



Extensible Resource Identifier (XRI) Syntax V2.0

Committee Specification 01

14 November 2005

Specification URIs:

This Version:

<http://docs.oasis-open.org/xri/xri-syntax/2.0/specs/cs01/xri-syntax-V2.0-cs.html>

<http://docs.oasis-open.org/xri/xri-syntax/2.0/specs/cs01/xri-syntax-V2.0-cs.pdf>

<http://docs.oasis-open.org/xri/xri-syntax/2.0/specs/cs01/xri-syntax-V2.0-cs.doc>

Previous Version:

<http://docs.oasis-open.org/xri/xri/V2.0/xri-syntax-V2.0-cd-01.pdf>

Latest Version:

<http://docs.oasis-open.org/xri/2.0/specs/xri-syntax-V2.0.html>

<http://docs.oasis-open.org/xri/2.0/specs/xri-syntax-V2.0.pdf>

<http://docs.oasis-open.org/xri/2.0/specs/xri-syntax-V2.0.doc>

Technical Committee:

OASIS eXtensible Resource Identifier (XRI) TC

Editors:

Drummond Reed, Cordance <drummond.reed@cordance.net>

Dave McAlpin, Epok <dave.mcalpin@epok.net>

Contributors:

Peter Davis, Neustar <peter.davis@neustar.biz>

Nat Sakimura, NRI <n-sakimura@nri.co.jp>

Mike Lindelsee, Visa International <mlindels@visa.com>

Gabe Wachob, Visa International <gwachob@visa.com>

Abstract:

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

Status:

This document was last revised or approved by the XRI Technical Committee on the above date. The level of approval is also listed above. Check the current location noted above for possible later revisions of this document. This document is updated periodically on no particular schedule.

36 Technical Committee members should send comments on this specification to the
37 Technical Committee's email list. Others should send comments to the Technical
38 Committee by using the "Send A Comment" button on the Technical Committee's web
39 page at <http://www.oasis-open.org/committees/xri>.

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

44 The non-normative errata page for this specification is located at [http://www.oasis-](http://www.oasis-open.org/committees/xri)
45 [open.org/committees/xri](http://www.oasis-open.org/committees/xri).

46

47 Table of Contents

48	Introduction	5
49	1.1 Overview of XRIs.....	5
50	1.1.1 Generic Syntax	5
51	1.1.2 URI, URL, URN, and XRI.....	6
52	1.2 Terminology and Notation	6
53	1.2.1 Keywords	6
54	1.2.2 Syntax Notation.....	7
55	2 Syntax	8
56	2.1 Characters.....	8
57	2.1.1 Character Encoding	8
58	2.1.2 Reserved Characters	8
59	2.1.3 Unreserved Characters.....	8
60	2.1.4 Percent-Encoded Characters	9
61	2.1.4.1 Encoding XRI Metadata.....	9
62	2.1.5 Excluded Characters.....	9
63	2.2 Syntax Components.....	10
64	2.2.1 Authority	11
65	2.2.1.1 XRI Authority.....	11
66	2.2.1.2 Global Context Symbol (GCS) Authority	11
67	2.2.1.3 IRI Authority	12
68	2.2.2 Cross-References	12
69	2.2.3 Path.....	13
70	2.2.4 Query	14
71	2.2.5 Fragment.....	14
72	2.3 Transformations	14
73	2.3.1 Transforming XRI References into IRI and URI References.....	14
74	2.3.2 Escaping Rules for XRI Syntax.....	15
75	2.3.3 Transforming IRI References into XRI References	16
76	2.4 Relative XRI References	17
77	2.4.1 Reference Resolution	17
78	2.4.2 Reference Resolution Examples	17
79	2.4.2.1 Normal Examples	17
80	2.4.2.2 Abnormal Examples.....	18
81	2.4.3 Leading Segments Containing a Colon	18
82	2.4.4 Leading Segments Beginning with a Cross-Reference	19
83	2.5 Normalization and Comparison.....	19
84	2.5.1 Case.....	19
85	2.5.2 Encoding, Percent-Encoding, and Transformations	19
86	2.5.3 Optional Syntax.....	19
87	2.5.4 Cross-References	20
88	2.5.5 Canonicalization.....	20
89	3 Security and Data Protection Considerations	21

90	3.1 Cross-References	21
91	3.2 XRI Metadata	21
92	3.3 Spoofing and Homographic Attacks.....	21
93	3.4 UTF-8 Attacks	22
94	3.5 XRI Usage in Evolving Infrastructure	22
95	4 References.....	23
96	4.1 Normative	23
97	4.2 Informative.....	23
98	Appendix A. Collected ABNF for XRI (Normative)	24
99	Appendix B. Transforming HTTP IRIs to XRIs (Non-Normative)	27
100	Appendix C. Glossary.....	28
101	Appendix D. Acknowledgments.....	33
102	Appendix E. Notices	34
103		

104 Introduction

105 1.1 Overview of XRIs

106 Extensible Resource Identifiers (XRIs) provide a standard means of abstractly identifying a
107 resource independent of any particular concrete representation of that resource—or, in the case
108 of a completely abstract resource, of any representation at all.

109 As shown in Figure 1, XRIs build on the foundation established by URIs (Uniform Resource
110 Identifiers) and IRIs (Internationalized Resource Identifiers) as defined by **[URI]** and **[IRI]**,
111 respectively.



112
113 Figure 1: The relationship of XRIs, IRIs, and URIs

114 The IRI specification created a new identifier by extending the unreserved character set to include
115 characters beyond those allowed in generic URIs. It also defined rules for transforming this
116 identifier into a syntactically legal URI. Similarly, this specification creates a new identifier, an
117 XRI, that extends the syntactic elements (but not the character set) allowed in IRIs. To
118 accommodate applications that expect IRIs or URIs, this specification also defines rules for
119 transforming an XRI reference into a valid IRI or URI reference.

120 Although an XRI is not a Uniform Resource Name (URN) as defined in *URN Syntax [RFC2141]*,
121 an XRI consisting entirely of persistent segments is designed to meet the requirements set out in
122 *Functional Requirements for Uniform Resource Names [RFC1737]*.

123 This document specifies the normative syntax for XRIs, along with associated normalization,
124 processing and equivalence rules. See also *An Introduction to XRIs [XRIIntro]* for a non-
125 normative introduction to XRI architecture.

126 1.1.1 Generic Syntax

127 XRI syntax follows the same basic pattern as IRI and URI syntax. A fully-qualified XRI consists of
128 the prefix “xri://” followed by the same four components as a generic authority-based IRI or URI.

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

130 The definitions of these components are, for the most part, supersets of the equivalent
131 components in the generic IRI or URI syntax. One advantage of this approach is that the vast
132 majority of HTTP URIs and IRIs, which derive directly from generic URI syntax, can be
133 transformed to valid XRIs simply by changing the scheme from “http” to “xri”. This transformation
134 is discussed in Appendix B, “Transforming HTTP IRIs to XRIs”.

135 XRI syntax extends generic IRI syntax in the following four ways:

- 136 1. *Persistent and reassignable segments.* Unlike generic URI syntax, XRI syntax allows the
137 internal components of an XRI reference to be explicitly designated as either persistent or
138 reassignable.

- 139 2. *Cross-references*. Cross-references allow XRI references to contain other XRI references
140 or IRIs as syntactically-delimited sub-segments. This provides syntactic support for
141 “compound identifiers”, i.e., the use of well-known, fully-qualified identifiers within the
142 context of another XRI reference. Typical uses of cross-references include using well-
143 known types of metadata in an XRI reference (such as language or versioning metadata),
144 or the use of globally-defined identifiers to mark parts of an XRI reference as having
145 application- or vocabulary-specific semantics.
- 146 3. *Additional authority types*. While XRI syntax supports the same generic syntax used in
147 IRIs for DNS and IP authorities, it also provides two additional options for identifying an
148 authority: a) global context symbols (GCS), shorthand characters used for establishing
149 the abstract global context of an identifier, and b) cross-references, which enable any
150 identifier to be used to specify an XRI authority.
- 151 4. *Standardized federation*. Federated identifiers are those delegated across multiple
152 authorities, such as DNS names. Generic URI syntax leaves the syntax for federated
153 identifiers up to individual URI schemes, with the exception of explicit support for IP
154 addresses. XRI syntax standardizes federation of both persistent and reassignable
155 identifiers at any level of the path.

156 1.1.2 URI, URL, URN, and XRI

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

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

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

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

170 This conclusion is shared by **[URI]** which states in section 1.1.3:

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

175 XRIs are consistent with this philosophy. Although XRIs are designed to fulfill the requirements of
176 abstract “names” that are resolved into concrete locators, XRI syntax does not distinguish
177 between identifiers that represent “names”, “locators” or “characteristics.”

178 1.2 Terminology and Notation

179 1.2.1 Keywords

180 The key words “MUST”, “MUST NOT”, “REQUIRED”, “SHALL”, “SHALL NOT”, “SHOULD”,
181 “SHOULD NOT”, “RECOMMENDED”, “MAY” and “OPTIONAL” in this document are to be
182 interpreted as described in **[RFC2119]**. When these words are not capitalized in this document,
183 they are meant in their natural language sense.

184 1.2.2 Syntax Notation

185 This specification uses the syntax notation employed in **[IRI]**: Augmented Backus-Naur Form
186 (ABNF), defined in **[RFC2234]**. Although the ABNF defines syntax in terms of the US-ASCII
187 character encoding, XRI syntax should be interpreted in terms of the character that the ASCII-
188 encoded octet represents, rather than the octet encoding itself, as explained in **[URI]**. As with
189 URIs, the precise bit-and-byte representation of an XRI reference on the wire or in a document is
190 dependent upon the character encoding of the protocol used to transport it, or the character set of
191 the document that contains it.

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

195 To simplify comparison between generic XRI syntax and generic IRI syntax, the ABNF
196 productions that are unique to XRIs are shown with light green shading, while those inherited
197 from **[IRI]** are shown with light yellow shading.

198 | This is an example of ABNF specific to XRI.

199 ; This is an example of ABNF inherited from IRI.

200 Lastly, because the prefix “xri://” is optional in absolute XRIs that use a global context symbol
201 (see section 2.2.1.2), some example XRIs are shown without this prefix.

