



# Extensible Resource Identifier (XRI) Syntax and Resolution Specification

## Working Draft 07, 29 July 2003

### Document identifier:

wd-xri-specification-07

### Location:

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

### Editors:

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

### Abstract:

This document is the normative technical specification for XRI syntax and resolution. For an introduction to the uses and features of XRIs, see the non-normative *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](mailto:xri@lists.oasis-open.org) list. Others should subscribe to and send comments to the [xri-comment@lists.oasis-open.org](mailto:xri-comment@lists.oasis-open.org) list. To subscribe, send an email message to [xri-comment-request@lists.oasis-open.org](mailto: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>.

---

## 32 Table of Contents

33	Introduction .....	4
34	1.1 Overview of XRIs.....	4
35	1.1.1 Generic Syntax .....	4
36	1.1.2 Examples .....	5
37	1.1.3 URI, URL, URN, and XRI.....	5
38	1.2 Design Considerations .....	6
39	1.2.1 Abstraction and Independence .....	6
40	1.2.2 Persistence and Reassignability.....	6
41	1.2.3 Human-friendliness and Machine-friendliness.....	6
42	1.2.4 Internationalization.....	6
43	1.2.5 Cross-Context Identification.....	7
44	1.2.6 Authority, Delegation, and Federation .....	7
45	1.2.7 Security and Privacy .....	7
46	1.2.8 Extensibility .....	7
47	1.3 Terminology and Notation .....	7
48	1.3.1 Keywords .....	7
49	1.3.2 Syntax Notation.....	7
50	1.3.3 Glossary.....	8
51	2 Syntax .....	11
52	2.1 Syntax Components.....	11
53	2.1.1 Authority.....	11
54	2.1.1.1 URI Authority .....	11
55	2.1.1.2 XRI Authority.....	12
56	2.1.1.2.1 Global Context Symbols (GCS) .....	13
57	2.1.1.2.2 Cross-References .....	13
58	2.1.2 Path.....	14
59	2.1.3 Query .....	14
60	2.1.4 Fragment.....	15
61	2.2 Characters.....	15
62	2.2.1 Reserved Characters.....	15
63	2.2.2 Unreserved Characters.....	15
64	2.2.3 Escaped Characters .....	16
65	2.2.3.1 Escaped Encoding.....	16
66	2.2.3.2 Converting XRIs to URIs.....	16
67	2.2.3.3 XRI-specific conversion for use in URIs.....	17
68	2.2.3.4 Converting URIs to XRIs.....	19
69	2.2.4 Excluded Characters.....	20
70	2.2.5 Legal Character Sequence .....	20
71	2.3 Character Encoding and Internationalization .....	20
72	2.4 Relative XRIs.....	20
73	2.4.1 Establishing a Base XRI .....	20
74	2.4.2 Obtaining the Referenced XRI.....	21

