding this metadata using a cross-reference

771 beginning with the GCS “\$” symbol. As defined in section 2.1.1.3, the top level authority for these
772 identifiers is the XRI Metadata Specification [tktk – need reference]. It defines special identifiers
773 for UTF-8 metadata, presentation metadata, and other standard types of identifier metadata
774 together with the rules governing their interpretation.

775 2.2.4.3 Transforming XRIs into IRIs and URIs

776 Although XRIs are intended to be used by applications that understand them natively, it may also
777 be desirable to use them:


- 778 • In contexts that expect a fully-conformant URI reference as defined by **[RFC2396]**.
- 779 • In contexts where there is already a predefined escaping procedure for characters that would
780 otherwise be illegal in a URI under **[RFC2396]**, for example the “anyURI” datatype defined in
781 **[XMLSchema2]**.
- 782 • In contexts where it is desirable to use an Internationalized Resource Identifier as described
783 in **[IRI]**. Note that while **[IRI]** defines the process for converting an IRI to a URI, this
784 conversion differs slightly from the conversion defined for “anyURI” in **[XMLSchema2]** in that
785 it includes an algorithm appropriate for internationalized domain names.

786 This section specifies a progression of steps for transforming an XRI into:

- 787 • A valid IRI (steps 1 – 3 below),
- 788 • A valid anyURI (steps 1 – 4 below), and
- 789 • A valid generic URI (steps 1 – 5 below).

790 Except for transformations specific to XRI syntax, these steps closely follow the algorithm
791 proposed in **[IRI]**.

792 Applications MUST transform XRIs to IRIs, anyURIs, or generic URIs using the following steps (or
793 an equivalent process that achieves exactly the same result). These steps assume that the XRI is
794 already in escaped normal form as defined in section 2.2.4.

- 795 1. If the XRI is not encoded in UTF-8, convert the XRI to a sequence of characters encoded
796 in UTF-8, normalized according to Normalization Form C (NFC) as defined in **[UTR15]**.
- 797 2. If necessary, add XRI metadata using cross-references as defined in section 2.2.4.2.
798 Note that the addition of XRI metadata may change the resulting IRI or URI for the
799 purposes of comparison. The significance or insignificance of specific types of XRI
800 metadata is defined in the XRI Metadata Specification [tktk – need reference].
- 801 3. Perform the XRI-specific conversion defined in section 2.2.4.4. Note that this step is not
802 idempotent (i.e., each time this step is applied, it may yield different results), so it is very
803 important that implementers not apply this step more than once to avoid changing the
804 semantics of the identifier. At the completion of this step, the escaped XRI may be used
805 as an IRI. This is referred to as *IRI normal form*.
- 806 4. If the XRI has a “hostname” component, replace it with the “hostname” component
807 converted using the “ToASCII” operation defined in section 4.1 of **[RFC3490]**, with the
808  “STD3ASCIIRules” flag set to true and the “AllowUnassigned” flag set to false. At this
809 point the XRI may be used as an anyURI as defined in **[XMLSchema2]** or in a
810 comparable context. This is referred to as *anyURI normal form*.
- 811 5. Replace each character that is disallowed in URI references with escaped triplet(s) as
812 described in section 2.2.4.1, one escaped triplet for each octet in the UTF-8 encoding of
813 the disallowed character. At this point the XRI may be used as a generic URI. This is
814 referred to as *URI normal form*.

815 The form of the XRI that results from each step in this transformation is equivalent to the result of
816 any other step. Applying this conversion does not change the equivalence of the identifier, with
817 the exception of language or font metadata additions as discussed in Step 2.

818 In general, an application SHOULD use the least escaped version appropriate for the context in
819 which the identifier appears. For example, if the context allows an XRI directly, the identifier
820 SHOULD be an XRI in escaped normal form as described in section 2.2.4. If the context allows
821 an IRI but not an XRI, the identifier SHOULD be in IRI normal form, and so on.

822 2.2.4.4 Special Escaping Rules for XRI Syntax

823 This section defines special rules for preventing misinterpretation of XRI syntax when an XRI is
824 evaluated by a non-XRI aware processor.

825 The first rule deals with cross-references as explained in section 2.1.1.4. Since a cross-reference
826 contains either a URI or an XRI value (which itself may contain further nested URIs or XRIs), it
827 may include characters that, if not escaped, would cause misinterpretation when the entire XRI is
828 transformed according to the steps in section 2.2.4.3. Consider the following XRI:

829

```
830 xri:@example/(xri:@example2/abc?id=1)
```

831

832 The generic parsing algorithm described in [RFC2396] would separate the above XRI into the
833 following components:

834

```
835 scheme = xri  
836 authority = <undefined>  
837 path = @example/(xri:@example2/abc  
838 query = id=1)
```

839

840 The desired separation is:

841

```
842 scheme = xri  
843 authority = <undefined>  
844 path = @example/(xri:@example2?id=1)  
845 query = <undefined>
```

846

847 To avoid this type of misinterpretation, certain characters in a cross-reference must be escaped
848 when transforming an XRI into IRI, anyURI, or URI normal form. In particular, the question mark
849 “?” character must be escaped as “%3F” and the number sign “#” character must be escaped as
850 “%28”.

851 Following this rule, the above example would be expressed as:

852

```
853 xri:@example/(xri:@example2%3Fid=1)
```

854

855 In addition, the slash “/” character in a cross-reference may also be misinterpreted by a non-XRI
856 aware processor. Consider:

857

```
858 xri://example.com/(@example/abc)
```

859

860 If this were used as a base URI as defined in section 5 of [RFC2396], the algorithm described in
861 section 5.2 of [RFC2396] would append a relative-path reference to:

862

```
863 xri://example.com/(@example/
```

864

865 instead of the intended:

866

```
867 xri://example.com/
```

868

869 This is because the algorithm is defined in terms of the last (right-most) slash character. This
870 problem is avoided by escaping slashes within cross-references as "%2F". Following this rule, the
871 above example would now be expressed as:

872

```
873 xri://example.com/(@example%2Fabc)
```

874

875 Ambiguity is also possible if an XRI in escaped normal form contains characters that have been
876 escaped to indicate that they should not be interpreted in their normal syntactical sense. For
877 example, consider the following XRI in escaped normal form:

878

```
879 xri://example.com/(@example/abc%2Fd/ef)
```

880

881 This slash character between "c" and "d" is escaped to show that it's not a syntactical element of
882 the XRI, i.e., that it should be interpreted literally and not as a path separator. To preserve this
883 type of distinction when converting an XRI to an IRI or URI, the percent "%" character must be
884 escaped as "%25". Following this rule, the above example fully converted would be:

885

```
886 xri://example.com/(@example%2Fabc%252Fd%2Fef)
```

887

888 To summarize, the four special escaping rules below **MUST** be applied during Step 3 of section
889 2.2.4.3. Before applying these rules, the XRI **MUST** be in escaped normal form and all URIs in
890 cross-references **MUST** be in an escaped form appropriate to their schemes.

- 891 1. Escape all percent "%" characters as "%25" across the entire XRI.
- 892 2. Escape all number sign "#" characters that appear within a cross-reference as "%23".
- 893 3. Escape all question mark "?" characters that appear within a cross-reference as "%3F".
- 894 4. Escape all slash "/" characters that appear within a cross-reference as "%2F".

895

896 2.2.4.5 Transforming URIs and IRIs Back into XRIs

897 Transformation of an XRI in IRI, anyURI, or URI normal form into an XRI in escaped normal form
898 **MUST** use the following steps (or an equivalent process that achieves the same result). Except
899 for the steps specific to XRI syntax, this procedure very closely follows the algorithm defined in
900 **[IRI]**.

901 If the XRI is in URI normal form, perform this sequence of steps:

- 902 1. If the identifier is not encoded in US-ASCII, convert it to a sequence of octets in US-
903 ASCII.
- 904 2. If the identifier has a "hostname" component, replace it with the UTF-8 encoded
905 "hostname" component converted using the "ToUnicode" operation defined in section 4.2
906 of **[RFC3490]**, with the "UseSTD3ASCIIRules" flag set to true and the
907 "AllowedUnassigned" flag set to false.
- 908 3. Convert all escaped characters (as defined in section 2.2.4) with their corresponding
909 octets, except for the percent "%" character, those characters in the "reserved"
910 production of **[RFC2396]** and US-ASCII characters disallowed in URIs by section 2.4.3 of
911 **[RFC2396]**.

- 912 4. Re-escape any octet produced in step 3 that is not part of a strictly legal UTF-8 octet
 913 sequence.
- 914 5. Perform the following special conversions for XRI syntax:
- 915 a. Convert all escaped slash "/" characters to their corresponding octets.
 - 916 b. Convert all escaped question mark "?" characters to their corresponding octets.
 - 917 c. Convert all escaped number sign "#" characters to their corresponding octets.
 - 918 d. Convert all escaped percent "%" characters to their corresponding octets.
- 919 6. Encode the resulting sequence in UTF-8 (except for that portion already converted by
 920 step 3).

921 If the XRI is in anyURI normal form, perform this sequence of steps:

- 922 1. If the XRI is not encoded in UTF-8, convert the XRI to a sequence of characters encoded
 923 in UTF-8, normalized according to Normalization Form C (NFC) as defined in [UTR15].
- 924 2. Perform Step 2 above.
- 925 3. Perform Step 5 above.

926 If the XRI is in IRI normal form, perform the same steps as with an XRI in anyURI normal form,
 927 except skip step 2.

928

929 2.2.5 Excluded Characters

930 XRI syntax excludes the same characters as URI syntax for the same reasons as described in
 931 section 2.5 of [RFC2396] and [RFC2396bis]. Data octets corresponding to these characters
 932 MUST be escaped in order to be represented within an XRI.

933

```

934 excluded = invisible / delims / unwise
935 invisible = CTL / SP / %x80-FF
936 delims = "<" / ">" / "%" / DQUOTE
937 unwise = "{" / "}" / "|" / "\" / "^" / "`"
938
```

939 As with IRIs, infrastructure responsible for accepting or presenting XRIs MAY deal with
 940 characters in the "excluded" set above, escaping them on input and/or unescaping them prior to
 941 rendering as described in section 2.2.4. A string that contains these characters in an unescaped
 942 form, however, is not technically a valid XRI.

943 Note that in certain contexts, presenting "space" or other whitespace characters in unescaped
 944 form may present special risks for several reasons. First, it is often difficult to visually determine
 945 the number of spaces or other characters composing a block of whitespace, leading to
 946 transcription errors. Second, the space character is often used to delimit an XRI, so unescaped
 947 spaces or whitespace characters can make it difficult or impossible to determine where the
 948 identifier ends. Finally, unescaped spaces or whitespace can be used to maliciously construct
 949 subtly different identifiers intended to mislead the reader. For these reasons, unescaped spaces
 950 or whitespace characters SHOULD be avoided in presentation.

951 [IRI] provides the following guidance concerning other characters that should be avoided. This
 952 guidance applies to XRIs as well.