202 2 Syntax

203 This section defines the normative syntax for XRIs. Note that additional constraints are inherited
204 from **[IRI]** and **[URI]**, as defined in section 2.2. Also note that some productions in the XRI ABNF
205 are ambiguous. As with IRIs and URIs, a “first-match-wins” rule is used to disambiguate
206 ambiguous productions. See **[URI]** for more details.

207 2.1 Characters

208 XRI character set and encoding are inherited from **[IRI]**, which is a superset of generic URI
209 syntax as defined in **[URI]**.

210 2.1.1 Character Encoding

211 The standard character encoding of XRI is UTF-8, as recommended by **[RFC2718]**. When an XRI
212 reference is presented as a human-readable identifier, the representation of the XRI reference in
213 the underlying document may use the character encoding of the underlying document. However,
214 this representation must be converted to UTF-8 before the XRI can be processed outside the
215 document. This encoding in UTF-8 **MUST** include normalization according to Normalization Form
216 KC (NFKC) as defined in **[UTR15]**. The stricter NFKC is specified rather than Normalization Form
217 C (NFC) used in IRI encoding **[IRI]** because NFKC reduces the number of UCS compatibility
218 characters allowed in an XRI and increases the probability of equivalence matches.

219 2.1.2 Reserved Characters

220 The overall XRI reserved character set is the same as the reserved character set defined by
221 **[URI]** and **[IRI]**. Due to the extended syntax of XRIs, however, the allocation of reserved
222 characters between the “general delimiters” and “sub-delimiters” productions is different. Those
223 characters that have defined semantics in generic XRI syntax appear in the xri-gen-delims
224 production. Those characters that do not have defined semantics but that are reserved for use as
225 implementation-specific delimiters appear in the xri-sub-delims production. The rgcs-char
226 production that appears in xri-gen-delims below is discussed in section 2.2.1.2.

```
227 xri-reserved = xri-gen-delims / xri-sub-delims
228 xri-gen-delims = ":" / "/" / "?" / "#" / "[" / "]" / "(" / ")"
229 / "*" / "!" / rgcs-char
230 xri-sub-delims = "&" / ";" / "," / "'"
```

231 If an XRI reserved character is used as a data character and not as a delimiter, the character
232 **MUST** be percent-encoded per the rules in section 2.1.4, “Percent-Encoded Characters”. XRI
233 references that differ in the percent-encoding of a reserved character are not equivalent.

234 2.1.3 Unreserved Characters

235 The characters allowed in XRI references that are not reserved are called unreserved. XRI has
236 the same set of unreserved characters as the "iunreserved" production in **[IRI]**.

```
237 iunreserved = ALPHA / DIGIT / "-" / "." / "_" / "~" / uchar
```



```
238 |      uchar          = %xA0-D7FF / %xF900-FDCF / %xFDF0-FFEF
239 |                    / %x10000-1FFFD / %x20000-2FFFD / %x30000-3FFFD
240 |                    / %x40000-4FFFD / %x50000-5FFFD / %x60000-6FFFD
241 |                    / %x70000-7FFFD / %x80000-8FFFD / %x90000-9FFFD
242 |                    / %xA0000-AFFFD / %xB0000-BFFFD / %xC0000-CFFFD
243 |                    / %xD0000-DFFFD / %xE1000-EFFFD
```

244 Percent-encoding unreserved characters in an XRI does not change what resource is identified
245 by that XRI. However, it may change the result of an XRI comparison (see section 2.5,
246 “Normalization and Comparison”), so unreserved characters SHOULD NOT be percent-encoded.

247 2.1.4 Percent-Encoded Characters

248 XRIs follow the same rules for percent-encoding as IRIs and URIs. That is, any *data* character in
249 an XRI reference MUST be percent-encoded if it does not have a representation using an
250 unreserved character but SHOULD NOT be percent-encoded if it does have a representation
251 using an unreserved character. Delimiters in an XRI reference that have a representation using a
252 reserved character MUST NOT be percent-encoded.

253 An XRI reference thus percent-encoded is said to be in *XRI-normal form*. Not all XRI references
254 in XRI-normal form are syntactically legal IRI or URI references. Rules for converting an XRI
255 reference to a valid IRI or URI reference are discussed in section 2.3.1. An XRI reference is in
256 XRI-normal form if it is minimally percent-encoded and matches the ABNF provided in this
257 document, but it is a valid IRI or URI reference only after it is percent-encoded according to the
258 transformation described in section 2.3.1.

259 A percent-encoded octet is a character triplet consisting of the percent character “%” followed by
260 the two hexadecimal digits representing that octet’s numeric value.

```
261 |      pct-encoded    = "%" HEXDIG HEXDIG
```

262 The uppercase hexadecimal digits “A” through “F” are equivalent to the lowercase digits “a”
263 through “f”, respectively. XRI references that differ only in the case of hexadecimal digits used in
264 percent-encoded octets are equivalent. For consistency, XRI generators and normalizers
265 SHOULD use uppercase hexadecimal digits for percent-encoded triplets.

266 Note that a % symbol used to represent itself in an XRI reference (i.e., as data and not to
267 introduce a percent-encoded triplet) must be percent-encoded.

268 2.1.4.1 Encoding XRI Metadata

269 In some cases, the transformation of an identifier in its native language and display format into an
270 XRI reference in XRI-normal form may lose information that cannot be retained through percent-
271 encoding. For example, in certain languages, displaying the glyph of a UTF-8 encoded character
272 requires additional language and font information not available in UTF-8. The loss of this
273 information during UTF-8 encoding might cause the resulting XRI to be ambiguous.

274 XRI syntax offers an option for encoding this language metadata using a cross-reference
275 beginning with the GCS “\$” symbol (see section 2.2.1.2). The top level authority for language
276 metadata is the *XRI Metadata Specification* published by the OASIS XRI Technical Committee.

277 2.1.5 Excluded Characters

278 Certain characters, such as “space”, are excluded from XRI syntax and must be percent-encoded
279 in order to be represented within an XRI. Systems responsible for accepting or presenting XRI
280 references may choose to percent-encode excluded characters on input and/or decode them
281 prior to display, as described in section 2.1.4. A string that contains these characters in a non-
282 percent-encoded form, however, is not a valid XRI.

283 Note that presenting “space” or other whitespace characters in a non-percent-encoded form is not
284 recommended for several reasons. First, it is often difficult to visually determine the number of
285 spaces or other characters composing a block of whitespace, leading to transcription errors.
286 Second, the space character is often used to delimit an XRI reference, so non-percent-encoded
287 whitespace characters can make it difficult or impossible to determine where the identifier ends.
288 Finally, non-percent-encoded whitespace can be used to maliciously construct subtly different
289 identifiers intended to mislead the reader. For these reasons, non-percent-encoded whitespace
290 characters SHOULD be avoided in presentation, and alternatives to whitespace as a logical
291 separator within XRIs (such as dots or hyphens) SHOULD be used whenever possible.

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

294 “The UCS contains many areas of characters for which there are strong visual
295 look-alikes. Because of the likelihood of transcription errors, these also should be
296 avoided. This includes the full-width equivalents of Latin characters, half-width
297 Katakana characters for Japanese, and many others. This also includes many
298 look-alikes of ‘space’, ‘delims’, and ‘unwise’, characters excluded in **[RFC3491]**.”

299 “Additional information is available from **[UniXML]**. **[UniXML]** is written in the
300 context of running text rather than in the context of identifiers. Nevertheless, it
301 discusses many of the categories of characters not appropriate for IRIs.”

302 Finally, although they are not excluded characters, special care should be taken by user agents
303 with regard to the display of UCS characters that are visual look-alikes (homographs) for XRI
304 delimiters (all characters in the xri-reserved production, section 2.1.2). See section 3.3, “Spoofing
305 and Homographic Attacks” for additional information.

306 2.2 Syntax Components

307 XRI syntax builds on generic IRI (and ultimately, URI) syntax. However because XRI syntax
308 includes syntactic elements other than those defined in **[IRI]** and **[URI]**, this specification defines
309 a new protocol element, “XRI”, along with rules for transforming XRI references into generic IRI or
310 URI references for applications that expect them (see section 2.3.1, “Transforming XRI
311 References into IRI and URI References”). An XRI reference MUST be constructed such that it
312 qualifies as a valid IRI as defined by **[IRI]** when converted to IRI-normal form and such that it
313 qualifies as a valid URI as defined by **[URI]** when converted to URI-normal form.

314 As with URIs, an XRI must be in absolute form, while an XRI reference may be either an XRI or a
315 relative XRI reference.

```
316 XRI = [ "xri://" ] xri-hier-part [ "?" iquery ]  
317 [ "#" ifragment ]  
  
318 xri-hier-part = ( xri-authority / iauthority ) xri-path-abempty  
  
319 XRI-reference = XRI / relative-XRI-ref  
  
320 absolute-XRI = [ "xri://" ] xri-hier-part [ "?" iquery ]  
  
321 relative-XRI-ref = relative-XRI-part [ "?" iquery ] [ "#" ifragment ]  
  
322 relative-XRI-part = xri-path-absolute  
323 / xri-path-noscheme  
324 / ipath-empty  
  
325 xri-value = xri-no-scheme / relative-XRI-ref  
  
326 xri-no-scheme = xri-hier-part [ "?" iquery ] [ "#" ifragment ]
```

327 An XRI begins with an optional prefix “xri://” followed by the same set of hierarchical components
328 as a URI – authority, path, query, and fragment. An XRI is always in absolute form. A relative XRI
329 reference consists of an XRI path followed by an optional XRI query and optional XRI fragment.
330 The absolute-XRI production is provided for contexts that require an XRI in absolute form but that
331 do not allow the fragment identifier.

332 Finally, in certain contexts where XRIs are used exclusively, the prefix “xri://” is redundant. These
333 contexts can use the xri-value production, which includes all levels of XRI paths.

334 2.2.1 Authority

335 XRIs support the same types of authorities as generic IRIs, called *IRI authorities*. XRIs also
336 support additional types of abstract identification authorities called *XRI authorities*.