75	2.5 Normalization and Comparison.....	21
76	3 Resolution.....	23
77	3.1 Introduction to Resolution Architecture.....	23
78	3.1.1 Assumptions.....	23
79	3.1.2 Phases of Resolution.....	23
80	3.2 Phase 1: Authority Resolution.....	24
81	3.2.1 General Description.....	24
82	3.2.2 DNS-specified Authority Resolution (DAR).....	25
83	3.2.3 XRI Authority Resolution Framework (XARF).....	26
84	3.2.3.1 Introduction.....	26
85	3.2.3.2 User Relative XRIs.....	26
86	3.2.3.3 Authority Descriptors.....	27
87	3.2.3.4 Authority Protocol Descriptors.....	27
88	3.2.3.5 Algorithm.....	28
89	3.2.3.6 XRI-HTTP Relative Lookup Mechanism (NA1).....	30
90	3.2.3.7 RDDS Relative Lookup Mechanism (NA2).....	31
91	3.2.3.7.1 Introduction.....	32
92	3.2.3.7.2 Algorithm.....	33
93	3.2.4 IP-Address Authority Resolution (IAR).....	34
94	3.3 Phase 2: Local Access.....	34
95	3.3.1 Format of Local Access Descriptors.....	34
96	3.3.2 Format of Local Access Protocol Descriptors.....	35
97	3.3.3 Local Access Service Descriptors.....	35
98	3.3.4 Requirements for Local Access Bindings.....	35
99	3.3.5 Local Access Bindings.....	36
100	3.3.5.1 THHTTP Local Access Binding.....	36
101	3.4 Flowchart of Authority Resolution.....	37
102	4 Security and Data Protection.....	38
103	4.1 XRI Usage in Legacy Infrastructure.....	38
104	4.2 Secure Resolution.....	38
105	4.3 XRI Usage in Evolving Infrastructure.....	38
106	5 References.....	39
107	5.1 Normative.....	39
108	5.2 Informative.....	39
109	Appendix A. Collected ABNF for XRI.....	40
110	Appendix B. Special Identifiers Assigned by the XRI Specification.....	43
111	Appendix C. Transforming HTTP URIs to XRIs.....	44
112	Appendix D. Acknowledgments.....	45
113	Appendix E. Revision History.....	46
114	Appendix F. Notices.....	47
115	Appendix G. Issues.....	48
116		

---

# 117 Introduction

## 118 1.1 Overview of XRIs

119 An Extensible Resource Identifier (XRI) provides a standard means of abstractly identifying a  
120 resource independent of any given concrete representation of that resource (or, in the case of a  
121 completely abstract resource, independent of any representation at all). XRIs are defined similarly  
122 to URIs in "*Uniform Resource Identifiers (URI): Generic Syntax*" [RFC2396] but contain additional  
123 syntactical elements and extend the unreserved character set to include characters beyond those  
124 allowed in generic URIs. To accommodate applications that expect generic URIs, rules are  
125 defined that allow an XRI to be transformed into a conformant URI as defined by [RFC2396].  
126 Since a revision of RFC 2396 is currently a work in progress, the XRI scheme also incorporates  
127 some simplifications and enhancements to generic URI syntax as proposed in [RFC2396bis].

128 In addition, XRI syntax is internationalized following the recommendations in "*Guidelines for New  
129 URL Schemes*" [RFC2718] and "*Extensible Markup Language (XML) 1.0 (Second Edition)*"  
130 [XML], and specifically the requirements of the "anyURI" datatype as specified in "XML Schema  
131 Part 2: Datatypes" [XMLSchema2]. To do this, the XRI scheme incorporates the syntax  
132 recommended in another work-in-progress, "*Internationalized Resource Identifiers (IRIs)*" [IRI].

133 Although an XRI is not a Uniform Resource Name (URN) as defined in "*URN Syntax*" [RFC2141],  
134 fully persistent XRIs are also designed to meet the requirements set out in "*Functional  
135 Requirements for Uniform Resource Names*" [RFC1737].

136 This document specifies the ABNF that defines the XRI scheme. A valid XRI MUST conform to  
137 the ABNF specified in this document. In addition this document specifies a resolution framework  
138 for XRIs. An XRI MAY be resolved using one or more of mechanisms specified by this framework.

### 139 1.1.1 Generic Syntax

140 URI syntax is designed to be simple and extensible, and XRI syntax is very similar. A fully-  
141 qualified XRI consists of the scheme name "xri:" followed by the same four optional components  
142 as a generic URI.  
143

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

145  
146 One advantage of this approach is that the vast majority of HTTP URIs, which inherit directly from  
147 generic URI syntax, can be transformed to valid XRIs simply by changing the scheme from "http"  
148 to "xri". The relationship of HTTP URIs and XRIs and rules for this transformation are further  
149 discussed in [ref to Appendix C, "Transforming HTTP URIs to XRIs"].