953 The UCS also contains many areas of characters for which there are strong
 954 visual look-alikes. Because of the likelihood of transcription errors, these also
 955 should be avoided in IRIs. These include the full-width equivalents of ASCII
 956 characters, half-width Katakana characters for Japanese, and many others. This
 957 also includes many look-alikes of "space", "delims", and "unwise", characters
 958 excluded in [RFC3491].

959 Additional information is available from [UniXML]. [UniXML] is written in the
960 context of running text rather than in the context of identifiers. Nevertheless, it
961 discusses many of the categories of characters not appropriate for IRIs.

962

963 **2.3 Relative XRI References**

964 The authority component of an XRI may be either a URI-authority (section 2.1.1.1) or an XRI-
965 authority (section 2.1.1.2). In this section, “authority” should be understood as defined by section
966 2.1.1 of this specification and not in the narrower sense of section 3.2 of [RFC2396].

967 For a relative XRI reference whose base XRI contains an authority component matching the URI-
968 authority production, the rules for resolving relative references defined in section 5.2 of
969 [RFC2396] apply. However, for a relative XRI reference whose base XRI contains an authority
970 component matching the XRI-authority production, the rules defined in section 5.2 of [RFC2396]
971 need modification because an XRI authority is considered opaque by generic URI syntax.

972 The following sections, therefore, define the process for resolving a relative XRI reference into a
973 fully qualified XRI regardless of the type of authority involved.

974 **2.3.1 Establishing a Base XRI**

975 A base XRI is established according to the rules defined in section 5.1 of [RFC2396]. Applying
976 these rules, however, may require the conversion of the XRI into URI normal form as described in
977 section 2.2.4.3. Once in URI normal form, there is no difference between establishing a base XRI
978 and establishing the base of any URI.

979 **2.3.2 Obtaining the Referenced XRI**

980 Section 5.2 of [RFC2396] describes rules for resolving relative references to absolute forms of
981 URIs. For XRIs, these rules apply with the following modifications:

- 982 • In step 1, the XRI reference is parsed using an XRI aware parser such that the “authority”
983 component is interpreted as the “authority-path” production defined in section 2.1.1 of this
984 specification.
- 985 • Step 4 states, “If the ‘authority’ component is defined, then the reference is a ‘network-path’
986 and we skip to step 7”. For XRIs, the presence of an “authority” component does not imply
987 that the reference is a “network-path” as defined by [RFC2396] because it may be an “XRI-
988 authority” component. However, the instruction to skip to step 7 is still valid for XRIs. In other
989 words, the processing instruction is correct, but the inference as to the type of reference is
990 invalid.
- 991 • In step 4, the base XRI is parsed using an XRI-aware parser such that the “authority”
992 component is interpreted as the “authority-path” production defined in section 2.1.1 of this
993 specification.

- 994 • In step 7, the block that reads:

```
995 if authority is defined then  
996     append "/" to result  
997     append authority to result
```

998 is replaced by

```
999 if authority is defined then  
1000     if type-of(authority) == URI-authority  
1001         append "/" to result  
1002         append authority to result
```

1003

1004 It is important to note that the algorithm described in section 5.2 of [RFC2396] will generally
1005 produce incorrect results when applied to relative XRI references where the authority component
1006 matches the “XRI-authority” production. This type of relative XRI reference, therefore, should only
1007 be used in contexts in which the algorithm specified in this section is known to be employed.

1008 The following are examples of resolving relative XRI references. These examples closely follow
1009 the examples for resolving relative references in URIs in appendix C of [RFC2396]. Starting with a
1010 base XRI of:

1011 `xri:@a.b.c/d.e/f;p?q`

1012 the following relative XRIs would be resolved as shown:

| | | | |
|------|----------------------|---|--|
| 1013 | <code>.g:h</code> | = | <code>xri:@a.b.c/d.e/.g:h</code> |
| 1014 | <code>./g:h</code> | = | <code>xri:@a.b.c/d.e/g:h</code> |
| 1015 | <code>g:h</code> | = | <code>g:h</code> (see section 2.3.3 below) |
| 1016 | <code>g</code> | = | <code>xri:@a.b.c/d.e/g</code> |
| 1017 | <code>./g</code> | = | <code>xri:@a.b.c/d.e/g</code> |
| 1018 | <code>g/</code> | = | <code>xri:@a.b.c/d.e/g/</code> |
| 1019 | <code>/g</code> | = | <code>xri:@a.b.c/g</code> |
| 1020 | <code>?y</code> | = | <code>xri:@a.b.c/d.e/?y</code> |
| 1021 | <code>g?y</code> | = | <code>xri:@a.b.c/d.e/g?y</code> |
| 1022 | <code>#s</code> | = | (current document)#s |
| 1023 | <code>g#s</code> | = | <code>xri:@a.b.c/d.e/g#s</code> |
| 1024 | <code>g?y#s</code> | = | <code>xri:@a.b.c/d.e/g?y#s</code> |
| 1025 | <code>;x</code> | = | <code>xri:@a.b.c/d.e;/x</code> |
| 1026 | <code>g;x</code> | = | <code>xri:@a.b.c/d.e/g;x</code> |
| 1027 | <code>g;x?y#s</code> | = | <code>xri:@a.b.c/d.e/g;x?y#s</code> |
| 1028 | <code>.</code> | = | <code>xri:@a.b.c/d.e/</code> |
| 1029 | <code>./</code> | = | <code>xri:@a.b.c/d.e/</code> |
| 1030 | <code>..</code> | = | <code>xri:@a.b.c/</code> |
| 1031 | <code>../</code> | = | <code>xri:@a.b.c/</code> |
| 1032 | <code>../g</code> | = | <code>xri:@a.b.c/g</code> |

1033 As with URIs, the “..” syntax cannot be used to change the authority component of an XRI.

| | | | |
|------|---------------------|---|-----------------------------|
| 1034 | <code>../..</code> | = | <code>xri:@a.b.c/..</code> |
| 1035 | <code>../..</code> | = | <code>xri:@a.b.c/..</code> |
| 1036 | <code>../..g</code> | = | <code>xri:@a.b.c/..g</code> |

1037 2.3.3 Leading Segments Containing a Colon

1038 [RFC2396] points out that relative URI references with an initial segment containing a colon may
1039 be subject to two interpretations:

1040 Authors should be aware that a path segment which contains a colon character
1041 cannot be used as the first segment of a relative URI path (e.g., “this:that”),
1042 because it would be mistaken for a scheme name.

1043 It is therefore necessary to precede such segments with other segments (e.g.,
1044 “./this:that”) in order for them to be referenced as a relative path.

1045 There are cases where relative XRI references would be similarly misinterpreted. Therefore if any
1046 segment prior to the first forward slash (“/”) character in a relative XRI reference contains a colon,
1047 the relative XRI reference must be rewritten to begin either with a “.” or a “./”. Thus, “foo:bar”
1048 becomes “.foo:bar” or “./foo:bar” and “foo.bar:baz” becomes “.foo.bar:baz” or “./foo.bar:baz”. Note
1049 that by the rules of sections 2.3.2 and 2.4.3, this transformation does not affect equivalence.

1050 2.4 Normalization and Comparison

1051 In general, the normalization and comparison rules for generic URIs specified in [RFC2396] apply
1052 to XRIs in URI normal form, namely that the scheme and hostname are case insensitive. This

1053 section describes a number of additional XRI-specific rules for normalization and comparison.
1054 To reduce the requirements imposed upon a minimally conforming processor, the majority of
1055 these rules are RECOMMENDED rather than REQUIRED. An implementation that fails to
1056 observe them, however, may frequently treat two XRIs as non-equal when in fact they are equal.
1057 In addition to these rules, Section 6 of [RFC2396bis] offers advice on more aggressive strategies
1058 for normalization. Although entirely non-normative, implementers may find this information useful
1059 in developing a strategy for establishing equivalence, particularly with respect to XRIs containing
1060 cross-references to URIs.

1061 Finally, each application that uses XRIs MAY define additional equivalence rules as appropriate.
1062 Due to the level of abstraction XRIs provide, such higher-order equivalence rules may be based
1063 on indirect comparisons or specified XRI-to-XRI mappings (for example, mappings of
1064 reassignable XRIs to persistent XRIs).



1065 2.4.1 Case

1066 The following rules regarding case sensitivity SHOULD be applied in XRI comparisons.

- 1067 • Comparison of the scheme component of XRIs and all URIs used as cross-references is
1068 case-insensitive.
- 1069 • Comparison of URI authority components as defined in section 2.1.1.1 is case-insensitive as
1070 defined in [RFC2396].
- 1071 • Comparison of XRI authority components as defined in section 2.1.1.2 is case-insensitive.
1072 Specifically, because an XRI authority component can contain a wide range of Unicode
1073 characters, two XRI authority components are equivalent if they match according to the
1074 compatibility caseless match operation defined in section 3.13 of [Unicode] after applying
1075 steps 1 and 3 of the transformation described in section 2.2.4.3.
- 1076 • As specified in section 2.2.4.1, comparison of percent-encoded characters is case-insensitive
1077 for the hexadecimal digits “A” through “F”.

1078 2.4 Encoding, Escaping, and Transformations

- 1079 • Two XRIs MUST be considered equivalent if they are character-for-character equivalent.
1080 Therefore, they are also equivalent if they are byte-for-byte equivalent and use the same
1081 character encoding.
- 1082 • Two XRIs that differ only in escaped unreserved characters SHOULD be considered
1083 equivalent. If one XRI escapes one or more unreserved characters, and another XRI is
1084 different only in that the same characters are not escaped, they are equivalent.
- 1085 • All forms of an XRI during the transformation process described in section 2.2.4.3 SHOULD
1086 be considered equivalent, assuming the same XRI metadata is inserted as described in
1087 section 2.2.4.2.

1088 2.4.3 Optional Syntax

- 1089 • An xri-segment (section 2.1.2) that omits the optional leading dot (“.”) is equivalent to the
1090 same xri-segment prefixed with the leading dot. For example the segment “/foo.bar” is
1091 equivalent to the segment “./foo.bar”.
- 1092 • 2.10A cross-reference (section 2.1.1.4) that begins with the GCS symbol for annotations (“!”)
1093 AND the delimiter that precedes the cross-reference SHOULD be ignored entirely for
1094 purposes of comparison. For example, “xri:@:A6B4.(!www.example.org):5E32” is equivalent
1095 to “xri:@:A6B4:5E32”. Note that because XRI annotations are explicitly designed to be
1096 ignored by XRI processors, failure to observe this rule will cause XRIs that are intended to be
1097 equivalent to be incorrectly evaluated.



1098

2.4.4 Cross-References

1099

- If an XRI contains a cross-reference, the rules in this section SHOULD be applied recursively to each cross-reference. For example, the following two XRI's should be considered equivalent:

1100

1101

1102

1103

1104

```
xri:@example/(+example/(+foo))
xri:@example/(+Example/(+FOO))
```

1105

1106

- From the standpoint of XRI syntax, all cross-references beginning with the GCS “\$” symbol SHOULD be considered significant unless stated otherwise in the XRI Metadata Specification. See section 2.2.4.2.

1107

1108

1109

1110

2.4.5 Canonicalization

1111

In general, XRIs do not have a single canonical form. This is particularly true for XRIs that contain URI cross-references, since many URI schemes, including the HTTP scheme, do not define a canonical form. Additionally, the authority for a particular segment of an XRI may define its own rules with respect to case-sensitivity, optional or implicit syntax, etc., making canonicalization of those segments outside the scope of this specification.