337 2.2.1.1 XRI Authority

338 There are two ways to express an XRI authority: using a global context symbol (GCS), or using a
339 cross-reference (abbreviated in the ABNF as *xref*). Cross-references are covered in section 2.2.2.

```
340 | xri-authority = gcs-authority / xref-authority
```

341 2.2.1.2 Global Context Symbol (GCS) Authority

342 XRIs offer a simple, compact syntax for indicating the logical global context of an identifier: a
343 single prefix character called a *global context symbol*.

```
344 | gcs-authority = pgcs-authority / rgcs-authority  
345 | pgcs-authority = "!" xri-subseg-pt-nz *xri-subseg  
346 | rgcs-authority = rgcs-char xri-segment  
347 | rgcs-char = "=" / "@" / "+" / "$"
```

348 The global context symbol characters were selected from the set of symbol characters that are
349 valid in a URI under **[URI]**. The bang character, “!”, which is used uniformly in XRI syntax to
350 indicate a persistent identifier segment, serves as the GCS character for global persistent
351 identifiers. The other GCS characters may be used to indicate the global context of either a
352 persistent or a reassignable identifier as shown in Table 1 below:

353

Symbol Character	Authority Type	Establishes Global Context For
=	Person	Identifiers for whom the authority is controlled by an individual person.
@	Organization	Identifiers for whom the authority is controlled by an organization or a resource in an organizational context.
+	General public	Identifiers for whom there is no specific controlling authority because they represent generic dictionary concepts or “tags” whose meaning is determined by consensus. (In the English language, for example, these would be the generic nouns.)
\$	Standards body	Identifiers for whom the authority is controlled by a specification from a standards body, for example, other XRI specifications from the OASIS XRI Technical Committee, other OASIS specifications, or (using cross-references) other standards bodies.

355

Table 1: XRI global context symbols.

356 2.2.1.3 IRI Authority

357 XRIs support the same type of authority defined by the “iauthority” production of **[IRI]**.

```

358 iauthority      = [ iuserinfo "@" ] ihost [ ":" port ]
359 iuserinfo       = *( iunreserved / pct-encoded / sub-delims / ":" )
360 ihost           = IP-literal / IPv4address / ireg-name
361 port           = *DIGIT

```

362 The syntax is inherited directly from **[IRI]**. First, the “iuserinfo” sub-component permits the
363 identification of a user in the context of a host. Next, the “ihost” sub-component has three options
364 for identifying the host: a registered name (such as a domain name), an IPv4 address, or an IPv6
365 literal.

366 A host identifier can be followed by an optional port number. The XRI syntax specification does
367 not define a default port because it is expected this will be inherited from the resolution protocol.
368 Therefore, if the port is omitted in an XRI, it is undefined.

369 Note that authority segments that begin with GCS characters or cross-references (see below)
370 may match both the “iauthority” and the “xri-authority” productions. For instance, “!1”,
371 “@example”, “=example”, “+example”, “\$example” and “(=example)” all match both productions.
372 As with all XRI syntax, the “first-match-wins” rule is used to resolve ambiguities. Consequently, all
373 the examples listed above would be considered XRI authorities, not IRI authorities.

374 2.2.2 Cross-References

375 Cross-references are the primary extensibility mechanism in XRI. They allow an identifier
376 assigned in one context to be reused in another context, permitting identifiers to be shared across
377 contexts. This simplifies identifying logically equivalent resources across hierarchies (a directory
378 concept referred to as “polyarchy”).

379 A cross-reference is syntactically delimited by enclosing it in parentheses, similar to the way an
380 IPv6 literal is encapsulated in square brackets as specified in [RFC2732]. A cross-reference may
381 contain either an XRI reference or an absolute IRI.

```
382 xref = "( ( XRI-reference / IRI ) )"
```

383 It is important that the value of a cross-reference be syntactically unambiguous, whether it is an
384 absolute IRI or one of the various forms of an XRI reference. Therefore special attention must be
385 paid to relative XRI references to avoid ambiguity, as discussed in section 2.4.3.

386 A cross-reference may appear at any node of any XRI except within an IRI authority segment. A
387 cross-reference as the very first sub-segment in an XRI is a valid top-level XRI authority.

```
388 xref-authority = xref *xri-subseg
```

389 This syntax allows any globally-unique identifier in any URI scheme (e.g., an HTTP URI, mailto
390 URI, URN etc.) to specify a global XRI authority.

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

393 2.2.3 Path

394 As with IRIs, the XRI path component is a hierarchal sequence of path segments separated by
395 slash ("/") characters and terminated by the first question-mark ("?") or number sign ("#")
396 character, or by the end of the XRI reference. But while an IRI path segment is considered
397 opaque by a generic URI processor, an XRI path segment can be parsed by an XRI processor
398 into two types of sub-segments: **segments* (pronounced "star segments") and *!segments*
399 (pronounced "bang segments").

```
400 xri-path = xri-path-abempty  
401           / xri-path-absolute  
402           / xri-path-noscheme  
403           / ipath-empty  
  
404 xri-path-abempty = *( "/" xri-segment )  
  
405 xri-path-absolute = "/" [ xri-segment-nz *( "/" xri-segment ) ]  
  
406 xri-path-noscheme = xri-subseg-od-nx *xri-subseg-nc  
407                   *( "/" xri-segment )  
  
408 xri-segment = xri-subseg-od *xri-subseg  
  
409 xri-segment-nz = xri-subseg-od-nz *xri-subseg  
  
410 xri-subseg = ( "*" / "!" ) (xref / *xri-pchar)  
  
411 xri-subseg-nc = ( "*" / "!" ) (xref / *xri-pchar-nc)  
  
412 xri-subseg-od = [ "*" / "!" ] (xref / *xri-pchar)  
  
413 xri-subseg-od-nz = [ "*" / "!" ] (xref / 1*xri-pchar)  
  
414 xri-subseg-od-nx = [ "*" / "!" ] 1*xri-pchar-nc  
  
415 xri-subseg-pt-nz = "!" (xref / 1*xri-pchar)
```

416 * segments are used to specify *reassignable identifiers*—identifiers that may be reassigned by an
417 identifier authority to represent a different resource at some future date. ! segments are used to
418 specify *persistent identifiers*—identifiers that are permanently assigned to a resource and will not
419 be reassigned at a future date. A ! segment SHOULD meet the requirements for persistent
420 identifiers set out in *Functional Requirements for Uniform Resource Names [RFC1737]*. The
421 default is a * segment, so a leading star (“*”) is optional for the first (or only) sub-segment if this
422 subsegment is reassignable.

423 An XRI path segment may contain the same characters as a URI path segment plus the
424 expanded UCS character set inherited from [IRI]. If a star (“*”) or bang (“!”) appears in a path of
425 an XRI reference, it will be interpreted as a sub-segment delimiter. If this interpretation is not
426 desired for these characters, or for any other special XRI delimiters, these characters MUST be
427 percent-encoded when they appear in the path segment. See section 2.1.4, “Percent-Encoded
428 Characters”.

```
429 xri-pchar      = iunreserved / pct-encoded / xri-sub-delims / ":"  
430 xri-pchar-nc   = iunreserved / pct-encoded / xri-sub-delims
```

431 With the exception of star (“*”), bang (“!”) and cross-reference delimiters, an XRI path segment is
432 considered opaque by generic XRI syntax. As with IRIs, XRI extensions or generating
433 applications may define special meanings for other XRI reserved characters for the purpose of
434 delimiting extension-specific or generator-specific sub-components.

435 2.2.4 Query

436 The XRI query component is identical to the IRI query component as described in section 2.2 of
437 [IRI].

```
438 iquery = *( ipchar / iprivate / "/" / "?" )
```

439 2.2.5 Fragment

440 XRI syntax also supports fragments as described in section 2.2 of [IRI].

```
441 ifragment = *( ipchar / "/" / "?" )
```

442 Since XRI federation syntax can inherently address attributes or sub-resources to any depth,
443 fragments are supported primarily for compatibility with generic URI syntax. XRIs can also employ
444 cross-references to identify media types or other alternative representations of a resource. See
445 section 2.2.2.

446 2.3 Transformations

447 2.3.1 Transforming XRI References into IRI and URI References

448 Although XRIs are intended to be used by applications that understand them natively, it may also
449 be desirable to use them in contexts that do not recognize an XRI reference but that allow an IRI
450 reference as described in [IRI], or a fully-conformant URI reference as defined by [URI].

451 This section specifies the steps for transforming an XRI reference into a valid IRI reference. At
452 the completion of these steps, the XRI reference is in *IRI-normal form*. An XRI reference in IRI-
453 normal form may then be mapped into a valid URI reference by following the algorithms defined
454 in section 3.1 of [IRI]. After that mapping, the XRI reference is in *URI-normal form*.

455 Applications transforming XRI references to IRI references MUST use the following steps (or a
456 process that achieves exactly the same result). Before applying these steps, the XRI reference
457 must be in XRI-normal form as defined in section 2.1.4.

- 458 1. If the XRI reference is not encoded in UTF-8, convert the XRI reference to a sequence of
459 characters encoded in UTF-8, normalized according to Normalization Form KC (NFKC)
460 as defined in **[UTR15]**.
- 461 2. If the XRI reference is not relative (i.e., if it matches the “XRI” ABNF production) and the
462 optional “xri://” prefix has been omitted, prepend “xri://” to the XRI reference.
- 463 3. Optionally add XRI metadata using cross-references as defined in section 2.1.4.1. Note
464 that the addition of XRI metadata may change the resulting IRI or URI reference for the
465 purposes of comparison as explained in section 2.5.4.
- 466 4. Apply the XRI escaping rules defined in section 2.3.2. Note that this step is not
467 idempotent (i.e., it may yield a different result if applied more than once), so it is very
468 important that implementers not apply this step more than once to avoid changing the
469 semantics of the identifier.

470 At the completion of step 4, the percent-encoded XRI reference is now in IRI-normal form and
471 may be used as an IRI reference conformant with **[IRI]**.

472 Applying this conversion does not change the equivalence of the identifier, with the possible
473 exception of the addition of XRI metadata as discussed in Step 3.

474 In general, an application SHOULD use the least-transformed version appropriate for the context
475 in which the identifier appears. For example, if the context allows an XRI reference directly, the
476 identifier SHOULD be an XRI reference in XRI-normal form as described in section 2.1.4. If the
477 context allows an IRI reference but not an XRI reference, the identifier SHOULD be in IRI-normal
478 form. Only when the context allows neither XRI nor IRI references should URI-normal form be
479 used.