150 XRI syntax extends this generic URI syntax in six ways by providing syntactic support for:

- 151 1. *Persistent and reassignable segments*. Generic URI syntax does not distinguish between  
152 persistent and reassignable identifiers. XRI syntax enables the top-level authority  
153 segment as well as any subsequent path segment to be expressed as either persistent or  
154 reassignable.
- 155 2. *Unlimited delegation*. Generic URI syntax supports delegated identifiers (i.e., DNS names  
156 or IP addresses) within the top-level authority segment. XRI syntax supports delegation  
157 of both persistent and reassignable identifiers at any level of the path.
- 158 3. *Cross-references*. Generic URI syntax does not provide for nesting of URIs in order to  
159 share identifiers across contexts. Since this is particularly useful with abstract identifiers  
160 (e.g., to establish the generic type of a resource, or to share identifier metadata such as  
161 versioning), XRI syntax allows URIs (including XRIs) to be nested inside parentheses.

- 162 4. *Internationalized character set*. Generic URI syntax limits legal characters to a subset of  
163 the repertoire of US-ASCII characters. XRI syntax allows the much wider repertoire of  
164 Unicode characters, greatly facilitating the use of XRIs in languages other than English.
- 165 5. *Global context symbols*. In addition to generic URI syntax for DNS and IP authorities, XRI  
166 syntax provides shorthand symbols for establishing the global context of an identifier.
- 167 6. *Non-resolvability*. Generic URI syntax does not provide a way to indicate whether or not a  
168 URI is resolvable. Since an XRI may itself be the full representation of a non-resolvable,  
169 abstract resource (e.g., a concept like "love", "honor", or "user-friendly") that is used only  
170 for the purposes of establishing equivalence, XRI syntax permits an XRI value to be  
171 expressed as explicitly non-resolvable.

## 172 1.1.2 Examples

173 The following examples illustrate XRI syntax. They have minimal annotation and are only  
174 intended to give a sense of the scope of XRI syntax. For details and the normative syntax, see  
175 Section 2.  
176

```
177 xri://www.example.com/pages/index.html
178     --standard HTTP URI used as an XRI
179
180 xri://[2010:836B:4179::836B:4179]/pages/index.html
181     --using an IPv6 authority per RFC 2732
182
183 xri://www.example.com/inventory.parts/widget.subwidget.foobarator
184     --delegation of reassignable identifiers
185
186 xri://www.example.com/:inventory:parts/:12:7:234
187     --delegation of persistent identifiers
188
189 xri:@ExampleCorp
190 xri:@ExampleCorp.website
191 xri:=JohnDoe
192 xri:=JohnDoe.home
193 xri:=JohnDoe.work
194 xri:+flowers
195 xri:+flowers.rose
196 xri:+flowers.daisy
197     --global context symbols
198
199 xri://www.example.com/(+management)/(+CEO)
200 xri:(urn:oasis:spec:2040)/(+tableofcontents)
201 xri:(mailto:john.doe@example.com)/(+email.address)
202 xri:=JohnDoe.home/(+email.address)
203 xri:=JohnDoe.home/(+email.address).($v/3)
204     --cross-references
205
206 xri:(+flowers.rose)
207 xri:(//www.example.com/dictionary/flowers/rose)
208     --non-resolvable XRIs
209
```

## 210 1.1.3 URI, URL, URN, and XRI

211 The evolution and interrelationships of the terms "URI", "URL", and "URN" are explained in a  
212 report from the Joint W3C/IETF URI Planning Interest Group, "*Uniform Resource Identifiers*  
213 *(URIs), URLs, and Uniform Resource Names (URNs): Clarifications and Recommendations*"  
214 **[RFC3305]**. This report states in Section 2.1:

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

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

220 "Over time, the importance of this additional level of hierarchy seemed to lessen; the view  
221 became that an individual scheme did not need to be cast into one of a discrete set of URI  
222 types, such as "URL", "URN", "URC", etc. Web-identifier schemes are, in general, URI  
223 schemes, as a given URI scheme may define subspaces."