1112

1113

1114

1115

1116

Nevertheless it is valuable to define guidelines for making XRIs reasonably canonical. XRIs that follow these guidelines will be more consistent in presentation, simpler to process, less prone to false-negative comparisons, and more easily cached. To that end, unless there is a compelling reason to do otherwise, XRIs should be provided them in a form in which:

1117

1118

1119

1120

- optional xri scheme is added,
- The scheme is provided in lowercase,
- The authority component is provided in lowercase,
- Percent-escaping uses uppercase A through F,
- If optional, the leading dot in xri-segments is omitted,
- Unnecessary escaping is removed,
- ./ and ../ are absent in absolute XRIs, and
- Cross-references are reasonably canonical with respect to their schemes.

1121

1122

1123

1124

1125

1126

1127

1128

1129

Table 3 illustrates the application of these rules. Although the XRIs in the first and second columns are equivalent, the form in the second column is recommended.

1130

1131

| Avoid | Recommended | Comment |
|-------------------|------------------|-----------------------------|
| @example | xri:@example | Add optional scheme |
| XRI:@example | xri:@example | Lowercase scheme |
| xri:@Example | xri:@example | Lowercase authority |
| xri:@example%2f | xri:@example%2F | Uppercase percent escaping |
| xri:@example/.abc | xri:@example/abc | Remove optional leading dot |
| xri:@ex%61mple | xri:@example | Remove unnecessary escaping |

| | | |
|--------------------|------------------|-----------------------------------|
| xri:@example/./abc | xri:@example/abc | Avoid ./ and ../ in absolute XRIs |
|--------------------|------------------|-----------------------------------|

1132

Table 3: Examples of XRI canonicalization recommendations.

1133

3 Resolution

1134

3.1 Introduction

1135 XRI resolution is the process of dereferencing an XRI to a network endpoint in order to
1136 communicate with the resource identified by the XRI. Because XRIs may be used across a wide
1137 variety of communities and applications, including as database keys, filenames, directory keys,
1138 object IDs, and XML IDs, no single resolution mechanism may be appropriate for all XRIs.
1139 However, in the interest of promoting interoperability, this specification defines a simple, flexible
1140 resolution protocol that relies exclusively on HTTP (or HTTPS) for network transport.

1141 Identifier management policies are defined on a community-by-community basis. With XRIs, the
1142 authoritative community is specified by the authority segment of the XRI (section 2.1.1). When a
1143 community chooses to create a new identifier authority, it SHOULD define a policy for assigning
1144 and managing identifiers under this authority. Furthermore, it SHOULD define what resolution
1145 protocol(s) can be used for resolving identifiers assigned by the authority.

1146

3.1.1 Assumptions

1147 This resolution protocol makes several minimal assumptions about the XRIs being resolved:

- 1148 • The endpoints representing the top-level authority for any globally unique XRI are
1149 identified with the “URI-authority” or “XRI-authority” segment of the XRI as defined in
1150 section 2.1.1. If the endpoint identified by an XRI authority begins with a cross-reference
1151 to another URI scheme (for example, the URN scheme), this identifier must be resolvable
1152 by the root authority in that community. (Although other protocols could be specified by
1153 that authority to resolve such URI cross-references, such protocols are outside the scope
1154 of this specification.)
- 1155 • Only absolute XRIs are resolved using this protocol. To resolve a relative XRI, it must be
1156 converted into an absolute XRI using the procedure in section 2.3.
- 1157 • The XRI being resolved has been converted into URI normal form, following the rules in
1158 section 2.2.4.3.
- 1159 • Data or metadata associated with a single XRI may be retrieved or manipulated by
1160 multiple protocols at multiple endpoints.
- 1161 • Each endpoint may present a different subset, type, or representation of data or
1162 metadata associated with the identified resource.

1163

3.1.2 Phases of Resolution

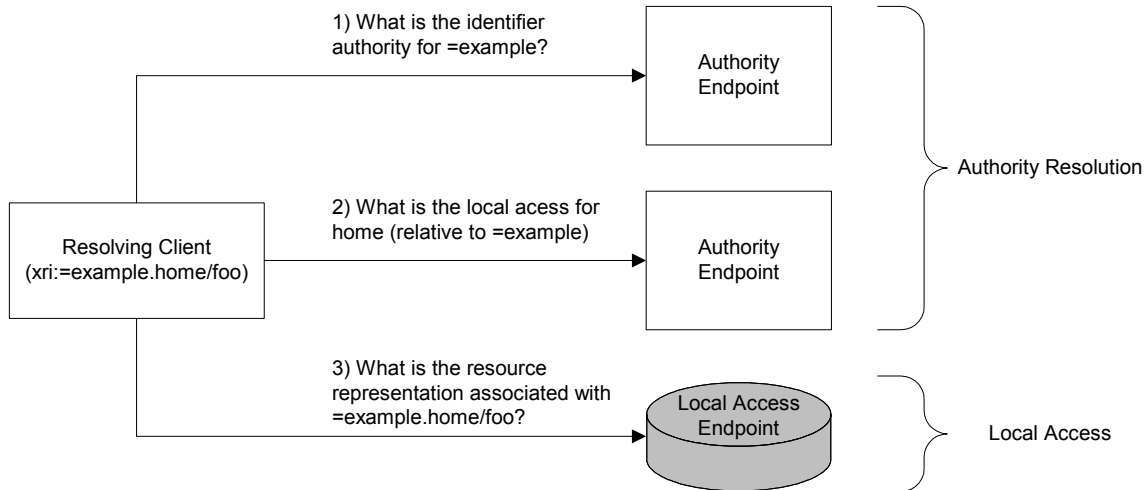
1164 The XRI resolution protocol is designed to be as simple and flexible as possible given the
1165 assumptions above. Based on the structure of XRIs, it consists of two phases:

- 1166 • Authority resolution
- 1167 • Local access

1168 Authority resolution is the process of finding the endpoint or endpoints representing the top-level
1169 identifier authority for the XRI. The result of authority resolution is a list of local access endpoints
1170 identified by one or more URIs and supporting at least one local access protocol. An XRI resolver
1171 chooses one of these endpoints and then accesses it using the desired local access protocol.

1172 Figure 1 illustrates these two phases of XRI resolution:

1173



1174
1175

Figure 1: Phases of Resolution

1176 **3.1.3 XRI vs. URI Authorities**

1177 As described in sections 2.1.1.1 and 2.1.1.2, XRI and URI authorities have different syntactic
1178 structures, partially due to the higher layer of abstraction represented by XRI authorities. For this
1179 reason, XRI authorities are resolved sub-segment by sub-segment as described in section 3.2,
1180 while URI authorities, since they are based on DNS names or IP addresses, are resolved by
1181 transforming the XRI to a HTTP URI as described in section 3.3.

1182 **3.1.4 XRI Metadata Reserved for XRI Resolution**

1183 As defined in section 2.1.1.3, the GCS symbol "\$" is reserved for XRI metadata, i.e., special
1184 identifiers assigned by this specification or the XRI Metadata Specification to describe or resolve
1185 other identifiers.

1186 Within the "\$" namespace, the identifier "\$r" is reserved for identifiers assigned by this resolution
1187 specification. Table 4 summarizes these identifiers.
1188

| Identifier | Use | See Section |
|------------|---|-------------|
| \$r.s | XML namespace for XRI resolution schemas | 3.2.2 |
| \$r.a | Namespace for local access protocol types | 3.4.1 |
| \$r.t | Namespace for resource representation types | 3.4.4 |

1189

Table 4: Special identifiers reserved for XRI resolution.

1190 **3.2 XRI Authority Resolution**

1191 **3.2.1 Overview**

1192 XRI authority resolution is an iterative process that resolves the sub-segments within the XRI
1193 authority segment from left to right. Each sub-segment is resolved in the context of the sub-
1194 segment immediately to the left. The first (or leftmost) sub-segment specifies the root of the
1195 identifier community. Each XRI community provides, by definition, one or more network endpoints

1196 (HTTP or HTTPS URIs) that answer resolution requests for identifiers in its context. This starting
1197 point is further discussed in section 3.2.3.

1198 After the first HTTP/HTTPS URI is determined based on the identifier community, the resolution
1199 process proceeds to the next sub-segment to the right. Each sub-segment is resolved to an
1200 Identifier Authority Descriptor as defined in section 3.2.2. This XML instance provides the data
1201 and metadata necessary to construct the URI for the next identifier authority as described in
1202 section 3.2.4. Once the final authority is reached, the Identifier Authority Descriptor provides the
1203 available local access service protocol(s) as discussed in section 3.4. In addition, the Identifier
1204 Authority Descriptor can provide a mapping of other XRIs used for this identifier authority.

1205 All three options—next authority, local access, or mapping—are available at every iteration. For
1206 example, the XRI authority identifier “@a.b.c” may be the prefix to another identifier authority with
1207 the XRI “@a.b.c.d”, in which case “@a.b.c” must resolve to a Identifier Authority Descriptor. Or
1208 “@a.b.c” may be a local access endpoint itself, in which case its Identifier Authority Descriptor will
1209 contain references to local access services. Finally, this Identifier Authority Descriptor can also
1210 assert that the identifier “xri:@a.b.c” maps to the identifier “xri:@:1:2:3” in order to provide
1211 resolvers or caches with a long-term persistent XRI.

1212 3.2.2 Identifier Authority Descriptors

1213 To provide a straightforward, flexible resolution mechanism, XRI authority endpoints are
1214 described using a simple XML document with a very flexible content model. Its purpose is only to
1215 provide the data and metadata necessary to support delegated resolution and access of XRI-
1216 identified authorities and resources.

1217 The formal XML Schema definition of an XDI Identifier Authority Descriptor is provided in
1218 Appendix B. The following example illustrates the fields defined in this schema:
1219

```
1220 <IdentifierAuthority xmlns="xri:$r.s/IdentifierAuthority">  
1221   <Resolved>.foo</Resolved>  
1222   <NextAuthority>  
1223     <URI>http://xri.example.com</URI>  
1224     <URI>https://xri.example.com</URI>  
1225   </NextAuthority>  
1226   <LocalAccess>  
1227     <Service> xri:$r.a/X2R</Service>  
1228     <Type>application/rddl+xml</Type>  
1229     <URI>http://xri.example.com</URI>  
1230   </LocalAccess>  
1231   <LocalAccess>  
1232     <Service> xri:$r.a/X2R</Service>  
1233     <Type>image/jpeg</Type>  
1234     <URI>http://pictures.xri.example.com</URI>  
1235   </LocalAccess>  
1236   <Mapping>xri:@:1:2:3</Mapping>  
1237 </IdentifierAuthority>
```

1238
1239 All schema elements are in the XML namespace “xri:\$r.s/IdentifierAuthority”. Following are the
1240 elements and attributes that comprise the IdentifierAuthority document type:

1241 **/IdentifierAuthority**

1242 Required. The outer element of the IdentifierAuthority document.

1243 **/IdentifierAuthority/Expires**

1244 0 or 1. The UTC time at which this document MUST no longer be relied upon. A resolver
1245 MAY discard this document before the time indicated in this result. If the HTTP transport
1246 caching semantics specify an expiry time which is earlier than the time expressed in this

1247 attribute, then the “IdentifierAuthority” document MUST no longer be relied upon after the
1248 expiry time declared in the HTTP headers per section 13.2 of [RFC2616].