480 **2.3.2 Escaping Rules for XRI Syntax**

481 This section defines rules for preventing misinterpretation of XRI syntax when an XRI reference is
482 evaluated by a non-XRI-aware parser.

483 The first rule deals with cross-references as explained in section 2.2.2. Since a cross-reference
484 contains either an IRI or an XRI reference (which itself may contain further nested IRIs or XRI
485 references), it may include characters that, if not escaped, would cause misinterpretation when
486 the XRI reference is used in a context that expects an IRI or URI reference. Consider the
487 following XRI:

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

489 The generic parsing algorithm described in **[URI]** would separate the above XRI into the following
490 components:

```
491 scheme = xri  
492 authority = @example  
493 path = /(xri://@example2/abc  
494 query = id=1)
```

495 The desired separation is:

```
496 scheme = xri  
497 authority = @example  
498 path = /(xri://@example2/abc?id=1)  
499 query = <undefined>
```

500 To avoid this type of misinterpretation, certain characters in a cross-reference must be percent-
501 encoded when transforming an XRI reference into IRI-normal form. In particular, the question
502 mark (“?”) character must be percent-encoded as “%3F” and the number sign “#” character must
503 be percent-encoded as “%28”.

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

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

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

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

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

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

512 instead of the intended:

```
513 xri://@example.com/
```

514 This is because the “merge” algorithm in section 5.2.3 of [URI] is defined in terms of the last
515 (right-most) slash character. This problem is avoided by encoding slashes within cross-references
516 as “%2F”. Following this rule, the above example would be expressed as:

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

518 Ambiguity is also possible if an XRI reference in XRI-normal form contains characters that have
519 been percent-encoded to indicate that they should not be interpreted as delimiters. For example,
520 consider the following XRI in XRI-normal form:

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

522 This slash character between “c” and “d” is percent-encoded to show that it’s not a syntactical
523 element of the XRI, i.e., that it should be interpreted as data and not as a delimiter. To preserve
524 this type of distinction when converting an XRI reference to an IRI reference, the percent “%”
525 character must be percent-encoded as “%25”. Following this rule, the above example fully
526 converted would be:

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

528 To summarize, the following four special rules MUST be applied during step 4 of section 2.3.1.
529 Before applying these rules, the XRI reference MUST be in XRI-normal form and all IRIs in cross-
530 references MUST be in a percent-encoded form appropriate to their schemes.

- 531 1. Percent-encode all percent “%” characters as “%25” across the entire XRI reference.
- 532 2. Percent-encode all number sign “#” characters that appear within a cross-reference as
533 “%23”.
- 534 3. Percent-encode all question mark “?” characters that appear within a cross-reference as
535 “%3F”.
- 536 4. Percent-encode all slash “/” characters that appear within a cross-reference as “%2F”.

537 2.3.3 Transforming IRI References into XRI References

538 Transformation of an XRI reference in IRI-normal form into an XRI reference in XRI-normal form
539 MUST use the following steps (or a process that achieves the same result).

- 540 1. If the XRI reference is not encoded in UTF-8, convert the XRI reference to a sequence of
541 characters encoded in UTF-8, normalized according to Normalization Form KC (NFKC)
542 as defined in [UTR15].

- 543 2. Perform the following special conversions for XRI syntax:
- 544 a. Convert all percent-encoded slash (“/”) characters to their corresponding octets.
- 545 b. Convert all percent-encoded question mark (“?”) characters to their
- 546 corresponding octets.
- 547 c. Convert all percent-encoded number sign (“#”) characters to their corresponding
- 548 octets.
- 549 d. Convert all percent-encoded percent (“%”) characters to their corresponding
- 550 octets.

551 Note that this process is not idempotent (i.e., it may yield a different result if applied more than

552 once), so it is very important that implementers only apply this process to XRI references in IRI-

553 normal form. If it is applied to an XRI reference in XRI-normal form, the resulting identifier may

554 not be equivalent to the XRI reference before transformation.

555 2.4 Relative XRI References

556 2.4.1 Reference Resolution

557 For XRI references in IRI-normal form or URI-normal form, resolving a relative XRI reference into

558 an absolute XRI reference is straightforward. If the base XRI and the relative XRI reference are in

559 IRI-normal form, section 6.5 of **[IRI]** applies. If the base XRI and the relative XRI reference are in

560 URI-normal form, section 5 of **[URI]** applies.

561 It is important that XRI references appear in a form appropriate to their context (i.e., in URI-

562 normal form in contexts that expect URI references and in IRI-normal form in contexts that expect

563 IRI references), since the algorithms described in **[IRI]** and **[URI]** may produce incorrect results

564 when applied to XRI references in XRI-normal form, particularly when those XRI references

565 contain cross-references.

566 In contexts that allow a native XRI reference (i.e., an XRI reference in XRI-normal form), it may

567 be useful to perform relative reference resolution without first converting to IRI- or URI-normal

568 form. In fact, it may be difficult or impossible to convert to IRI- or URI-normal form without first

569 resolving the relative XRI reference to an absolute XRI. The algorithms described in section 5 of

570 **[URI]** apply to XRI references in XRI-normal form provided that the processor:

- 571 • treats the characters allowed in IRI references but not in URI references the same as it
- 572 treats unreserved characters in URI references (as required by section 5 of **[IRI]**) and
- 573 • treats all characters within all cross-references the same as unreserved characters in URI
- 574 references (i.e., treats cross-references as opaque with respect to relative reference
- 575 resolution).

576 2.4.2 Reference Resolution Examples

577 The following are examples of relative XRI reference resolution. These examples are very similar

578 to the examples for resolving relative references in **[URI]**. Starting with the following base XRI in

579 XRI-normal form:

580 `xri://@a*a/!b!b/c*c/(xri://@d*d/e)?q`

581 a relative reference is transformed to its target XRI as shown in the following examples.

582 2.4.2.1 Normal Examples

583	<code>!g!g</code>	=	<code>xri://@a*a/!b!b/c*c/!g!g</code>
584	<code>./!g!g</code>	=	<code>xri://@a*a/!b!b/c*c/!g!g</code>
585	<code>!g!g/</code>	=	<code>xri://@a*a/!b!b/c*c/!g!g/</code>
586	<code>!/g!g</code>	=	<code>xri://@a*a/!g!g</code>
587	<code>//@!g!g</code>	=	Not a legal relative XRI reference
588	<code>?y</code>	=	<code>xri://@a*a/!b!b/c*c/(xri://@d*d/e)?y</code>
589	<code>!g!g?y</code>	=	<code>xri://@a*a/!b!b/c*c/!g!g?y</code>

```

590 #s = xri://@a*a/!b!b/c*c/(xri://@d*d/e)?q#s
591 !g!g#s = xri://@a*a/!b!b/c*c/!g!g#s
592 !g!g?y#s = xri://@a*a/!b!b/c*c/!g!g?y#s
593 !x = xri://@a*a/!b!b/c*c/!x
594 !g!g;x = xri://@a*a/!b!b/c*c/!g!g;x
595 !g!g;x?y#s = xri://@a*a/!b!b/c*c/!g!g;x?y#s
596 = xri://@a*a/!b!b/c*c/(xri://@d*d/e)?q
597 . = xri://@a*a/!b!b/c*c/
598 ./ = xri://@a*a/!b!b/c*c/
599 .. = xri://@a*a/!b!b/
600 ../ = xri://@a*a/!b!b/
601 ../!g!g = xri://@a*a/!b!b/!g!g
602 ../.. = xri://@a*a/
603 ../.. / = xri://@a*a/
604 ../.. /!g!g = xri://@a*a/!g!g

```

2.4.2.2 Abnormal Examples

As in IRIs and URIs, the ".." syntax cannot be used to change the authority component of an XRI.

```

607 ../../..!g!g = xri://@a*a/!g!g
608 ../../.. /!g!g = xri://@a*a/!g!g

```

As in IRIs and URIs, "." and ".." have a special meaning only when they appear as complete path segments.

```

611 ../!g!g = xri://@a*a/!g!g
612 ../!g!g = xri://@a*a/!g!g
613 !g!g. = xri://@a*a/!b!b/c*c/!g!g.
614 !g!g = xri://@a*a/!b!b/c*c/!g!g
615 !g!g.. = xri://@a*a/!b!b/c*c/!g!g..
616 ..!g!g = xri://@a*a/!b!b/c*c/..!g!g

```

XRI parsers, like IRI and URI parsers, must be prepared for superfluous or nonsensical uses of "." and "..".

```

619 ../!g!g = xri://@a*a/!b!b/!g!g
620 ../!g!g/. = xri://@a*a/!b!b/c*c/!g!g/
621 !g!g/./h = xri://@a*a/!b!b/c*c/!g!g/h
622 !g!g/.. /h = xri://@a*a/!b!b/c*c/h
623 !g!g;x=1 /./y = xri://@a*a/!b!b/c*c/!g!g;x=1/y
624 !g!g;x=1 /.. /y = xri://@a*a/!b!b/c*c/y

```

XRI parsers, like IRI and URI parsers, must take care to separate the reference's query and/or fragment components from the path component before merging it with the base path and removing dot-segments.

```

628 !g!g?y /./x = xri://@a*a/!b!b/c*c/!g!g?y /./x
629 !g!g?y /.. /x = xri://@a*a/!b!b/c*c/!g!g?y /.. /x
630 !g!g#s /./x = xri://@a*a/!b!b/c*c/!g!g#s /./x
631 !g!g#s /.. /x = xri://@a*a/!b!b/c*c/!g!g#s /.. /x

```

2.4.3 Leading Segments Containing a Colon

[URI] points out that relative URI references with an initial segment containing a colon may be subject to misinterpretation:

"A path segment that contains a colon character (e.g., 'this:that') cannot be used as the first segment of a relative-path reference because it would be mistaken for

637 a scheme name. Such a segment must be preceded by a dot-segment (e.g.,
638 './this:that') to make a relative-path reference.”