224 This conclusion is shared by [RFC2396bis], which states in Section 1.1.3:

225 "An individual [URI] scheme does not need to be classified as being just one of "name" or  
226 "locator". Instances of URIs from any given scheme may have the characteristics of names or  
227 locators or both, often depending on the persistence and care in the assignment of identifiers  
228 by the naming authority, rather than any quality of the scheme."

229 The XRI scheme expressly embraces this precept. As an abstract URI, an XRI is explicitly  
230 intended to be used as a persistent identifier or long-term "name" for a resource. However XRIs  
231 are also resolvable and can be used a method of locating a resource (including another XRI).  
232 Since in certain contexts it may be important to distinguish whether an XRI is intended to be  
233 resolved vs. being used only for identification, the XRI scheme includes syntax for expressing this  
234 difference. See [ref to "Cross-references"].

## 235 1.2 Design Considerations

236 The full set of requirements for XRI syntax and resolution is documented in "*XRI Requirements  
237 and Glossary v1.0*" [XRIReqs]. A synopsis of the major design considerations is included here.

### 238 1.2.1 Abstraction and Independence

239 The preeminent requirement is that XRI syntax be fully abstract, i.e., independent of resource  
240 location, network, application, transport protocol, type, or security method. Although XRI syntax  
241 may be extended for specific uses, the generic XRI syntax is designed to represent pure UML-  
242 describable associations between resources (see [UML]) and thus to allow portability across all  
243 networks, directories, domains, and applications.

### 244 1.2.2 Persistence and Reassignability

245 As noted in Section 1.1.3 above, XRI syntax and resolution is designed to express and resolve  
246 fully persistent identifiers, fully reassignable identifiers, or any combination of persistent and  
247 reassignable identifier segments.

### 248 1.2.3 Human-friendliness and Machine-friendliness

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

### 253 1.2.4 Internationalization

254 XRIs are designed to be rendered in the natural language of their intended consumer. They allow  
255 the Unicode range of characters [Unicode] and provide syntactical support for expressing  
256 optional language-dependent context metadata. As a result, XRIs extend the virtues of human  
257 readability, memorability, and usability to non-English speaking audiences.

## 258 **1.2.5 Cross-Context Identification**

259 XRI syntax and resolution is designed to allow the use of an absolute identifier in the context of  
260 another absolute identifier, i.e., for a URI (including an XRI) to be contained within another XRI.  
261 Such embedded identifiers are called *cross-references*, and they are key to XRI extensibility.

## 262 **1.2.6 Authority, Delegation, and Federation**

263 XRI syntax and resolution are designed to allow any resource to serve as a root authority, and for  
264 any authority to delegate to any other authority at any level of the path. Thus XRI design imposes  
265 no specific delegation model, network topology, or federation structure.

## 266 **1.2.7 Security and Privacy**

267 XRI syntax and resolution is designed to be adapted to any security model, method, or  
268 infrastructure, as well as to any privacy policy or framework. XRI design does not require sensitive  
269 data to be included in an identifier, and if such data is needed in an XRI, the syntax permits  
270 encryption and obfuscation of identifier segments for enhanced security and privacy.

## 271 **1.2.8 Extensibility**

272 Like XML, the XRI scheme is designed to be extended and specialized by different identifier  
273 authorities, and also like XML, these extensions and specializations are designed to be  
274 interoperable.

## 275 **1.3 Terminology and Notation**

### 276 **1.3.1 Keywords**

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

### 281 **1.3.2 Syntax Notation**

282 This specification uses the same syntax notation as [RFC2396], namely, Augmented Backus-  
283 Naur Form (ABNF) as defined in [RFC2234]. As explained in RFC 2396, although the ABNF  
284 defines syntax in terms of the US-ASCII character encoding, XRI syntax should be interpreted in  
285 terms of the character that the ASCII-encoded octet represents, rather than the octet encoding  
286 itself. Like other URIs, how an XRI is represented in terms of bits and bytes on the wire is  
287 dependent upon the character encoding of the protocol used to transport it, or the charset of the  
288 document that contains it.