1249 **/IdentifierAuthority/Resolved**

1250 0 or 1. Expresses the XRI identifier sub-segment whose resolution results in this
1251 IdentifierAuthority document. This field can be used in conjunction with Digital Signatures
1252 to provide secure resolution (functionality that is not specified in this document, but which
1253 will be part of a future deliverable of the OASIS XRI TC). This field may also be useful for
1254 debugging or auditing purposes.

1255 **/IdentifierAuthority/NextAuthority**

1256 0 or 1. Indicates the next identifier authority to query if the just-resolved sub-segment
1257 represents another identifier authority. If the just-resolved sub-segment does not identify
1258 another identifier authority, but rather a resource in the context of the current identifier
1259 authority, then the “NextAuthority” element may not be present.

1260 **/IdentifierAuthority/NextAuthority/URI**

1261 1 or more. Required if “NextAuthority” element is present. Indicates the transport level
1262 URI where the next identifier authority can be contacted. Required to be an HTTP or
1263 HTTPS URI by this specification. Future extensions may use other transport protocols.

1264 **/IdentifierAuthority/LocalAccess**

1265 0 or more. Indicates that the just-resolved sub-segment specifies a identifier authority
1266 where local access service is available.

1267 **/IdentifierAuthority/LocalAccess/Service**

1268 0 or 1. Indicates the type of local access service. The service type is specified by a URI
1269 (including the URI normal form of an XRI). This specification defines one service: “X2R”
1270 which is identified with the URI “xri:\$r.a/X2R” (see section 3.4.1.) An X2R service
1271 converts the XRI into a digital representation of that resource. No more specific
1272 semantics are defined. If this element is absent, then the service associated with this
1273 local access endpoint is X2R.

1274 **/IdentifierAuthority/LocalAccess/Type**

1275 0 or more. The media type of content available at this service. If this element is not
1276 present, then no assumption can be made about the type of data available at this
1277 endpoint. The content of this attribute must be of the form of a media type as defined in
1278 [RFC2046]. This element may appear multiple times to indicate multiple media types
1279 available through this local access service.

1280 **/IdentifierAuthority/LocalAccess/URI**

1281 1 or more. Required if “LocalAccess” element is present. Indicates the transport-level URI
1282 at which the local access service can be requested.

1283 **/IdentifierAuthority/Mapping and /IdentifierAuthority/NextAuthority/Mapping**

1284 0 or more. Represents another XRI that the Identifier authority specified by the resolution
1285 step may also be known as. This must be an absolute XRI (“absolute-xri” in the ABNF,
1286 section 2.1.)

1287 Identifier authority mapping may be used, for example, to assert that a identifier authority
1288 known by a reassignable XRI may also be known by one or more persistent XRIs.
1289 Mapping may also be used to express that a particular authority may be known by a
1290 different reassignable XRI than the one which is being resolved. Both cases may be
1291 particularly useful in populating or querying a cache.
1292

1293 Identifier Authority Descriptor documents have an “open schema” that allows other elements and
1294 attributes from other namespaces to be added throughout. These points of extensibility can be
1295 used to deploy new identifier authority or local access resolution schemes.

1296 Another possible extension is the attachment of XML Digital Signatures and SAML assertions to
 1297 support secure resolution. The current protocol does not specify such security features, although
 1298 a secure resolution protocol is a future deliverable of the OASIS XRI TC and can be
 1299 accommodated within the current Identifier Authority Descriptor schema.

1300 3.2.3 Initiating Resolution

1301 With an XRI authority, the first sub-segment corresponding to the community root may be a global
 1302 context symbol (GCS) or a cross-reference. In either case, the associated community must have
 1303 published an Identifier Authority Descriptor that contains one or more HTTP or HTTPS URIs
 1304 declaring the root resolvers for the community. This Identifier Authority Descriptor is known a
 1305 *priori* and is part of the configuration of a resolver, not unlike the configuration of root DNS
 1306 servers in a DNS resolver.

1307 It is important to note that global context symbols are considered a segment by themselves.
 1308 Therefore, if the sub-segment following the GCS does not begin with a colon (meaning it is not a
 1309 persistent identifier), then a dot is implied. Table 5 and Table 6 demonstrate the parsing of such a
 1310 sub-segment in the case of a GCS and a cross-reference, respectively.

1311

| | |
|-----------------------------------|---------------------------|
| XRI | xri:@example.internal/foo |
| Identifier Authority | @example.internal |
| Identifier Community | @ |
| First Sub-segment Resolved | .example |

1312 Table 5: Parsing the first sub-segment of an XRI that begins with a global context symbol.

1313

| | |
|-----------------------------------|---|
| XRI | xri:(http://www.example.com).internal/foo |
| Identifier Authority | (http://www.example.com).internal |
| Identifier Community | (http://www.example.com) |
| First Sub-segment Resolved | .internal |

1314 Table 6: Parsing the first sub-segment of an XRI that begins with a cross-reference.

1315

1316 3.2.4 Iterating Resolution

1317 Once the Identifier Authority Descriptor representing the community root authority is known, the
 1318 resolution process begins an iteration by constructing the Next Resolution URI. With each
 1319 iteration, the Next Resolution URI is constructed by concatenating the same three parts:

- 1320 1. The Next Authority URI extracted from the Identifier Authority Descriptor corresponding to
 1321 the current context,
- 1322 2. The XRI syntax delimiter (“.” or “:”) preceding the next sub-segment, and
- 1323 3. The next sub-segment value itself (see the clarification regarding cross-references in
 1324 section 3.2.6).
 1325

1326 The URI which forms the base of the Next Resolution URI is the value of a URI element found at
 1327 XPath **/IdentifierAuthority/NextAuthority/URI** in the Identifier Authority Descriptor. If this URI
 1328 does not end with a “/” character, one must be appended before proceeding. Then this URI is

1329 concatenated with the separator and the URI normal form (section 2.2.4.3) of the XRI sub-
1330 segment being resolved. As noted above, if there is no separator character preceding the sub-
1331 segment, a "." MUST be used for the separator.

1332 For example, when resolving the "c" sub-segment of "xri:@a.b.c", if the Next Authority URI
1333 resulting from the resolution of "xri:@a.b" is "http://example.com/xri-authority/", then the Next
1334 Resolution URI is the concatenation of "http://example.com/xri-authority/" with "." and "c", yielding
1335 "http://example.com/xri-authority/.c". An HTTP request is made to this URI, and the next Identifier
1336 Authority Descriptor for the context "xri:@a.b.c" is retrieved.

1337 Construction of the Next Resolution URI is more formally described in this pseudo-code:

1338

```
1339 ia-uri = authority-uri  
1340  
1341 if (authority-uri doesn't end in "/"):  
1342     ia-uri = ia-uri + "/"  
1343  
1344 if (current-sub-segment isn't preceded with "." or ":" separator):  
1345     ia-uri = ia-uri + "."  
1346 else:  
1347     ia-uri = ia-uri + separator  
1348  
1349 ia-uri = ia-uri + uri-escape(sub-segment)
```

1350

1351 Once the Next Resolution URI is constructed, an HTTP or HTTPS GET request is made using
1352 this URI. Each GET request results in a 2XX or 304 HTTP response. The HTTP/HTTPS response
1353 should either contain the next Identifier Authority Descriptor or, with a 304 response, signify that
1354 the cached version on the client is still valid (depending on the client's HTTP request). HTTP
1355 caching semantics should be leveraged as much as possible to support the efficiency and
1356 scalability of this HTTP-based resolution system. The recommended use of HTTP caching
1357 headers is described in more detail in section 3.5.1.

1358 Any ultimate response besides a HTTP 2XX or 304 should be considered an error in the
1359 resolution process. There is no restriction on intermediate redirects (i.e., 3XX result codes) or
1360 other result codes (e.g., a 100 HTTP response) that eventually result in a 2XX or 304 response
1361 through normal operation of [RFC2616]. The content of this ultimate response will be a new
1362 Identifier Authority Descriptor for the context of the sub-segment being resolved.

1363 If there are more sub-segments in the XRI authority segment, the process iterates with the next
1364 sub-segment. If there are no more sub-segments, the final context (as described by the final
1365 Identifier Authority Descriptor retrieved) can be used for local access services as described in
1366 section 3.4.

1367 3.2.5 Examples

1368 Following is an example of resolving the authority portion of this XRI:

```
1369 xri:=example.home.base/foo.bar
```

1370 Assume that the URI for the "=" global context symbol is "http://equals.xri.oasis-open.org/xri-
1371 resolve" (found in **/IdentifierAuthority/NextAuthority/URI** of the Identifier Authority Descriptor
1372 for this community). As explained in 3.2.3, this information, which provides a starting point for
1373 resolution, is known *a priori* and is part of the configuration of the resolver.

1374

1375 **Resolving "=example"**

1376 The following HTTP request is made to "equals.xri.oasis-open.org":

```
1377 GET /xri-resolve/.example HTTP/1.1  
1378 If-Modified-Since: Fri, 31 Oct 2003 19:43:31 GMT  
1379
```



1380 <other HTTP headers>

1381

1382 The following HTTP response is received from “equals.xri.oasis-open.org” (the content has
1383 changed since “Fri, 31 Oct 2003 19:43:31 GMT”):

```
1384 200 OK HTTP/1.1
1385 Content-Type: application/xriad+xml
1386 Expires: Fri, 7 Nov 2003 19:43:31 GMT
1387 <other HTTP headers>
1388
1389 <IdentifierAuthority xmlns="...">
1390 <Resolved>.example</Resolved>
1391 <NextAuthority>
1392 <URI>
1393 http://xri.example.com/xri-resolve/
1394 </URI>
1395 </NextAuthority>
1396 <LocalAccess>...</LocalAccess>
1397
1398 </IdentifierAuthority>
```

1399

1400 **Resolving “=example.home”**

1401 Appending the next sub-segment “home” to the URI “http://xri.example.com/xri-resolve/” yields
1402 the URI “http://xri.example.com/xri-resolve/.home”, and the following HTTP request is made to
1403 xri.example.com:

```
1404 GET /xri-resolve/.home HTTP/1.1
1405 <other HTTP headers>
```

1406

1407 The following HTTP response is received from xri.example.com:

```
1408 200 OK HTTP/1.1
1409 Content-Type: application/xriad+xml
1410 If-Modified-Since: Fri, 31 Oct 2003 19:43:32 GMT
1411 <other HTTP headers>
1412
1413 <IdentifierAuthority xmlns="...">
1414 <Resolved>.home</Resolved>
1415 <NextAuthority>
1416 <URI>
1417 http://xri.home.example.com/xri-resolve/
1418 </URI>
1419 </NextAuthority>
1420 <LocalAccess>...</LocalAccess>
1421 ...
1422 </IdentifierAuthority>
```

1423

1424 **Resolving “=example.home.base”**

1425 Appending the next sub-segment “base” to the URI
1426 “http://xri.home.example.com/xri-resolve/” gives the URI
1427 “http://xri.home.example.com/xri-resolve/.base”:

```
1428 GET /xri-resolve/.base HTTP/1.1
1429 If-Modified-Since: Fri, 31 Oct 2003 19:43:32 GMT
1430
1431 <other HTTP headers>
```

1432

1433 The following HTTP response is received from xri.home.example.com:

```

1434 200 OK HTTP/1.1
1435 Content-type: application/xriad+xml
1436 Expires: Fri, 7 Nov 2003 19:43:33 GMT
1437
1438 <other HTTP headers>
1439
1440 <IdentifierAuthority xmlns="...">
1441 <Resolved>.base</Resolved>
1442 <NextAuthority>...</NextAuthority>
1443 <LocalAccess>
1444 <URI>
1445 http://xrilocal.home.example.com/xri-local/
1446 </URI>
1447 <URI>
1448 https://xrilocal.home.example.com/xri-local/
1449 </URI>
1450 </LocalAccess>
1451 ...
1452 </IdentifierAuthority>

```

1453
1454 The result of the final Identifier authority resolution step is the set of HTTP and HTTPS URIs
1455 shown above that can be used for local access services.