639 Relative XRI references can be similarly misinterpreted. If any segment prior to the first slash (“/”)
640 character in a relative XRI reference contains a colon, the relative XRI reference must be
641 rewritten to begin either with “*”, if appropriate, or “./”. Thus, “a:b” becomes either “*a:b” or “./a:b”.

642 **2.4.4 Leading Segments Beginning with a Cross-Reference**

643 A path segment that begins with a cross-reference cannot be used as the first segment of a
644 relative reference because it would be mistaken for an xref-authority. As with a leading segment
645 containing a colon, such a segment must be preceded with either a “*” or a “./” to make it a
646 relative XRI reference.

647 **2.5 Normalization and Comparison**

648 In general, the normalization and comparison rules for generic IRIs and URIs specified in Section
649 5 of [IRI] and Section 6 of [URI] apply to XRIs. This section describes a number of additional XRI-
650 specific rules for normalization and comparison. To reduce the requirements imposed upon a
651 minimally conforming processor, the majority of these rules are RECOMMENDED rather than
652 REQUIRED. An implementation that fails to observe them, however, may frequently treat two
653 XRIs as non-equal when in fact they are equal.

654 Each application that uses XRI references MAY define additional equivalence rules as
655 appropriate. Due to the level of abstraction XRIs provide, such higher-order equivalence rules
656 may be based on indirect comparisons or specified XRI-to-XRI mappings (for example, mappings
657 of reassignable XRIs to persistent XRIs).

658 **2.5.1 Case**

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

- 660 • Comparison of the scheme component of XRIs and all IRIs used as cross-references is case-
661 insensitive.
- 662 • Comparison of authority components (section 2.2.1) is case-insensitive as defined in [IRI].
- 663 • As specified in section 2.1.4, comparison of characters in a percent-encoding construction is
664 case-insensitive for the hexadecimal digits “A” through “F”, i.e. “%ab” is equivalent to “%AB”.

665 **2.5.2 Encoding, Percent-Encoding, and Transformations**

- 666 • Two XRIs MUST be considered equivalent if they are character-for-character equivalent.
667 Therefore, they are also equivalent if they are byte-for-byte equivalent and use the same
668 character encoding.
- 669 • Two XRIs that differ only in whether unreserved characters are percent-encoded SHOULD be
670 considered equivalent. If one XRI percent-encodes one or more unreserved characters, and
671 another XRI differs only in that the same characters are not percent-encoded, they are
672 equivalent.
- 673 • All forms of an XRI during the transformation process described in section 2.3.1 SHOULD be
674 considered equivalent, assuming the same XRI metadata is inserted as described in section
675 2.3.1.

676 **2.5.3 Optional Syntax**

- 677 • An “xri-segment” (section 2.2.3) that omits the optional leading star (“*”) SHOULD be
678 considered equivalent to the same “xri-segment” prefixed with an star. For example the
679 segment “/foo*bar” is equivalent to the segment “/*foo*bar”.

680 2.5.4 Cross-References

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

```
684 xri://example/(+example/(+foo))  
685 xri://example/(+Example/(+FOO))
```

- 686 • While cross-references beginning with the GCS “\$” symbol MAY be considered significant in
687 all cases, the specification governing a particular \$ namespace MAY declare that cross-
688 references in that namespace should be ignored for purposes of comparison. Failure to follow
689 such a rule may lead to false negatives. See section 2.1.4.1.

690 2.5.5 Canonicalization

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

696 It is nevertheless useful to define guidelines for making XRI references reasonably canonical. XRI
697 references that follow these guidelines will be more consistent in presentation, simpler to process,
698 less prone to false-negative comparisons, and more easily cached. To that end, unless there is a
699 compelling reason to do otherwise, XRI references SHOULD be provided in a form in which:

- 700 • The optional “xri://” prefix is included,
701 • The scheme is specified in lowercase,
702 • The authority component is specified in lowercase,
703 • Percent-encoding uses uppercase A through F,
704 • If optional, the leading star in xri-segments is omitted,
705 • Unnecessary percent-encoding is not present,
706 • ./ and ../ are absent in absolute XRIs, and
707 • Cross-references are reasonably canonical with respect to their schemes.

708 Table 2 illustrates the application of these rules. Although the XRIs in the first and second
709 columns are equivalent, the form in the second column is recommended.
710

Avoid	Recommended	Comment
@example	xri://example	Add optional “xri://”
XRI://example	xri://example	Lowercase “xri”
xri://Example	xri://example	Lowercase authority
xri://example%2f	xri://example%2F	Uppercase percent-encoding
xri://example/*abc	xri://example/abc	Remove optional leading star
xri://ex%61mple	xri://example	Remove unnecessary percent-encoding
xri://example/./abc	xri://example/abc	Avoid ./ and ../ in absolute XRIs

711

Table 2: Examples of XRI canonicalization recommendations.

712 3 Security and Data Protection Considerations

713 To a great extent, XRI syntax has the same security considerations as [IRI] and [URI]. In
714 particular the material in [URI], section 7, *Security Considerations*, includes a discussion of the
715 following topics:

- 716 • Reliability and Consistency
- 717 • Malicious Construction
- 718 • Back-End Transcoding
- 719 • Rare IP Address Formats
- 720 • Sensitive Information
- 721 • Semantic Attacks

722 This material notes that “a URI does not in itself pose a direct security threat.” The same is true
723 of an XRI. However infrastructure and applications that use XRIs may have special security and
724 data protection considerations as noted in this section.

725 3.1 Cross-References

726 Since cross-references in an XRI can reference other URI schemes, implementation must
727 carefully consider the relevant security considerations for those referenced schemes.

728 3.2 XRI Metadata

729 The use of cross-references employing the GCS “\$” symbol for encoding XRI metadata in an XRI
730 (section 2.1.4.1) may involve other security and data protection considerations that are outside
731 the scope of this specification. These considerations SHOULD be addressed in the relevant \$
732 namespace specification.

733 3.3 Spoofing and Homographic Attacks

734 One particularly important security consideration is spoofing, covered first in [URI] and more
735 thoroughly in [IRI] Section 7.5. Spoofing is a semantic attack in which an identifier is deliberately
736 constructed to deceive the user into believing it represents one resource when in fact it
737 represents another. With IRIs in particular, a common example of such an attack is using
738 characters from different scripts that are visual lookalikes (“homographs”), e.g., the Latin “A”, the
739 Greek “Alpha”, and the Cyrillic “A”. Another common attack is using homographs of the delimiter
740 character “/” to deceive the user about the true contents of an IRI authority segment.

741 Spoofing has already been used extensively in email “phishing” attacks. As more browsers add
742 support for Internationalized Domain Names (IDN), it is also beginning to appear in online Web
743 links (“pharming”). Not only are some users less suspicious of URIs on the Web, but the attacker
744 may even obtain a corresponding SSL/TLS certificate for the deceptive URI or IRI to make the
745 fraudulent site look completely secure and legitimate.

746 To help prevent this problem, XRI registries SHOULD institute policies preventing the registration
747 of deceptive XRIs. In addition, XRIs that use an XRI authority (section 2.2.1.1) are subject to a
748 particular semantic attack: spoofing the leading GCS character (section 2.2.1.2) with a
749 homograph from the Unicode character set. Such a character may cause users to believe they
750 are dealing with an XRI authority when in fact their user agent interprets the authority segment as
751 an IRI authority (section 2.2.1.3).

752 To help prevent this or any other attack based on spoofing legitimate XRI delimiters (all
753 characters in the xri-reserved production, section 2.1.2), user agents SHOULD employ one or
754 more of the following safeguards, particularly with regard to the authority segment of an XRI: a)

755 visually distinguish the defined XRI delimiter characters using special color, size, font, or other
756 mechanism that enables users to clearly understand when a legitimate XRI delimiter character is
757 being displayed, b) do not display any homograph of any XRI delimiter character in unencoded
758 form, and/or c) warn the user when an XRI contains a potentially deceptive homographic
759 character.

760 **3.4 UTF-8 Attacks**

761 Since XRIs incorporate the use of UTF-8 as specified by [IRI], they can also be subject to UTF-8
762 parsing attacks as described in section 10 of [RFC3629]:

763 “Implementers of UTF-8 need to consider the security aspects of how they
764 handle illegal UTF-8 sequences. It is conceivable that in some circumstances an
765 attacker would be able to exploit an incautious UTF-8 parser by sending it an
766 octet sequence that is not permitted by the UTF-8 syntax.”

767 For more information on these attacks, see section 10 of [RFC3629].

768 **3.5 XRI Usage in Evolving Infrastructure**

769 As XRIs are adopted as abstract identifiers, it is anticipated that new services will be developed
770 that take advantage of their extensibility. In particular, XRIs may enable new solutions to security
771 and data protection challenges at the resource identifier level that are not possible using existing
772 URI schemes.

773 For example, XRI cross-reference syntax permits the inclusion of identifier metadata such as an
774 encrypted or integrity-checked path, query or fragment. Cross-references can also be used to
775 indicate methods of obfuscating, proxying or redirecting resolution to prevent the exposure of
776 private or sensitive data.

777 A complete discussion of this topic is beyond the scope of this document. However, as a
778 consequence of XRI extensibility, it is not possible to make definitive statements regarding all
779 security and data protection considerations related to XRIs. New XRI-producing or consuming
780 applications should include independent security reviews for the specific contexts in which they
781 will be used.