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

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

295

296 `This is an example of ABNF specific to XRI.`

297

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

299

300 In addition, productions inherited from the IRI proposal [IRI] are prefixed with the letter "i" as they  
301 are in that document.

### 302 1.3.3 Glossary

303 A complete glossary of XRI-related terms is included in XRI Requirements and Glossary v1.0  
304 [XRIReqs]. Following are the definitions central to this specification.

305 [Note: Are terms that need to be defined specific to internationalization?]

#### 306 **Absolute Identifier**

307 An identifier that refers to a resource independent of the current context, i.e., using a  
308 global context. Mutually exclusive with "Relative Identifier".

#### 309 **Abstract Identifier**

310 An identifier that is not directly resolvable to a resource, but is either: a) non-resolvable  
311 because it abstractly represents a non-network resource (see "Non-Resolvable  
312 Identifier"), or b) must be resolved to another identifier first (which may in turn be another  
313 abstract identifier, or a concrete identifier). A URN as described in [RFC2141] is an  
314 example of an abstract identifier. Abstract identifiers provide for additional levels of  
315 indirection in referencing resources which can be useful for a variety of purposes,  
316 including persistence, equivalence, human-friendliness, and data protection.

#### 317 **Authority (or Identifier Authority)**

318 A resource that assigns identifiers to other resources. Note that in URI ABNF (and in the  
319 equivalence sections of XRI ABNF), the "authority" production refers explicitly to the top-  
320 level authority, i.e., the community root. However elsewhere in this specification the term  
321 "authority" refers more generally to the entity responsible for assigning and resolving  
322 identifiers at any level of delegation.

#### 323 **Community (or Identifier Community)**

324 The set of resources that share a common identifier authority, typically a common root  
325 authority. Technically, the set of resources whose identifiers form a directed acyclic graph  
326 or tree.

#### 327 **Concrete Identifier**

328 An identifier that can be directly resolved to a resource, rather than indirectly to another  
329 identifier. Examples include the MAC address of a networked computer, a phone number  
330 (that rings directly to a specific device), and a postal address (that is not a forwarding  
331 address). All concrete identifiers are intended to be resolvable identifiers. Contrast with  
332 "Abstract Identifier".

#### 333 **Context (or Identifier Context)**

334 The backpointer of an identifier, i.e., the resource of which the identifier is an attribute.  
335 Context is the parent resource that assigns the identifier for the target resource. Since  
336 multiple resources may assign an identifier for a target resource, the resource can be  
337 said to be identified in multiple contexts. For absolute identifiers, the context is global,  
338 i.e., they have a known starting point. For relative identifiers, the context is local, i.e., it  
339 depends on the resource resolving the identifier.