1456 **3.2.6 Resolving Cross-References in XRI Authorities**

1457 A sub-segment within an XRI authority segment may be a cross-reference (see sections 2.1.1.4
1458 and 2.1.2). Resolving a cross-reference is identical to resolving any other sub-segment because,
1459 from the standpoint of generic XRI resolution, the cross-reference is considered opaque. In other
1460 words, the value of the cross-reference (including the parentheses) is considered the literal value
1461 of the sub-segment for the purpose of constructing the Next Resolution URI as described in
1462 section 3.2.4.

1463 Table 7 provides several examples. In each of these examples, sub-segment “b” resolves to a
1464 Next Authority URI of “http://example.com/xri-authority/”.

| Cross-reference type | Example XRI | Next Resoluton URI after resolving “xri:@:a:b” |
|----------------------|---------------------------------------|---|
| Absolute XRI | xri:@:a:b:(@:1:2:3).e/f | http://example.com/xri-authority/:(@:1:2:3) |
| Absolute URI | xri:@:a:b.(mailto:jd@example.com).e/f | http://example.com/xri-authority/.(mailto:jd@example.com) |
| Relative XRI | xri:@:a:b:(.c.d).e/f | http://example.com/xri-authority/:(.c.d) |

1466 Table 7: Examples of the Next Authority URIs constructed using different types of cross-references.

1467 Note that specific identifier communities may specify special resolution rules for specific types of
1468 cross-references, but such extensions are out of scope for this specification.

1469 **3.2.7 User Relative XRIs**

1470 A special case of XRI authority resolution is the user-relative context symbol (“*”). This symbol
1471 means the XRI authority is defined by the user of the XRI rather than being specified in the XRI
1472 itself. For example, a frequent XRI user could use the “*” symbol to enter their own XRI
1473 “shortcuts” or “speed names” that will be resolved into absolute XRIs using a mapping stored at
1474 the user’s preferred identifier authority. Therefore, these XRIs are not resolvable without the a

1475 *priori* establishment of an authority designated by the user, typically in the configuration of the
1476 user's XRI resolver.

1477 An XRI beginning with the user-relative context symbol **MUST** be transformed into an absolute
1478 XRI that does **NOT** begin with a user-relative context symbol before it can be resolved using the
1479 resolution protocol defined in this specification. To perform this transformation, the XRI value
1480 following the "*" symbol **MUST** be treated as a relative XRI reference and resolved relative to a
1481 base XRI as defined in section 2.3. The mapping of the "*" symbol to this base XRI is
1482 implementation-dependent; however, the configuration of such mapping **SHOULD** be easily
1483 available to the user.

1484 Note that in most cases, mapping requires simply replacing the "*" character with a prefix
1485 corresponding to an absolute identifier authority. For example, if XRI resolver is pre-configured
1486 with a default community of "@employer/", then the XRI "xri:*Mary/workstation" would be
1487 converted into "xri:@employer/Mary/workstation".

1488 **3.3 URI Authority Resolution**

1489 A URI-authority segment (section 2.1.1.1) includes either a DNS name or an IP address that
1490 specifies the location of the endpoint with which to perform local access. Thus, the process for
1491 converting XRIs with URI authorities into local access URIs is simple. First, the XRI must be
1492 converted into URI-escaped form (section 2.2.4.3). Then the scheme is converted from "xri:" to
1493 "http:", and an HTTP request is performed on the resulting URI, as described in section 3.4,
1494 "Local Access", below.

1495 For example, the XRI "xri://www.example.com/foo.bar" is transformed to the HTTP URI
1496 "http://www.example.com/foo.bar".

1497 The use of URI authorities provides backwards compatibility with the large installed base of DNS-
1498 and IP-identifiable resources. Because URI authorities do not support the additional layer of
1499 abstraction and extensibility represented by XRI authority syntax, however, URI authorities are
1500 not recommended for new deployments of XRI identifiers.

1501 **3.4 Local Access**

1502 Local access is the process of interacting with a network endpoint to retrieve a representation of a
1503 network resource identified by an XRI.

1504 **3.4.1 Local Access Service Types**

1505 Any number of protocols may be used for local access. This specification defines an
1506 HTTP/HTTPS local access protocol. An LDAP or DSML local access protocol could be defined by
1507 specifying the appropriate transformation of the XRI local part into an LDAP distinguished name
1508 (including normalization of the XRI local path to the LDAP distinguished name syntax.)

1509 Work on such protocols is left to future specifications. To accommodate such work, this
1510 specification reserved namespace, "\$r.a", for enumerating local access service types. One
1511 enumeration, "X2R", is defined in section 3.2.2 under "/IdentifierAuthority/LocalAccess/Service".

1512 **3.4.2 HTTP/HTTPS Local Access**

1513 The HTTP/HTTPS local access protocol does not specify the semantics of the local access
1514 interaction, nor the form of the local access requests. The only semantics defined are those in
1515 **[RFC2616]**. Special attention should be paid to the semantics of the four main HTTP verbs: GET,
1516 PUT, POST, and DELETE. For example, clients performing local access typically would use GET
1517 when wishing to retrieve representations of a resource on the network.

1518 This specification does not impose particular semantics beyond what is defined in **[RFC2616]**, but
1519 users of this specification are encouraged to review the **[REST]** architecture when building
1520 applications using XRIs. Local access is not limited to the REST model of interaction, however.

1521 For example, HTTP local access could be leveraged for the delivery of SOAP messages over
1522 HTTP POST, or via use of the GET HTTP verb as a generic read-only resolution infrastructure.

1523 The HTTP/HTTPS local access binding defined in this section is flexible enough to be used for a
1524 variety of resources. It makes no assumptions about the type of resource identified by the XRI
1525 being resolved. The resource type must be established through the context in which the XRI was
1526 originally used (e.g. an XML document) or discovered through use of the HTTP local access
1527 protocol (e.g., through the HTTP Content-Type header).

1528 **3.4.3 Constructing a Local Access HTTP/HTTPS URI**

1529 This section defines the construction of URIs for local access to resources identified with XRI
1530 authorities. The construction of URIs for local access to resources identified with URI authorities
1531 is defined in section 3.3.

1532 The HTTP/HTTPS URI with which to perform local access is constructed by concatenating the
1533 Local Access URI from the Identifier Authority Descriptor (section 3.2.2) with the local part of the
1534 XRI. Specifically, the URI from the element identified with the relative XPATH
1535 **/IdentifierAuthority/LocalAccess/URI** is concatenated with the URI normal form (section
1536 2.2.4.3) version of the remaining relative-path (section 2.1). If the LocalAccess URI does not
1537 terminate in a “/”, it MUST be inserted before the relative-path.

1538 The following pseudocode describes the process for creating the concrete HTTP/HTTPS URI to
1539 which a local access request is made:

```
1540 concrete-http-uri = localaccess-uri  
1541  
1542 if (localaccess-uri does not end in “/”):  
1543     concrete-http-uri = localaccess-uri + “/”  
1544  
1545 concrete-http-uri = localaccess-uri + uri-escape(relative-path)
```

1546

1547 The verb used in the resulting HTTP/HTTPS request may be any of the verbs defined in
1548 **[RFC2616]**, though not all verbs may be supported at every endpoint. All local access endpoints
1549 SHOULD support at least the GET verb, and this should return either a representation of the
1550 identified resource or metadata about the resource.

1551 The full suite of HTTP content negotiation features is available to clients when performing local
1552 access. For example, if the local access service URI is “http://xri.example.com/xri-local”, then the
1553 following local access HTTP request for “xri:=example.home/foo.bar” could be made to
1554 “xri.example.com”:

```
1555 GET /xri-local/foo.bar HTTP/1.1  
1556 If-Modified-Since: Fri, 31 Oct 2003 19:43:33 GMT  
1557 <other HTTP headers>  
1558
```

1559

1560 The following HTTP response should then be received from xri.example.com:

```
1561 200 OK HTTP/1.1  
1562 Expires: Sat, 1 Nov 2003 19:43:33 GMT  
1563 Content-Type: text/plain  
1564 <other HTTP headers>  
1565  
1566 This is the result of a local access request.
```

1567

1568 **3.4.4 Using a Cross-Reference to Specify a Representation Type**

1569 A cross-reference MAY be used to specify a desired resource representation type when
1570 performing local access. The namespace “\$r.t” is reserved for this purpose. This specification

1571 does not enumerate such types; they are further defined in the XRI Metadata Specification [tktk
1572 need reference].

1573 To specify a particular resource representation type using “\$.t” metadata, a “\$.t” cross-reference
1574 is appended to the XRI during a local access request. For example, an RDDDL document could be
1575 specified by appending the cross-reference “(\$.t/RDDL)”.

1576 The following example illustrates this technique. Assuming the original XRI being resolved is
1577 “xri:=example.home/foo.bar” and the local access URI is “http://xri.example.com/xri-local/”, the
1578 following HTTP request would request the RDDDL document describing this resource:

1579

```
1580 GET /xri-local/foo.bar/%28$.t%2FRDDL%29 HTTP/1.1  
1581 <other HTTP headers>
```

1582

1583 Note that the cross-reference is escaped per the rules for the URI normal form of an XRI in
1584 section 2.2.4.3.

1585 The resulting HTTP response would be:

```
1586 200 OK HTTP/1.1  
1587 <cache-headers>  
1588 <other HTTP headers>  
1589  
1590 <content of representation of RDDDL for xri:=example.home/foo.bar>
```

1591

1592 3.5 HTTP Headers

1593 3.5.1 Caching

1594 The full caching capabilities of **[RFC2616]** should be leveraged during both identifier authority
1595 resolution and local access. Specifically, implementations of XRI resolution SHOULD implement
1596 the caching model described section 13 of **[RFC2616]**. In particular, the “Expiration Model” of
1597 section 13.2 SHOULD be used, as this requires the fewest round-trip network connections.