782

4 References

783

4.1 Normative

- 784 **[IRI]** M. Dürst, M. Suignard, *Internationalized Resource Identifiers (IRIs)*,
785 <http://www.ietf.org/rfc/rfc3987.txt>, RFC 3987, January 2005.
- 786 **[RFC1737]** K. Sollins, L. Masinter, *Functional Requirements for Uniform Resource*
787 Names, <http://www.ietf.org/rfc/rfc1737.txt>, RFC 1737, December 1994.
- 788 **[RFC2119]** S. Bradner, *Key words for use in RFCs to Indicate Requirement Levels*,
789 <http://www.ietf.org/rfc/rfc2119.txt>, RFC 2119, March 1997.
- 790 **[RFC2141]** R. Moats, *URN Syntax*, <http://www.ietf.org/rfc/rfc2141.txt>, IETF RFC
791 2141, May 1997.
- 792 **[RFC2234]** D. H. Crocker and P. Overell, *Augmented BNF for Syntax Specifications:*
793 *ABNF*, <http://www.ietf.org/rfc/rfc2234.txt>, RFC 2234, November 1997.
- 794 **[RFC2718]** L. Masinter, H. Alvestrand, D. Zigmond, R. Petke, *Guidelines for New*
795 *URL Schemes*, <http://www.ietf.org/rfc/rfc2718.txt>, RFC 2718, November
796 1999.
- 797 **[RFC2732]** R. Hinden, B. Carpenter, L. Masinter, *Format for Literal IPv6 Addresses*
798 *in URL's*, <http://www.ietf.org/rfc/rfc2732.txt>, RFC 2732, December,
799 1999.
- 800 **[RFC3305]** M. Mealing, R. Denenberg, *Uniform Resource Identifiers (URIs), URLs,*
801 *and Uniform Resource Names (URNs): Clarifications and*
802 *Recommendations*, <http://www.ietf.org/rfc/rfc3305.txt>, RFC 3305,
803 August 2002.
- 804 **[RFC3491]** P. Hoffman, M. Blanchet, *Nameprep: A Stringprep Profile for*
805 *Internationalized Domain Names (IDN)*, <http://www.ietf.org/rfc/rfc3491>,
806 RFC 3491, March 2003.
- 807 **[RFC3629]** F. Yergeau, *UTF-8, A Transformation Format of ISO 10646*,
808 <http://www.faqs.org/rfcs/rfc3629.html>, RFC 3629, November, 2003.
- 809 **[UniXML]** M. Dürst, A. Freytag, *Unicode in XML and other Markup Languages*,
810 Unicode Technical Report #20, World Wide Web Consortium Note,
811 February 2002.
- 812 **[URI]** T. Berners-Lee, R. Fielding, L. Masinter, *Uniform Resource Identifier*
813 *(URI): Generic Syntax*, <http://www.ietf.org/rfc/rfc3986.txt>, STD 66, RFC
814 3986, January 2005.
- 815 **[UTR15]** M. Davis, M. Dürst, *Unicode Normalization Forms*,
816 <http://www.unicode.org/unicode/reports/tr15/tr15-23.html>, April 17,
817 2003.

818

4.2 Informative

- 819 **[XRIIntro]** D. Reed, D. McAlpin, *Introduction to XRIs*, [http://docs.oasis-](http://docs.oasis-open.org/committees/xri)
820 [open.org/committees/xri](http://docs.oasis-open.org/committees/xri), Work-In-Progress.
- 821 **[XRIReqs]** G. Wachob, D. Reed, M. Le Maitre, D. McAlpin, D. McPherson, *Extensible*
822 *Resource Identifier (XRI) Requirements and Glossary v1.0*,
823 [http://www.oasis-](http://www.oasis-open.org/apps/org/workgroup/xri/download.php/2523/xri-requirements-and-glossary-v1.0.doc)
824 [open.org/apps/org/workgroup/xri/download.php/2523/xri-requirements-](http://www.oasis-open.org/apps/org/workgroup/xri/download.php/2523/xri-requirements-and-glossary-v1.0.doc)
825 [and-glossary-v1.0.doc](http://www.oasis-open.org/apps/org/workgroup/xri/download.php/2523/xri-requirements-and-glossary-v1.0.doc), June 2003.

826 Appendix A. Collected ABNF for XRI (Normative)

827 This section contains the complete ABNF for XRI syntax. XRI productions use green shading,
828 while productions inherited from IRI use yellow shading. A valid XRI MUST conform to this ABNF.
829

```
830 XRI           = [ "xri://" ] xri-hier-part [ "?" iquery ]
831               [ "#" ifragment ]

832 xri-hier-part = ( xri-authority / iauthority ) xri-path-abempty

833 XRI-reference = XRI
834               / relative-XRI-ref

835 absolute-XRI = [ "xri://" ] xri-hier-part [ "?" iquery ]

836 relative-XRI-ref = relative-XRI-part [ "?" iquery ] [ "#" ifragment ]

837 relative-XRI-part = xri-path-absolute
838                   / xri-path-noscheme
839                   / ipath-empty

840 xri-value      = xri-no-scheme / relative-XRI-ref

841 xri-no-scheme = xri-hier-part [ "?" iquery ]
842               [ "#" ifragment ]

843 xri-authority = gcs-authority
844               / xref-authority

845 gcs-authority = pgcs-authority / rgcs-authority

846 pgcs-authority = "!" xri-subseg-pt-nz *xri-subseg

847 rgcs-authority = rgcs-char xri-segment

848 rgcs-char      = "=" / "@" / "+" / "$"

849 xref-authority = xref *xri-subseg

850 xref           = "(" ( XRI-reference / IRI ) ")"

851 xri-path       = xri-path-abempty
852                 / xri-path-absolute
853                 / xri-path-noscheme
854                 / ipath-empty

855 xri-path-abempty = *( "/" xri-segment )

856 xri-path-absolute = "/" [ xri-segment-nz *( "/" xri-segment ) ]

857 xri-path-noscheme = xri-subseg-od-nx *xri-subseg-nc *( "/" xri-segment )

858 xri-segment     = xri-subseg-od *xri-subseg

859 xri-segment-nz = xri-subseg-od-nz *xri-subseg
```



```

860 xri-subseg      = ( "*" / "!" ) (xref / *xri-pchar)
861 xri-subseg-nc   = ( "*" / "!" ) (xref / *xri-pchar-nc)
862 xri-subseg-od   = [ "*" / "!" ] (xref / *xri-pchar)
863 xri-subseg-od-nz = [ "*" / "!" ] (xref / 1*xri-pchar)
864 xri-subseg-od-nx = [ "*" / "!" ] 1*xri-pchar-nc
865 xri-subseg-pt-nz = "!" (xref / 1*xri-pchar)
866 xri-pchar       = iunreserved / pct-encoded / xri-sub-delims / ":"
867 xri-pchar-nc    = iunreserved / pct-encoded / xri-sub-delims
868 xri-reserved    = xri-gen-delims / xri-sub-delims
869 xri-gen-delims  = ":" / "/" / "?" / "#" / "[" / "]" / "(" / ")"
870                / "*" / "!" / rgcs-char
871 xri-sub-delims  = "&" / ";" / "," / "'"

872 IRI             = scheme ":" ihier-part [ "?" iquery ]
873                [ "#" ifragment ]

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

875 ihier-part      = "/" iauthority ipath-abempty
876                / ipath-abs
877                / ipath-rootless
878                / ipath-empty

879 iauthority      = [ iuserinfo "@" ] ihost [ ":" port ]

880 iuserinfo        = *( iunreserved / pct-encoded / sub-delims / ":" )

881 ihost           = IP-literal / IPv4address / ireg-name

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

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

884 IPv6address     =
885                /
886                / [
887                / [ *1( h16 ":" ) h16 ] "::" 3( h16 ":" ) ls32
888                / [ *2( h16 ":" ) h16 ] "::" 2( h16 ":" ) ls32
889                / [ *3( h16 ":" ) h16 ] "::"   h16 ":"   ls32
890                / [ *4( h16 ":" ) h16 ] "::"   ls32
891                / [ *5( h16 ":" ) h16 ] "::"   h16
892                / [ *6( h16 ":" ) h16 ] "::"

893 ls32            = ( h16 ":" h16 ) / IPv4address

894 h16             = 1*4HEXDIG

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

```

```

896  dec-octet      = DIGIT           ; 0-9
897                / %x31-39 DIGIT     ; 10-99
898                / "1" 2DIGIT        ; 100-199
899                / "2" %x30-34 DIGIT  ; 200-249
900                / "25" %x30-35      ; 250-255

901  ireg-name     = *( iunreserved / pct-encoded / sub-delims )

902  port          = *DIGIT

903  ipath-abempty = *( "/" isegment )

904  ipath-abs     = "/" [ isegment-nz *( "/" isegment ) ]

905  ipath-rootless = isegment-nz *( "/" isegment )

906  ipath-empty   = 0<ipchar>

907  isegment     = *ipchar

908  isegment-nz  = 1*ipchar

909  iquery       = *( ipchar / iprivate / "/" / "?" )

910  iprivate     = %xE000-F8FF / %xF0000-FFFFD / %x100000-10FFFFD

911  ifragment    = *( ipchar / "/" / "?" )

912  ipchar       = iunreserved / pct-encoded / sub-delims / ":" / "@"

913  iunreserved  = ALPHA / DIGIT / "-" / "." / "_" / "~" / ucschar

914  pct-encoded  = "%" HEXDIG HEXDIG

915  ucschar      = %xA0-D7FF / %xF900-FDCF / %xFDF0-FFEF
916                / %x10000-1FFFFD / %x20000-2FFFFD / %x30000-3FFFFD
917                / %x40000-4FFFFD / %x50000-5FFFFD / %x60000-6FFFFD
918                / %x70000-7FFFFD / %x80000-8FFFFD / %x90000-9FFFFD
919                / %xA0000-AFFFFD / %xB0000-BFFFFD / %xC0000-CFFFFD
920                / %xD0000-DFFFFD / %xE1000-EFFFFD

921  reserved     = gen-delims / sub-delims

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

923  sub-delims   = "!" / "$" / "&" / "'" / "(" / ")"
924                / "*" / "+" / "," / ";" / "="

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

```

926 **Appendix B. Transforming HTTP IRIs to XRIs**
927 **(Non-Normative)**

928 To leverage existing infrastructure, it may sometimes be useful to convert HTTP IRIs into XRIs.
929 Because XRI syntax is, for the most part, a superset of generic IRI syntax, the majority of HTTP
930 IRIs can be converted to valid XRIs simply by replacing the scheme name “http” with “xri”.
931 Generally the authority component of the resulting XRI will be properly interpreted as an IRI
932 authority. There may be some cases, however, in which a legal authority component in an IRI will
933 be interpreted as an XRI authority after this conversion. For example,

934 `http://!!1/example`

935 is a legal IRI. Converted to an XRI, it would become

936 `xri://!!1/example`

937 Because the authority segment “!!1” matches both the “xri-authority” and the “iauthority” ABNF
938 productions, it would be interpreted as an XRI authority based on the “first-match-wins” rule used
939 to resolve ambiguities in the ABNF. Section 2.2.1.2 provides other examples of legal IRI
940 authorities that would be interpreted as XRI authorities when used in an XRI. However these
941 cases are unlikely to arise in practice since they typically result in an invalid URI when converted
942 from an IRI.