#### 340 **Cross-reference**

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



- 348 **Delegated Identifier**
- 349 A multi-segment identifier in which different segments are assigned by different identifier  
350 authorities. Mutually exclusive with "Local Identifier".
- 351 **Identifier**
- 352 Per **[RFC2396bis]**, anything that "embodies the information required to distinguish what  
353 is being identified from all other things within its scope of identification". In UML terms, an  
354 identifier is an attribute of a resource (the identifier context) that forms an association with  
355 another resource (the identifier target). The general term "identifier" does not specify  
356 whether the identifier is abstract or concrete, persistent or reassignable, human-friendly  
357 or machine-friendly, absolute or relative, local or delegated, or resolvable or non-  
358 resolvable.
- 359 **Local Identifier**
- 360 A single identifier, or any set of segments in a multi-segment identifier, that are assigned  
361 by the same identifier authority. Mutually exclusive with "Delegated Identifier".
- 362 **Non-Resolvable Identifier**
- 363 An identifier that does not directly reference a network resource or resource  
364 representation, but only abstractly represents a resource. A non-resolvable identifier is  
365 always an abstract identifier and does not have any corresponding data or metadata  
366 describing the resource it represents, thus it cannot be resolved in the conventional  
367 sense. From a machine perspective, the purpose of non-resolvable identifiers is to  
368 establish equivalence across contexts. Mutually exclusive with "Resolvable Identifier."
- 369 **Persistent Identifier**
- 370 An identifier that is permanently assigned to a resource and that is intended never to be  
371 reassigned to another resource even if the original resource goes off the network, is  
372 terminated, or no longer exists. A URN as described in **[RFC2141]** is a persistent  
373 identifier. Mutually exclusive with "Reassignable Identifier".
- 374 **Reassignable Identifier**
- 375 An identifier that may be reassigned from one resource to another. Example: the domain  
376 name "example.com" may be reassigned from ABC Company to XYZ Company, or the  
377 email address "john@example.com" may be reassigned from John Smith to John Jones.  
378 Reassignable identifiers tend to be human-friendly identifiers because they often  
379 represent the mapping of semantic relationships onto network resources or resource  
380 representations. Mutually exclusive with "Persistent Identifier".
- 381 **Relative Identifier**
- 382 An identifier that refers to a resource only in relationship to the current context, i.e., the  
383 context in which the identifier is being resolved. Mutually exclusive with "Absolute  
384 Identifier".
- 385 **Resolvable Identifier**
- 386 An identifier that references a network resource or resource representation and that can  
387 be resolved into data or metadata describing the target resource. Mutually exclusive with  
388 "Non-Resolvable Identifier."
- 389 **Resource**
- 390 Per **[RFC2396bis]**, "anything that can be named or described". Resources are of two  
391 types: network resources (those that are network addressable) and non-network  
392 resources (those that exist entirely independent of a network). Network resources in turn  
393 contain a subtype, resource representations. A resource representation may represent  
394 either a network resource or a non-network resource.
- 395 **Resource Representation**

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

399 **Target (or Identifier Target)**

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

---

## 402 2 Syntax

### 403 2.1 Syntax Components

404 Generic XRI syntax consists of the scheme name "xri:" follow by the same hierarchical sequence  
405 of components as generic URI syntax. [ Need to highlight that we are not defining a URI scheme,  
406 but still using some of the URI ABNF productions ] Taken as a whole this sequence is referred to  
407 as the *XRI value*.

408

```
409 XRI = "xri:" xri-value  
410 xri-value = [ xri-path ] [ "?" xri-query ] [ "#" xri-fragment ]
```

411

412 The path component can be hierarchical to any depth. A path can be globally absolute, relative to  
413 the local community, or relative to the current context as discussed in [ref Relative XRI section].

414

```
415 xri-path = global-path / local-path / relative-path  
416 global-path = authority-part [ local-path ]  
417 local-path = "/" relative-path  
418 relative-path = *( [ "." ] "./" ) xri-segments
```

419

#### 420 2.1.1 Authority

421 XRI syntax supports the same set of authorities as generic URI syntax, called a *URI authority*. In  
422 addition it supports an *XRI authority* which provides two other mechanisms of specifying the  
423 global context of an identifier, as defined in section 2.1.1.2.

424

```
425 authority-part = URI-authority / XRI-authority
```

426

##### 427 2.1.1.1 URI Authority

428 In the context of an XRI, a URI authority is distinguished by the starting double slash ("//").

429

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

431

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

435

```
436 userinfo = *( unreserved / escaped / ";" /  
437 ":" / "&" / "=" / "+" / "$" / ", " )
```

438

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

441

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

443

444 Note that the host identifier may be omitted; if so a default may be defined by the semantics of a  
445 specific URI scheme. No default is specified by the XRI scheme.