1598 All servers providing identifier authority lookup responses SHOULD send the Cache-Control or
1599 Expires headers per section 13.2 of **[RFC2616]**, unless there are overriding security or policy
1600 reasons that dictate otherwise.

1601 3.5.2 Location

1602 During identifier authority resolution, “Location” headers may be present per the **[RFC2616]**
1603 specification (i.e., during 3XX redirects). Redirects SHOULD be made cacheable through
1604 appropriate HTTP headers.

1605 During the local access phase, redirects may be returned, and the “Location” field may contain an
1606 HTTP/HTTPS URI or an XRI in URI normal form. This use of redirects constitutes a mapping
1607 facility that allows one XRI to resolve into another during local access. If the local access server is
1608 aware of the HTTP/HTTPS URI where the XRI may be accessed, it can provide a “Location”
1609 header containing an HTTP/HTTPS URI. In this case, it SHOULD provide an “X-XRI-Canonical”
1610 header (see below) to describe the XRI to which the redirection is targeting. If the local access
1611 server knows only of the target XRI, then it MUST return a redirection header (3XX code) with the
1612 “Location” field containing an XRI.

1613 3.5.3 Content-Location

1614 “Content-Location” may be used during local access where the resource being accessed is an
1615 “attribute” or “view” of another resource. This usually would occur in the case where metadata is
1616 being accessed using a trailing cross reference to an XRI value under the “\$.t” namespace (see

1617 section 3.4.4). Such a “Content-Location” header would specify where the resource itself may be
1618 accessible (rather than the metadata). This is not required and MUST NOT be required by
1619 resolving clients for proper operation. The content-location SHOULD be an HTTP/HTTPS URI if
1620 the local access server is aware of the HTTP/HTTPS location, otherwise it MAY be an XRI.

1621 3.5.4 Content-Type

1622 “Content-type” is required in the HTTP/HTTPS response during identifier authority resolution,
1623 both when returning an Identifier Authority Descriptor and for the HTTP/HTTPS responses during
1624 local access.

1625 The “Content-type” header in the 2XX responses in identifier authority resolution for each sub-
1626 segment MUST contain the value “application/xri+xml”, specifying that the content is an
1627 Identifier Authority Descriptor (section 3.2.2).

1628 In local access, clients and servers MAY negotiate content type using standard HTTP content
1629 negotiation features. Whether or not this feature is used, however, the server MUST respond with
1630 an appropriate media type in the “Content-type” header.

1631 3.5.5 X-XRI-Canonical

1632 This header is present only in HTTP/HTTPS redirects from local access servers. Its purpose is to
1633 notify a resolving client that the redirect is occurring because the original XRI is a mapping to
1634 another XRI. The value of this header is the target XRI in URI normal form (section 2.2.4.3). This
1635 header MAY be present even when the Location: header is present and contains an XRI. This
1636 header SHOULD be present when the Location: header is present and contains a HTTP/HTTPS
1637 or other URI.

1638 Form:

1639 `X-XRI-Canonical: <xri-in-uri-normal-form>`

1640 3.6 Other HTTP Features

1641 HTTP provides a number of other features including transfer-coding, proxying, validation-model
1642 caching, etc. All of these features may be used insofar as they do not conflict with the required
1643 uses of HTTP described in this document.

1644 3.7 Caching and Efficiency

1645 Resolution clients are encouraged to perform caching above the HTTP level in addition to at the
1646 HTTP level. For best results, however, resolution clients SHOULD be conservative with caching
1647 expiration semantics, including cache expiration dates. This implies that in a series of HTTP
1648 redirects, for example, the results of the entire process should only be cached as long as the
1649 shortest period of time allowed by any of the intermediate HTTP responses.

1650 Because not all HTTP client libraries expose caching expiration to applications, identifier
1651 authorities and local access servers SHOULD NOT use cacheable redirects with expiration times
1652 which are relatively short compared to the expiration times of other HTTP responses in the
1653 resolution or local access chain. In general, all XRI deployments should be mindful of limitations
1654 in current HTTP clients and proxies.

1655 For Identifier Authority Descriptors, the cache expiration time may also be shortened by the
1656 expiration time provided in the Identifier Authority Descriptor at **//IdentifierAuthority/Expires** (if
1657 present). That is, if the expiration time in **//IdentifierAuthority/Expires** is sooner than the
1658 expiration time calculated from the HTTP caching semantics, then the Identifier Authority
1659 Descriptor SHOULD be discarded before the expiration time in **//IdentifierAuthority/Expires**.

1660 With both application-level and HTTP-level caching, the resolution process is designed to have
1661 minimal overhead. In particular, because each sub-segment of an abstractly-identified XRI
1662 authority is resolved separately, each step of that resolution is a completely independent,

1663 cacheable HTTP request. For this reason, resolution of top-level (leftmost) sub-segments, which
1664 are common to more identifiers, will naturally result in a greater number of cache hits than
1665 resolution of sub-segments further to the right.

1666 **3.8 Points of Extensibility**

1667 The XRI resolution scheme described here leverages extensible mechanisms such as HTTP
1668 and XML to provide maximum flexibility. Specifically, changes or additions can be made at the
1669 following points of extensibility:

- 1670 • HTTP negotiation of content types, language, encoding, etc.
- 1671 • Use of HTTP verbs such as POST, PUT and DELETE during local access.
- 1672 • Use of HTTP redirects (3XX) or other response codes during identifier authority
1673 resolution or local access.
- 1674 • Insertion of new elements or attributes in the Identifier Authority Descriptor.
- 1675 • Use of cross-references within XRIs, particularly for associating new types of metadata
1676 with a resource (see section 3.4.4 for an example).

1677

1678

1679 4 Security and Data Protection

1680 4.1 Secure Resolution

1681 The resolution protocol described in section 3 is not intrinsically trustworthy. It is expected that, in
1682 practice, some combination of DNSSEC, SSL, TLS, and other existing technologies will be
1683 employed to increase the security of the resolution process.

1684 While such enhancements are outside the scope of this specification, an XRI Secure Resolution
1685 Specification is a future deliverable of the OASIS XRI TC. Additional follow-on work is also
1686 expected to define best practices and facilitate interoperability.

1687 4.2 XRI Metadata

1688 The use of "\$" cross-references for encoding XRI metadata in an XRI may involve other security
1689 and data protection considerations that are outside the scope of this specification. These
1690 considerations are addressed in the XRI Metadata Specification [tktk – need reference].

1691 4.3 XRI Usage in Legacy Infrastructure

1692 Where XRIs are used within the Internet and other computing infrastructure, the security and data
1693 protection considerations relating to XRIs are similar to those of other URI schemes. In this
1694 context the material in **[RFC2396bis]**, section 7, *Security Considerations*, is particularly
1695 informative. It includes a discussion of the following topics:

- 1696 • Reliability and Consistency
- 1697 • Malicious Construction
- 1698 • Rare IP Address Formats
- 1699 • Sensitive Information
- 1700 • Semantic Attacks

1701 This material notes that “a URI does not in itself pose a direct security threat.” In the case of
1702 XRIs, this statement remains true only in legacy environments. As noted below, it may not be true
1703 for new infrastructure that builds on the extensibility of XRI architecture. Such applications must
1704 be developed with independent security reviews for the specific scenarios in which XRIs are
1705 used.

1706 4.4 XRI Usage in Evolving Infrastructure

1707 As XRIs are adopted as abstract identifiers, it is anticipated that new services will be developed
1708 that take advantage of their extensibility. In particular, XRIs may enable new solutions to security
1709 and data protection problems that are not possible using existing URI schemes.

1710 For example, XRI cross-reference syntax permits the inclusion of identifier metadata such as an
1711 encrypted or integrity-checked path, query, or fragment. Cross-references can also be used to
1712 indicate methods of obfuscating, proxying, or redirecting resolution to prevent the exposure of
1713 private or sensitive data. These capabilities may enable new security and data protection features
1714 at the fundamental level of resource identifiers.

1715 A complete discussion of this topic is beyond the scope of this document. However, as a
1716 consequence of the extensibility of XRIs, it is not possible to make definitive statements regarding
1717 security and data protection considerations relating to XRIs.

5 References

5.1 Normative

- 1720 **[RFC1737]** K. Sollins, L. Masinter, *Functional Requirements for Uniform Resource Names*,
1721 <http://www.ietf.org/rfc/rfc1737.txt>, RFC 1737, December 1994.
- 1722 **[RFC2046]** N. Borenstein, N. Freed, *Multipurpose Internet Mail Extensions (MIME) Part Two:*
1723 *Media Types*, <http://www.ietf.org/rfc/rfc2046.txt>, RFC 2046, November 1996.
- 1724 **[RFC2119]** S. Bradner, *Key words for use in RFCs to Indicate Requirement Levels*,
1725 <http://www.ietf.org/rfc/rfc2119.txt>, RFC 2119, March 1997.
- 1726 **[RFC2141]** R. Moats, *URN Syntax*, <http://www.ietf.org/rfc/rfc2141.txt>, IETF RFC 2141, May
1727 1997.
- 1728 **[RFC2234]** D. H. Crocker and P. Overell, *Augmented BNF for Syntax Specifications: ABNF*,
1729 <http://www.ietf.org/rfc/rfc2234.txt>, RFC 2234, November 1997.
- 1730 **[RFC2396]** T. Berners-Lee, R. Fielding, L. Masinter, *Uniform Resource Identifiers (URI):*
1731 *Generic Syntax*, <http://www.ietf.org/rfc/rfc2396.txt>, RFC 2396, August 1998.
- 1732 **[RFC2616]** R. Fielding, J. Gettys, J. Mogul, H. Frystyk, L. Masinter, P. Leach, T. Berners-Lee,
1733 *Hypertext Transfer Protocol -- HTTP/1.1*, <http://www.ietf.org/rfc/rfc2616.txt>, RFC 2616, June
1734 1999.
- 1735 **[RFC2718]** L. Masinter, H. Alvestrand, D. Zigmond, R. Petke, *Guidelines for New URL*
1736 *Schemes*, <http://www.ietf.org/rfc/rfc2718.txt>, RFC 2718, November 1999.
- 1737 **[RFC2732]** R. Hinden, B. Carpenter, L. Masinter, *Format for Literal IPv6 Addresses in URL's*,
1738 <http://www.ietf.org/rfc/rfc2732.txt>, RFC 2732, December, 1999.
- 1739 **[RFC3066]** H. Alvestrand, *Tags for the Identification of Languages*,
1740 <http://www.ietf.org/rfc/rfc3066.txt>, RFC 3066, January, 2001.
- 1741 **[RFC3305]** M. Mealing, R. Denenberg, *Uniform Resource Identifiers (URIs), URLs, and*
1742 *Uniform Resource Names (URNs): Clarifications and Recommendations*,
1743 <http://www.ietf.org/rfc/rfc3305.txt>, RFC 3305, August 2002.
- 1744 **[RFC3490]** P. Faltstrom, P. Hoffman, A. Costello, *Internationalizing Domain Names in*
1745 *Applications (IDNA)*, <http://www.ietf.org/rfc/rfc3490>, RFC 3490, March 2003.
- 1746 **[RFC3491]** P. Hoffman, M. Blanchet, *Nameprep: A Stringprep Profile for Internationalized*
1747 *Domain Names (IDN)*, <http://www.ietf.org/rfc/rfc3491>, RFC 3491, March 2003.
- 1748 **[UML]** Object Management Group, *Unified Modeling Language (UML) Version 1.5*,
1749 <http://www.omg.org/technology/documents/formal/uml.htm>, March 1, 2003.
- 1750 **[Unicode]** The Unicode Consortium. The Unicode Standard, Version 4.0.0, defined by: *The*
1751 *Unicode Standard, Version 4.0* (Boston, MA, Addison-Wesley, 2003. ISBN 0-321-18578-1)
- 1752 **[UniXML]** Duerst, M. and A. Freytag, *Unicode in XML and other Markup Languages*,
1753 Unicode Technical Report #20, World Wide Web Consortium Note, February 2002.
- 1754 **[UTR15]** M. Davis, M. Duerst, *Unicode Normalization Forms*,
1755 <http://www.unicode.org/unicode/reports/tr15/tr15-23.html>, April 17, 2003.
- 1756 **[XML]** T. Bray, J. Paoli, C.M. Sperberg-McQueen, E. Maler, *Extensible Markup Language*
1757 *(XML) 1.0 (Second Edition) W3C Recommendation*, <http://www.w3.org/TR/REC-xml>, October
1758 2000.
- 1759 **[XMLSchema2]** P. Biron, A. Malhotra, *XML Schema Part 2: Datatypes W3C*
1760 *Recommendation*, <http://www.w3.org/TR/xmlschema-2/>, May 2001.