943 Special consideration must also be given to HTTP IRIs employing those characters in common to
944 both the “sub-delims” production of **[IRI]** and the “xri-gen-delims” production of this specification,
945 namely opening parenthesis (“(“), closing parenthesis (“)”), star (“*”), bang (“!”), dollar sign (“\$”),
946 plus sign (“+”) and equals sign (“=”). These characters are reserved as delimiters in HTTP IRIs
947 but have no scheme-specific meaning (i.e., they are only used as delimiters in a manner defined
948 by a local authority). In XRIs, however, these characters do have defined semantics that may or
949 may not match the meaning intended by an IRI author. Conversion of such IRIs to XRIs must be
950 handled on a case-by-case basis.

951 Appendix C. Glossary

952 The following definitions are used in specifications from the OASIS XRI Technical Committee
953 Note that this glossary supercedes the glossary in [XRIReqs].

954 **Absolute Identifier**

955 An identifier that refers to a resource independent of the current context, i.e., one that
956 establishes a global context. Mutually exclusive with “Relative Identifier.”

957 **Abstract Identifier**

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

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

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

963 A URN as described in [RFC2141] is one kind of abstract identifier. Compared to
964 concrete identifiers, abstract identifiers permit additional levels of indirection in
965 referencing resources, which can be useful for a variety of purposes, including
966 persistence, equivalence, human-friendliness, and data protection.

967 **Authority (or Identifier Authority)**

968 In the context of identifiers, an authority is a resource that assigns identifiers to other
969 resources. Note that in URI syntax as defined in [URI], the “authority” production refers
970 explicitly to the top-level authority identified by the segment beginning with “//”. Since XRI
971 syntax supports unlimited federation, the term “authority” can technically refer to an
972 identifier authority at any level. However, in the “xri-authority” and “iauthority” productions
973 (section 2.2.1), it explicitly refers to the top-level identifier authority. See also “IRI
974 Authority” and “XRI Authority”

975
976 In the context of identifier resolution, an authority is a resource (typically a server) that
977 responds to resolution requests from another resource (typically a client). From this
978 perspective, each sub-segment in the authority segment of an XRI identifies a separate
979 authority.

980 **Base Identifier**

981 An absolute identifier that identifies a context for a relative identifier. Changing the base
982 identifier changes the context of the relative identifier. See “Relative Identifier.”

983 **Canonical Form**

984 The form of an identifier after applying transformation rules for the purpose of determining
985 equivalence. See also “Normal Form”.

986 **Community (or Identifier Community)**

987 A set of resources that share a common identifier authority, often (but not always) a
988 common root authority. Technically, a set of resources whose identifiers form a directed
989 graph or tree.

990 **Concrete Identifier**

991 An identifier that can be directly resolved to a resource or resource representation, rather
992 than to another identifier. Examples include the MAC address of a networked computer
993 and a phone number that rings directly to a specific device. All concrete identifiers are
994 intended to be resolvable. Contrast with “Abstract Identifier.”

995 **Context (or Identifier Context)**

996 The resource of which an identifier is an attribute. For example, in the string of identifiers
997 “a/b/c”, the context of the identifier “b” is the resource identified by “a/”, and the context of
998 the identifier “c” is the resource identified by “a/b/”. Since multiple resources may assign
999 an identifier for a target resource, the resource can be said to be identified in multiple
1000 contexts. For absolute identifiers, the context is global, i.e., there is a known starting
1001 point, or root. For relative identifiers, the context is implicit. See also “Base Identifier.”

1002 **Cross-reference**

1003 An identifier assigned in one context that is reused in another context. Cross-references
1004 enable the expression of polyarchical relationships (relationships that cross multiple
1005 hierarchies – see “Polyarchy”.) Cross-references can be used to identify logically
1006 equivalent resources in different domains, authorities, or physical locations. For example,
1007 a cross-reference may be used to identify the same logical invoice stored in two
1008 accounting systems (the originating system and the receiving system), the same logical
1009 Web document stored on multiple proxy servers, the same logical datatype used in
1010 multiple databases or XML schemas, or the same logical concept used in multiple
1011 taxonomies or ontologies.

1012 In XRI syntax, cross-references are syntactically delimited by enclosing them in
1013 parentheses. This is analogous to enclosing a word or phrase in quotation marks in a
1014 natural language, such as English, to indicate that the author is referring to it independent
1015 of the current context. For example, the phrase “love bird” is quoted in this sentence to
1016 indicate that we are *mentioning*, rather than *using*, the phrase - that is, we are referring to
1017 it independent of the context of this glossary.

1018 **Delegated Identifier**

1019 A multi-segment identifier in which segments are assigned by more than one identifier
1020 authority. Namespace authority is delegated from one identifier authority to the next.
1021 Mutually exclusive with “Local Identifier.”

1022 **Federated Identifier**

1023 A delegated identifier that spans multiple independent identifier authorities. See also
1024 “Delegated Identifier.”

1025 **Global Context Symbol (GCS)**

1026 A reserved character used at the start of the authority segment of an XRI to establish the
1027 global context of an XRI authority. See section 2.2.1.2.

1028 **Hierarchy**

1029 A branching tree structure in which all primary relationships are parent-child. (Sibling
1030 relationships in a hierarchy are secondary, derived from the parent-child relationships.)
1031 URI and IRI syntax has explicit support for hierarchical paths. XRI syntax supports both
1032 hierarchical and polyarchical paths. See “Polyarchy” and “Cross-reference.”

1033 **Human-Friendly Identifier (HFI)**

1034 An identifier containing words or phrases intended to convey meaning in a specific
1035 human language and therefore be easy for people to remember and use. Contrast with
1036 “Machine-Friendly Identifier.”

1037 **Identifier**

1038 Per [URI], anything that “embodies the information required to distinguish what is being
1039 identified from all other things within its scope of identification.” In UML terms, an
1040 identifier is an attribute of a resource (the identifier context) that forms an association with
1041 another resource (the identifier target). The general term “identifier” does not specify
1042 whether the identifier is abstract or concrete, absolute or relative, persistent or

- 1043 reassignable, human-friendly or machine-friendly, delegated or local, hierarchical or
1044 polyarchical, or resolvable or self-referential.
- 1045 **I-name**
- 1046 An informal term used to refer to a reassignable XRI; more specifically, an XRI in which
1047 at least one sub-segment is reassignable.
- 1048 **I-number**
- 1049 An informal term used to refer to a persistent XRI; more specifically, an XRI in which all
1050 sub-segments are persistent. Note that a persistent XRI is not required to be numeric—it
1051 may be any text string meeting the XRI ABNF requirements.
- 1052 **IRI (Internationalized Resource Identifier)**
- 1053 IRI is a specification for internationalized URIs developed by the W3C. IRIs specify how
1054 to include characters from the Universal Character Set (Unicode/ISO10646) in URIs. The
1055 IRI specification **[IRI]** provides a mapping from IRIs to URIs, which allows IRIs to be used
1056 instead of URIs where appropriate. This XRI specification defines a similar transformation
1057 from XRIs to IRIs for the same reason.
- 1058 **IRI Authority**
- 1059 An identifier authority (see “Authority”) represented by the authority segment of an XRI
1060 that does not match the “xri-authority” production but matches the “iauthority” production
1061 from **[IRI]**. See section 2.2.1.3. Mutually exclusive with “XRI Authority”.
- 1062 **Local Identifier**
- 1063 Any identifier, or any set of segments in a multi-segment identifier, that is assigned by the
1064 same identifier authority. Each of these segments is local to that authority. Mutually
1065 exclusive with “Delegated Identifier.”
- 1066 **Machine-Friendly Identifier (MFI)**
- 1067 An identifier containing digits, hexadecimal values, or other character sequences
1068 optimized for efficient machine indexing, searching, routing, caching, and resolvability.
1069 MFIs generally do not contain human semantics. Compare with “Human-Friendly
1070 Identifier.”
- 1071 **Normal Form**
- 1072 The character-by-character format of an identifier after encoding, escaping, or other
1073 character transformation rules have been applied in order to satisfy syntactic
1074 requirements. Three normal forms are defined for XRIs—XRI-normal form, IRI-normal
1075 form, and URI-normal form. See section 2.3.1 for details. See also “Canonical Form”.
- 1076 **Path**
- 1077 The relationships between resources defined by a multi-segment identifier. In less strict
1078 contexts, the word “path” often refers to the multi-segment identifier itself, or to the
1079 resources it represents (such as filesystem directories).
- 1080 **Persistent Identifier**
- 1081 An identifier that is permanently assigned to a resource and intended never to be
1082 reassigned to another resource - even if the original resource goes off the network, is
1083 terminated, or ceases to exist. A URN as described in **[RFC2141]** is an example of a
1084 persistent identifier. Persistent identifiers tend to be machine-friendly identifiers, since
1085 human-friendly identifiers often reflect human semantic relationships that may change
1086 over time. Mutually exclusive with “Reassignable Identifier.”
- 1087 **Polyarchy**
- 1088 A treelike structure composed of multiple intersecting hierarchies in which primary
1089 relationships can cross hierarchies. A polyarchy allows one member to be connected or