446 A hostname, after the transformation described in step 4 of section 2.2.3.2, MUST meet the rules  
447 defined in section 3.2.2 of [RFC2396]. The productions for idomainlabel, qualified and hostname,  
448 therefore, have additional restrictions not reflected in the ABNF.

449

```
450 hostname = idomainlabel qualified
451 qualified = *( "." idomainlabel ) [ "." ]
452 idomainlabel = 1*ucschar
453 domainlabel = alphanum [ 0*61( alphanum / "-" ) alphanum ]
454 alphanum = ALPHA / DIGIT
```

455

```
456 IPv4address = dec-octet "." dec-octet "." dec-octet "." dec-octet
457 dec-octet = DIGIT ; 0-9
458 / %x31-39 DIGIT ; 10-99
459 / "1" 2DIGIT ; 100-199
460 / "2" %x30-34 DIGIT ; 200-249
461 / "25" %x30-35 ; 250-255
```

462

463 Support for an IPv6 address literal was added by [RFC2396bis] following the syntax originally  
464 specified in [RFC2732]. Note that because IPv6 literals use colons as delimiters, they must be  
465 encapsulated within square brackets. This is similar to the use of parentheses in XRI cross-  
466 references (see [ref Xref section]).

467

```
468 IPv6reference = "[" IPv6address "]"
469 IPv6address = 6( h4 ":" ) 1s32
470 / "::" 5( h4 ":" ) 1s32
471 / [ h4 ] ":" 4( h4 ":" ) 1s32
472 / [ *1( h4 ":" ) h4 ] ":" 3( h4 ":" ) 1s32
473 / [ *2( h4 ":" ) h4 ] ":" 2( h4 ":" ) 1s32
474 / [ *3( h4 ":" ) h4 ] ":" h4 ":" 1s32
475 / [ *4( h4 ":" ) h4 ] ":" 1s32
476 / [ *5( h4 ":" ) h4 ] ":" h4
477 / [ *6( h4 ":" ) h4 ] ":"
478 1s32 = ( h4 ":" h4 ) / IPv4address
479 ; least-significant 32 bits of address
480 h4 = 1*4HEXDIG
```

481

482 Lastly, a host identifier can be followed by an optional port number. XRI does not define a default  
483 port, so if the port is omitted in an XRI it is undefined.

484

```
485 port = *DIGIT
```

486

### 487 2.1.1.2 XRI Authority

488 In addition to the authorities supported in generic URI syntax, XRIs support two other  
489 mechanisms for specifying the global context of an identifier. The first is via global context  
490 symbols (GCS) and the second is via cross-references (abbreviated in the ABNF as "xref").

491

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

493

494 **2.1.1.2.1 Global Context Symbols (GCS)**

495 In support of the human-friendly identifier (HFI) requirements, XRIs offer a compact syntax for  
496 indicating the global context of an identifier. This approach uses the minimal possible metadata—  
497 a single prefix character—to provide the context for an XRI authority segment.

498

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

500

501 The global context symbol characters were selected from the set of symbol characters that are  
502 valid in a URI under **[RFC2396]** in order to represent the following global contexts:  
503

Symbol Character	Authority Type	Establishes global context for
+	General public	Identifiers for which there is no specific authority, i.e., that are established by public convention (e.g., in the English language, these would be the generic nouns).
=	Person	Identifiers that represent an individual person.
@	Organization	Identifiers that represent any authority other than the general public or an individual person.
\$	OASIS XRI TC	Identifiers established by the XRI specification for specific types of identifier metadata (e.g., language, version syntax, query syntax, etc.). See [ref Appendix B: Special Identifiers Assigned by the XRI-specification] for a list of these identifiers.
*	User-relative	Identifiers for which the authority is relative to the current user (i.e., "user-shortcut XRIs").