1761

5.2 Informative

- 1762 **[IRI]** M. Duerst, M. Suignard, *Internationalized Resource Identifiers (IRIs)*,
1763 <http://www.ietf.org/internet-drafts/draft-duerst-iri-05.txt>, Work-In-Progress, October 2003.
- 1764 **[REST]** <http://internet.conveyor.com/RESTwiki/moin.cgi/FrontPage>
- 1765 **[RFC2396bis]** R. Fielding, *Uniform Resource Identifiers (URI): Generic Syntax*, Internet Draft
1766 draft-fielding-uri-rfc2396bis-03, <http://www.apache.org/~fielding/uri/rev-2002/rfc2396bis.html>,
1767 Work-In-Progress, June 2003.
- 1768 **[XRIReqs]** G. Wachob, D. Reed, M. Le Maitre, D. McAlpin, D. McPherson, *Extensible
1769 Resource Identifier (XRI) Requirements and Glossary v1.0*, [http://www.oasis-
1770 open.org/apps/org/workgroup/xri/download.php/2523/xri-requirements-and-glossary-v1.0.doc](http://www.oasis-open.org/apps/org/workgroup/xri/download.php/2523/xri-requirements-and-glossary-v1.0.doc),
1771 June 2003.
- 1772

1773

Appendix A. Collected ABNF for XRI (Normative)

1774 This section contains the complete ABNF for XRI syntax, which includes the complete ABNF for
 1775 URI from **[RFC2396bis]** of which XRI syntax is a superset. XRI productions use green shading,
 1776 and productions inherited from URI use yellow shading. A valid XRI MUST conform to this ABNF.

1777

1778

```
abs-path      = "/" path-segments
```

1779

1780

```
absolute-xri = "xri:" global-path
```

1781

1782

```
alphanum     = ALPHA / DIGIT
```

1783

1784

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

1785

1786

```
authority-path = URI-authority / XRI-authority
```

1787

1788

```
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
```

1789

1790

1791

1792

1793

1794

```
delims      = "<" / ">" / "%" / DQUOTE
```

1795

1796

```
domainlabel = alphanum [ 0*61( alphanum / "-" ) alphanum ]
```

1797

1798

```
escaped     = "%" HEXDIG HEXDIG
```

1799

1800

```
excluded    = invisible / delims / unwise
```

1801

1802

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

1803

1804

```
gcs-char    = "+" / "=" / "@" / "$" / "*" / "!"
```

1805

1806

```
global-path = authority-path [ local-path ] [ query-frag ]
```

1807

1808

```
global-xri  = global-path [ "?" xri-query ] [ "#" xri-fragment ]
```

1809

1810

```
h4          = 1*4HEXDIG
```

1811

1812

```
hier-part   = net-path / abs-path / rel-path
```

1813

1814

```
host        = [ hostname / IPv4address / IPv6reference ]
```

1815

1816

```
hostname    = idomainlabel qualified
```

1817

1818

```
idomainlabel = 1*ucschar
```

1819

1820

```
invisible   = CTL / SP / %x80-FF
```

1821

1822

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

1823

1824

```
IPv6address = 6( h4 ":" ) 1s32
              / [ "::" 5( h4 ":" ) 1s32
                  / [ "h4" h4 ] "::" 4( h4 ":" ) 1s32
                  / [ *1( h4 ":" ) h4 ] "::" 3( h4 ":" ) 1s32
                  / [ *2( h4 ":" ) h4 ] "::" 2( h4 ":" ) 1s32
```

1825

1826

1827

1828

1829

```

1830 / [ *3( h4 ":" ) h4 ] "::" h4 ":" ls32
1831 / [ *4( h4 ":" ) h4 ] "::" ls32
1832 / [ *5( h4 ":" ) h4 ] "::" h4
1833 / [ *6( h4 ":" ) h4 ] "::"
1834
1835 IPv6reference = "[" IPv6address "]"
1836
1837 local-path = "/" relative-path
1838
1839 ls32 = ( h4 ":" h4 ) / IPv4address
1840 ; least-significant 32 bits of address
1841
1842 mark = "-" / "_" / "." / "!" / "~" / "*" / "'" / "(" / ")"
1843
1844 net-path = "/" authority [ abs-path ]
1845
1846 path-segments = segment *( "/" segment )
1847
1848 pchar = unreserved / escaped / ";" /
1849 ":" / "@" / "&" / "=" / "+" / "$" / ","
1850
1851 port = *DIGIT
1852
1853 qualified = *( "." idomainlabel ) [ "." ]
1854
1855 query = *( pchar / "/" / "?" )
1856
1857 query-frag = [ "?" xri-query ] [ "#" xri-fragment ]
1858
1859 relative-path = *( [ "." ] "/" ) xri-segments
1860
1861 relative-xri = ( local-path / relative-path ) [ query-frag ]
1862
1863 rel-path = path-segments
1864
1865 reserved = "/" / "?" / "#" / "[" / "]" / ";" /
1866 ":" / "@" / "&" / "=" / "+" / "$" / ","
1867
1868 scheme = ALPHA *( ALPHA / DIGIT / "+" / "-" / "." )
1869
1870 segment = *pchar
1871
1872 sub-segment = *xri-pchar / xref
1873
1874 uchar = %xA0-D7FF / %xF900-FDCF / %xFDF0-FFEF /
1875 %x10000-1FFFD / %x20000-2FFFD / %x30000-3FFFF /
1876 %x40000-4FFFF / %x50000-5FFFF / %x60000-6FFFF /
1877 %x70000-7FFFF / %x80000-8FFFF / %x90000-9FFFF /
1878 %xA0000-AFFFF / %xB0000-BFFFF / %xC0000-CFFFF /
1879 %xD0000-DFFFF / %xE1000-EFFFF
1880
1881 unreserved = ALPHA / DIGIT / mark
1882
1883 unwise = "{" / "}" / "|" / "\" / "^" / "`"
1884
1885 URI = scheme ":" hier-part [ "?" query ] [ "#" fragment ]
1886
1887 URI-authority = "/" [ userinfo "@" ] host [ ":" port ]
1888
1889 uric = reserved / unreserved / escaped
1890
1891 userinfo = *( unreserved / escaped / ";" /
1892 ":" / "&" / "=" / "+" / "$" / "," )

```

```

1893
1894 xref          = "(" ( xri-value / URI ) ")"
1895
1896 xref-authority = xref ( "." sub-segment / ":" sub-segment ) *( "."
1897 sub-segment / ":" sub-segment )
1898
1899 XRI           = absolute-xri / relative-xri
1900
1901 XRI-authority = ( gcs-char xri-segment ) / xref-authority
1902
1903 xri-characters = xri-reserved / xri-unreserved / escaped
1904
1905 xri-fragment  = [ xref ] * ( pchar / "/" / "?" )
1906
1907 xri-mark      = "-" / "_" / "~" / "'"
1908
1909 xri-path      = global-path / local-path / relative-path
1910
1911 xri-pchar     = xri-unreserved / escaped / ";" / "!" / "*"
1912 "@ " / "&" / "=" / "+" / "$" / ","
1913
1914 xri-query     = [ xref ] * ( pchar / "/" / "?" )
1915
1916 xri-reserved  = "/" / "?" / "#" / "[" / "]" / "(" / ")" / ";" / ":" /
1917 ", " / "." / "&" / "@" / "=" / "+" / "*" / "$" / "!"
1918
1919 xri-segment   = ( [ "." ] sub-segment / ":" sub-segment )
1920 *( "." sub-segment / ":" sub-segment )
1921
1922 xri-segments  = xri-segment *( "/" xri-segment )
1923
1924 xri-unreserved = ALPHA / DIGIT / ucschar / xri-mark
1925
1926 xri-value     = [ global-path / local-path / relative-path ]
1927 [ query-frag ]
1928
1929

```

1930

Appendix B. XML Schema for XRI Identifier Authority Descriptor (Normative)