- 1090 linked to any other. In contrast to a web, however, the overall structure tends to remain
1091 strongly hierarchical. XRI support polyarchic paths through the use of cross-references.
1092 See also “Cross-reference” and “Hierarchy”.
- 1093 **Reassignable Identifier**
- 1094 An identifier that may be reassigned from one resource to another. Example: the domain
1095 name “example.com” may be reassigned from ABC Company to XYZ Company, or the
1096 email address “mary@example.com” may be reassigned from Mary Smith to Mary Jones.
1097 Reassignable identifiers tend to be human-friendly because they often represent the
1098 potentially transitory mapping of human semantic relationships onto network resources or
1099 resource representations. Mutually exclusive with “Persistent Identifier.”
- 1100 **Relative Identifier**
- 1101 An identifier that refers to a resource only in relationship to a particular context (for
1102 example, the current community, the current document, or the current position in a
1103 delegated identifier). If the context changes, the identifier’s meaning also changes. A
1104 relative identifier can be converted into an absolute identifier by combining it with a base
1105 identifier (an absolute identifier that is used to identify a context). See “Base Identifier”.
1106 Mutually exclusive with “Absolute Identifier.”
- 1107 **Resolvable Identifier**
- 1108 An identifier that references a network resource or resource representation and that can
1109 be dereferenced using a resolution protocol or other mechanism into a network endpoint
1110 for communicating with the target resource. Mutually exclusive with “Self-Reference.”
- 1111 **Resource**
- 1112 Per [URI], “anything that can be named or described.” Resources are of two types:
1113 network resources (those that are network-addressable) and non-network resources
1114 (those that exist entirely independent of a network). Network resources are themselves of
1115 two types: physical resources (resources physically attached to or operating on the
1116 network) or resource representations (see “Resource Representation”).
- 1117 **Resource Representation**
- 1118 A network resource that represents the attributes of another resource. A resource
1119 representation may represent either another network resource (such as a machine,
1120 service, application, file, or digital object) or a non-network resource (such as a person,
1121 organization, or concept).
- 1122 **Segment (or Identifier Segment)**
- 1123 Any syntactically delimited component of an identifier. In generic URI syntax, all
1124 segments after the authority portion are delimited by forward slashes
1125 (“/segment1/segment2/...”). In XRI syntax, slash segments can be further subdivided into
1126 sub-segments called *star segments* (for reassignable identifiers) and *bang segments* (for
1127 persistent identifiers). See section 2.2.3. XRI also supports another type of segment
1128 called a cross-reference, which is enclosed in parentheses. See “Cross-Reference”.
- 1129 **Self-Reference (or Self-Referential Identifier)**
- 1130 An identifier which is itself the representation of the resource it references. Self-
1131 references are typically used to represent non-network resources (e.g., “love”, “Paris”,
1132 “the planet Jupiter”) in contexts where an identifier is not intended to be resolved to a
1133 separate network representation of that resource. The primary purpose of self-references
1134 is to establish equivalence across contexts (see “Cross-References”). Mutually exclusive
1135 with “Resolvable Identifier.”
- 1136 **Sub-segment**
- 1137 A syntactically delimited component of an identifier segment (see “Segment”). While URI
1138 and IRI syntax define only segments, XRI syntax defines both segments and sub-

1139 segments. XRI sub-segments are used to distinguish between persistent identifiers,
1140 reassignable identifiers, and cross-references. See sections 2.2.2 and 2.2.3.

1141 **Synonym (or Identifier Synonym)**

1142 An identifier that is asserted by an identifier authority to be equivalent to another identifier
1143 not because of strict literal equivalence, but because it resolves to the same resource.

1144 **Target (or Identifier Target)**

1145 The resource referenced by an identifier. A target may be either a network resource
1146 (including a resource representation) or a non-network resource.

1147 **URI (Uniform Resource Identifier)**

1148 The standard identifier used in World Wide Web architecture. Starting in 1998, RFC 2396
1149 has been the authoritative specification for URI syntax. In January 2005 it was
1150 superseded by RFC 3986 [URI].

1151 **XDI (XRI Data Interchange)**

1152 A generalized, extensible service for sharing, linking, and synchronizing XML data and
1153 metadata associated with XRI-identified resources. XDI is being developed by the OASIS
1154 XDI Technical Committee (<http://www.oasis-open.org/committees/xdi>).

1155 **XRI Authority**

1156 An identifier authority (see “Authority”) represented by the authority segment of an XRI
1157 that begins with either a global context symbol or a cross-reference. See section 2.2.1.1.
1158 Mutually exclusive with “IRI Authority.”

1159 **XRI Reference**

1160 A term that includes both absolute and relative XRIs. Used in the same way as “URI
1161 reference” and “IRI reference.” Note that to transform an XRI reference into an XRI, it
1162 must first be converted to absolute form, which in the case of a relative XRI requires the
1163 use of a base XRI to establish context.

1164

Appendix D. Acknowledgments

1165 The editors would like to acknowledge the contributions of the OASIS XRI Technical Committee,
1166 whose voting members at the time of publication were:

- 1167 • Geoffrey Strongin, Advanced Micro Devices
- 1168 • Ajay Madhok, AmSoft Systems
- 1169 • William Barnhill, Booz Allen and Hamilton
- 1170 • Drummond Reed, Cordance Corporation
- 1171 • Marc Le Maitre, Cordance Corporation
- 1172 • Dave McAlpin, Epok
- 1173 • Loren West, Epok
- 1174 • Chetan Sabnis, Epok
- 1175 • Paul Trevithick
- 1176 • Les Chasen, NeuStar
- 1177 • Peter Davis, NeuStar
- 1178 • Trung Tran, NeuStar
- 1179 • Ning Zhang, NeuStar
- 1180 • Masaki Nishitani, Nomura Research
- 1181 • Nat Sakimura, Nomura Research
- 1182 • Tetsu Watanabe, Nomura Research
- 1183 • Owen Davis, PlaNetwork
- 1184 • Victor Grey, PlaNetwork
- 1185 • Fen Labalme, PlaNetwork
- 1186 • Mike Lindelsee, Visa International
- 1187 • Gabriel Wachob, Visa International
- 1188 • Dave Wentker, Visa International
- 1189 • Bill Washburn, XDI.ORG

1190 The editors also would like to acknowledge the following people for their contributions to previous
1191 versions of the OASIS XRI specifications (affiliations listed for OASIS members):

1192 Thomas Bikeev, EAN International; Krishna Sankar, Cisco; Winston Bumpus, Dell; Joseph
1193 Moeller, EDS; Steve Green, Epok; Lance Hood, Epok; Adarbad Master, Epok; Davis McPherson,
1194 Epok; Phillipe LeBlanc, GemPlus; Jim Schreckengast, Gemplus; Xavier Serret, Gemplus; John
1195 McGarvey, IBM; Reva Modi, Infosys; Krishnan Rajagopalan, Novell; Tomonori Seki, NRI; James
1196 Bryce Clark, OASIS; Marc Stephenson, TSO; Mike Mealling, Verisign; Rajeev Maria, Visa
1197 International; Terence Spielman, Visa International; John Veizades, Visa International; Lark Allen,
1198 Wave Systems; Michael Willett, Wave Systems; Matthew Dovey; Eamonn Neylon; Mary
1199 Nishikawa; Lars Marius Garshol; Norman Paskin; Bernard Vatant.

1200 A special acknowledgement to Jerry Kindall (Epok) for a full editorial review.

1201 Also, the authors of and contributors to the following documents and specifications are
1202 acknowledged for the intellectual foundations of the XRI specification:

- 1203 • RFC 1737
- 1204 • RFC 2616
- 1205 • RFC 2718
- 1206 • RFC 3986 (STD 66) and its predecessor, RFC 2396
- 1207 • RFC 3987
- 1208 • XNS
- 1209

1210

Appendix E. Notices

1211 Copyright © OASIS® 1993–2005. All Rights Reserved.

1212 All capitalized terms in the following text have the meanings assigned to them in the OASIS
1213 Intellectual Property Rights Policy (the "OASIS IPR Policy"). The full Policy may be found at the
1214 OASIS website.

1215 This document and translations of it may be copied and furnished to others, and derivative works
1216 that comment on or otherwise explain it or assist in its implementation may be prepared, copied,
1217 published, and distributed, in whole or in part, without restriction of any kind, provided that the
1218 above copyright notice and this section are included on all such copies and derivative works.
1219 However, this document itself may not be modified in any way, including by removing the
1220 copyright notice or references to OASIS, except as needed for the purpose of developing any
1221 document or deliverable produced by an OASIS Technical Committee (in which case the rules
1222 applicable to copyrights, as set forth in the OASIS IPR Policy, must be followed) or as required to
1223 translate it into languages other than English.

1224 The limited permissions granted above are perpetual and will not be revoked by OASIS or its
1225 successors or assigns.

1226 This document and the information contained herein is provided on an "AS IS" basis and OASIS
1227 DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO
1228 ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE
1229 ANY OWNERSHIP RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR
1230 FITNESS FOR A PARTICULAR PURPOSE.

1231 OASIS requests that any OASIS Party or any other party that believes it has patent claims that
1232 would necessarily be infringed by implementations of this OASIS Committee Specification or
1233 OASIS Standard, to notify OASIS TC Administrator and provide an indication of its willingness to
1234 grant patent licenses to such patent claims in a manner consistent with the IPR Mode of the
1235 OASIS Technical Committee that produced this specification.

1236 OASIS invites any party to contact the OASIS TC Administrator if it is aware of a claim of
1237 ownership of any patent claims that would necessarily be infringed by implementations of this
1238 specification by a patent holder that is not willing to provide a license to such patent claims in a
1239 manner consistent with the IPR Mode of the OASIS Technical Committee that produced this
1240 specification. OASIS may include such claims on its website, but disclaims any obligation to do
1241 so.

1242 OASIS takes no position regarding the validity or scope of any intellectual property or other rights
1243 that might be claimed to pertain to the implementation or use of the technology described in this
1244 document or the extent to which any license under such rights might or might not be available;
1245 neither does it represent that it has made any effort to identify any such rights. Information on
1246 OASIS' procedures with respect to rights in any document or deliverable produced by an OASIS
1247 Technical Committee can be found on the OASIS website. Copies of claims of rights made
1248 available for publication and any assurances of licenses to be made available, or the result of an
1249 attempt made to obtain a general license or permission for the use of such proprietary rights by
1250 implementers or users of this OASIS Committee Specification or OASIS Standard, can be
1251 obtained from the OASIS TC Administrator. OASIS makes no representation that any information
1252 or list of intellectual property rights will at any time be complete, or that any claims in such list are,
1253 in fact, Essential Claims.

1254 The names "OASIS", "Extensible Resource Identifier", and "XRI" are trademarks of OASIS, the
1255 owner and developer of this specification, and should be used only to refer to the organization
1256 and its official outputs. OASIS welcomes reference to, and implementation and use of,
1257 specifications, while reserving the right to enforce its marks against misleading uses. Please see
1258 <http://www.oasis-open.org/who/trademark.php> for above guidance.