1931

```
1932 <xs:schema targetNamespace="xri:$r.s/IdentifierAuthority" xmlns:xs="http://www.w3.org/2001/XMLSchema"
1933 xmlns="xri:$r.s/IdentifierAuthority" elementFormDefault="qualified" attributeFormDefault="unqualified">
1934   <xs:complexType name="IdentifierAuthorityType">
1935     <xs:sequence>
1936       <xs:element name="Resolved" type="ResolvedType" minOccurs="0" maxOccurs="unbounded"/>
1937       <xs:element name="Expires" type="ExpiresType" minOccurs="0"/>
1938       <xs:element name="NextAuthority" type="NextAuthorityType" minOccurs="0"/>
1939       <xs:element name="LocalAccess" type="LocalAccessType" minOccurs="0" maxOccurs="unbounded"/>
1940       <xs:element name="Mapping" type="MappingType" minOccurs="0" maxOccurs="unbounded"/>
1941       <xs:any namespace="##other" processContents="lax" minOccurs="0" maxOccurs="unbounded"/>
1942     </xs:sequence>
1943     <xs:anyAttribute namespace="##other" processContents="lax"/>
1944   </xs:complexType>
1945   <xs:complexType name="NextAuthorityType">
1946     <xs:sequence>
1947       <xs:element name="URI" type="URIType" maxOccurs="unbounded"/>
1948       <xs:element name="Mapping" type="MappingType" minOccurs="0" maxOccurs="unbounded"/>
1949       <xs:any namespace="##other" processContents="lax" minOccurs="0" maxOccurs="unbounded"/>
1950     </xs:sequence>
1951     <xs:anyAttribute namespace="##other" processContents="lax"/>
1952   </xs:complexType>
1953   <xs:complexType name="LocalAccessType">
1954     <xs:sequence>
1955       <xs:element name="Service" type="ServiceType" minOccurs="0"/>
1956       <xs:element name="Type" type="TypeType" minOccurs="0" maxOccurs="unbounded"/>
1957       <xs:element name="URI" type="URIType" maxOccurs="unbounded"/>
1958       <xs:any namespace="##other" processContents="lax" minOccurs="0" maxOccurs="unbounded"/>
1959     </xs:sequence>
1960     <xs:anyAttribute namespace="##other" processContents="lax"/>
1961   </xs:complexType>
1962   <xs:element name="IdentifierAuthority" type="IdentifierAuthorityType"/>
1963   <xs:complexType name="ResolvedType">
1964     <xs:simpleContent>
1965       <xs:extension base="xs:string">
1966         <xs:anyAttribute namespace="##other" processContents="lax"/>
1967       </xs:extension>
1968     </xs:simpleContent>
1969   </xs:complexType>
1970   <xs:complexType name="URIType">
1971     <xs:simpleContent>
1972       <xs:extension base="xs:anyURI">
1973         <xs:anyAttribute namespace="##other" processContents="lax"/>
1974       </xs:extension>
1975     </xs:simpleContent>
1976   </xs:complexType>
1977   <xs:complexType name="ExpiresType">
1978     <xs:simpleContent>
1979       <xs:extension base="xs:dateTime">
1980         <xs:anyAttribute namespace="##other" processContents="lax"/>
1981       </xs:extension>
1982     </xs:simpleContent>
1983   </xs:complexType>
1984   <xs:complexType name="ServiceType">
1985     <xs:simpleContent>
1986       <xs:extension base="xs:anyURI">
1987         <xs:anyAttribute namespace="##other" processContents="lax"/>
1988       </xs:extension>
1989     </xs:simpleContent>
```



```
1990 </xs:complexType>
1991 <xs:complexType name="TypeType">
1992 <xs:simpleContent>
1993 <xs:extension base="xs:string">
1994 <xs:anyAttribute namespace="##other" processContents="lax"/>
1995 </xs:extension>
1996 </xs:simpleContent>
1997 </xs:complexType>
1998 <xs:complexType name="MappingType">
1999 <xs:simpleContent>
2000 <xs:extension base="xs:string">
2001 <xs:anyAttribute namespace="##other" processContents="lax"/>
2002 </xs:extension>
2003 </xs:simpleContent>
2004 </xs:complexType>
2005 </xs:schema>
2006
```

2007
2008

Appendix C. Transforming HTTP URIs to XRI (Non-Normative)

2009 To leverage existing infrastructure, it may sometimes be useful to convert HTTP URIs into XRIs.
2010 Because XRI syntax is, for the most part, a superset of generic URI syntax, the majority of HTTP
2011 URIs can be converted to valid XRIs simply by replacing the scheme “http” with “xri”. Special
2012 consideration, however, must be given to HTTP URIs employing the characters in the “xri-
2013 reserved” production of this specification that differ from those in the “reserved” production of
2014 **[RFC2396]** (as amended by **[RFC2732]**). These include opening parenthesis (“(“), closing
2015 parenthesis (“)”), dot (“.”), asterisk (“*”), and exclamation point (“!”).

2016 Typically, characters in the “reserved” production of **[RFC2396]** that appear in an HTTP URI as
2017 normal characters (i.e. not as syntactic delimiters) are escaped encoded. However, this is not
2018 required in all cases. **[RFC2396]** says

2019 “Characters in the ‘reserved’ set are not reserved in all contexts. The set of characters
2020 actually reserved within any given URI component is defined by that component. In general, a
2021 character is reserved if the semantics of the URI changes if the character is replaced with its
2022 escaped US-ASCII encoding.”

2023 Characters in the “xri-reserved” set that are properly left un-escaped in an HTTP URI may be
2024 semantically significant when the HTTP URI is converted to an XRI. For example,

2025 `http://www.example.com/example1:example2`

2026 is a valid HTTP URI even though it contains an unescaped reserved character – a colon (“:”) –
2027 because section 3.3 of **[RFC2396]** explicitly omits this character from the reserved set for “path”
2028 components. The same unescaped character in an XRI, however, will be interpreted as a
2029 delimiter. If the colon character should not be understood as a delimiter in the resulting XRI, it
2030 must be escaped during conversion. The same applies to the other characters mentioned above.

2031 Generally, any character not in the “xri-pchar” set that appears in the “abs_path”, “query”, or
2032 “fragment” components of the HTTP URI will need to be escaped when converting to an XRI. This
2033 avoids misinterpretation in the resulting XRI following the guidance in section 2.2.4 of this
2034 specification.

2035 Exceptions are possible. For example, if the author of the above HTTP URI intended the colon
2036 character to be interpreted as described in this specification, or if its use would not be
2037 misinterpreted, then it may be left in its unescaped form.

2038 In addition, it may be beneficial to escape other characters like the percent (“%”) character,
2039 particularly if it may be necessary to convert the resulting XRI back to an HTTP URI. Whether
2040 such additional escaping is desirable or not depends on the intended use of the resulting XRI, the
2041 context in which it will appear, how it is intended to be resolved, etc.

2042 It is worth noting that some rare forms of HTTP URIs can result in XRIs that are misleading to the
2043 reader. For example, the following unusual HTTP URI is valid per **[RFC2396]**.

2044 `http://@example.com/example1`

2045 When converted to an XRI, as

2046 `xri://@example.com/example1`

2047 a casual reader could easily misinterpret the “uri-authority” component as an “xri-authority”.
2048 Similarly, a URI with an authority segment like

2049 `http://=bob@example.com/example1`


2050 could be similarly misinterpreted.

2051

2052

Appendix D. Acknowledgments

2053 The editors would also like to thank the following people who participated in the XRI TC and/or
2054 provided input and review of this specification (affiliations listed for OASIS members):
2055

2056 Thomas Bikeev (EAN International), Winston Bumpus (formerly of Novell), James Bryce Clark
2057 (OASIS), Matthey Dovey (Individual), Lars Marius Garshol, Steve Green (Epok), Lance Hood
2058 (Epok), Phillipe LeBlanc (GemPlus), Marc LeMaitre (OneName), Rajeev Maria (Visa
2059 International), Adarbad Master (Epok), John McGarvey (IBM), Davis McPherson (Epok), Mike
2060 Mealling (Verisign), Reva Modi (Infosys), Joseph Moeller (EDS), Brian Nimmo (Epok), Mary
2061 Nishikawa (Individual), Eamonn Neylon (Individual), Masaki Nishitani (NRI), Norman Paskin,
2062 Krishnan Rajagopalan (Novell), Chetan Sabnis (Epok), Jim Schreckengast (formerly of Gemplus),
2063 Tomonori Seki (NRI), Xavier Serret (Gemplus), Terence Spielman (Visa International), Marc
2064 henson (TSO), Geoffrey Strongin (AMD), Bernard Vatant, John Veizades (Visa
2065 International), Bill Washburn (XNSORG), Tetsu Watanabe (NRI), Dave Wentker (Visa
2066 International), Loren West (Epok), and Michael Willett (Wave Systems).

2067

2068 A special acknowledgement to Jerry Kindall (Epok) for a full editorial review.

2069

2070 Also, the authors of and contributors to the following documents and specifications are
2071 acknowledged for the intellectual foundations of the XRI specification:

- 2072 • RFC 1737
- 2073 • RFC 2396 (and RFC 2396bis)
- 2074 • RFC 2616
- 2075 • RFC 2718
- 2076 • RFC 3401-3405 (DDDS)
- 2077 • REST Architecture
- 2078 • IRI – Internationalized Resource Identifiers draft
- 2079 • XNS
- 2080

2081

Appendix E. Revision History

2082

[This appendix should be removed for specifications that are at OASIS Standard level.]

| Rev | Date | By Whom | What |
|---------|------------|--|---|
| wd-01 | 2003-06-24 | Drummond Reed | Initial version to review structure and section 1 with other editors |
| wd-02 | 2003-06-25 | Drummond Reed | Reorganized overall structure and drafted first portion of section 2 |
| wd-03 | 2003-06-30 | Drummond Reed | Reorganized level two headings; edited section 1; drafted all ABNF portions of section 2; added collected ABNF to Appendix A; added Appendix B with initial \$ identifiers; added Appendix C. |
| wd-04 | 2003-07-02 | Dave McAlpin | Editorial changes; new text in 2.2.3.2, 2.4, 2.4.*, 2.5, 2.5.*, 4.* |
| wd-05 | 2003-07-03 | Dave McAlpin | Editorial changes; added resolution text (section 3) |
| wd-06 | 2003-07-03 | Dave McAlpin | Minor edits; removed inline notes and created issues section as Appendix G. |
| wd-07 | 2003-07-24 | Dave McAlpin | Internationalization. Major revisions to 2.2 – 2.5. Harmonization of section 3. |
| wd-08 | 2003-09-30 | Gabe Wachob | New resolution section, addressed many issues |
| wd-09 | 2003-10-28 | All members of XRI Editors SC, Jerry Kindall | See sequence of 09 (a-m) drafts posted to XRI Editors SC. |
| wd-rc-1 | 2003-11-07 | Drummond Reed | Converted all final proofing action items from 09m to comments. |
| wd-rc-2 | 2003-11-19 | All Editors, Drummond Reed | Incorporated all comments and feedback, resolved all action items, updated all comments. |

2083

2084

Appendix F. Notices

2085 OASIS takes no position regarding the validity or scope of any intellectual property or other rights
2086 that might be claimed to pertain to the implementation or use of the technology described in this
2087 document or the extent to which any license under such rights might or might not be available;
2088 neither does it represent that it has made any effort to identify any such rights. Information on
2089 OASIS's procedures with respect to rights in OASIS specifications can be found at the OASIS
2090 website. Copies of claims of rights made available for publication and any assurances of licenses
2091 to be made available, or the result of an attempt made to obtain a general license or permission
2092 for the use of such proprietary rights by implementors or users of this specification, can be
2093 obtained from the OASIS Executive Director.

2094 OASIS invites any interested party to bring to its attention any copyrights, patents or patent
2095 applications, or other proprietary rights which may cover technology that may be required to
2096 implement this specification. Please address the information to the OASIS Executive Director.

2097 **Copyright © OASIS Open 2003. All Rights Reserved.**

2098 This document and translations of it may be copied and furnished to others, and derivative works
2099 that comment on or otherwise explain it or assist in its implementation may be prepared, copied,
2100 published and distributed, in whole or in part, without restriction of any kind, provided that the
2101 above copyright notice and this paragraph are included on all such copies and derivative works.
2102 However, this document itself does not be modified in any way, such as by removing the
2103 copyright notice or references to OASIS, except as needed for the purpose of developing OASIS
2104 specifications, in which case the procedures for copyrights defined in the OASIS Intellectual
2105 Property Rights document must be followed, or as required to translate it into languages other
2106 than English.

2107 The limited permissions granted above are perpetual and will not be revoked by OASIS or its
2108 successors or assigns.

2109 This document and the information contained herein is provided on an "AS IS" basis and OASIS
2110 DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO
2111 ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE
2112 ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A
2113 PARTICULAR PURPOSE.