504

505 Note that because the global context symbol precedes an xri-segment and the xri-segment  
506 production allows cross-references (below), the global context symbols can be used with any type  
507 of authority specified under any URI scheme.

508 **2.1.1.2.2 Cross-References**

509 Cross-references are the primary extensibility mechanism in XRI. A cross-reference is either: a)  
510 an absolute URI, or b) a global XRI value. Note these are syntactically distinct because the  
511 former must start with a legal URI scheme, and consequently an ALPHA, while the latter must  
512 start with a symbol character. In either case, a cross-reference is enclosed in parentheses the  
513 same way an IPv6 literal is encapsulated in square brackets as specified in **[RFC2732]** (see  
514 section 2.1.1.1).

515

516 `xref-authority = xref ( "." sub-segment / ":" sub-segment ) * ( "."  
517 sub-segment / ":" sub-segment )  
518 xref = "(" ( global-xri / URI ) )"  
519 global-xri = global-path [ "?" xri-query ] [ "#" xri-fragment ]`

520

521 A cross-reference may appear at any node of any XRI except within a URI authority segment.  
522 When a cross-reference is used as the very first segment in an XRI, it enables any globally-

523 unique identifier in any URI scheme to specify an authority, e.g., an HTTP URI, mailto URI, URN,  
524 etc.

525 A cross-reference is also the means by which a XRI can be expressed as non-resolvable. To do  
526 this, the entire XRI is enclosed in parentheses. Note that this is the equivalent in the English  
527 language of putting a word or phrase in quotes to express that the author is referring to the word  
528 or phrase itself and not to its normal meaning. Examples:

529

```
530 The term "user-friendly" is used frequently in computing.  
531 --English-language usage of a quoted term  
532  
533 xri:(+user-friendly)  
534 --XRI equivalent of expressing this abstract concept
```

535

## 536 2.1.2 Path

537 As with URIs in general, the XRI path component is a hierarchal sequence of path segments  
538 separated by a slash ("/") character and terminated by the first question-mark ("?") or number  
539 sign ("#") character, or by the end of the XRI. The key difference is that while a URI path segment  
540 is considered opaque, an XRI path segment can have two types of sub-segments: dot-sub-  
541 segments and colon-sub-segments.

542

```
543 xri-segments = xri-segment *( "/" xri-segment )  
544 xri-segment = ( [ "." ] sub-segment / ":" sub-segment )  
545 *( "." sub-segment / ":" sub-segment )  
546 sub-segment = *xri-pchar / xref
```

547

548 Dot-sub-segments specify reassignable identifiers and colon-sub-segments specify persistent  
549 identifiers (following the lead of URN syntax in **[RFC2141]**). The default is a reassignable  
550 identifier, so no leading dot is required if this is the first (or only) sub-segment. **[ Does this  
551 distinction between reassignable and persistent "segments" need to be spun on out a bit more? ]**

552 An XRI path segment can contain the same characters as a URI path segment with the exception  
553 of the dot (".") and the colon (":"), which if used will be interpreted as described above. If this  
554 interpretation is not desired for these characters, or for any other special XRI delimiters, these  
555 characters **MUST** be escaped when they appear in the path segment. See [Ref to Escaping  
556 section].

557

```
558 xri-pchar = xri-unreserved / escaped / ";" / "!" / "*" /  
559 "@" / "&" / "=" / "+" / "$" / ", "
```

560

561 Other than dot-sub-segments and colon-sub-segments (and cross-references within these), an  
562 XRI path segment is considered opaque by generic XRI syntax. As with URIs in general, XRI  
563 extensions or generating applications may define special meanings for other URI reserved  
564 characters for the purpose of delimiting extension-specific or generator-specific sub-components.  
565 For example, section 3.4 of **[RFC2396]** specifies the set of URI reserved characters that can be  
566 used within a query segment.

## 567 2.1.3 Query

568 The XRI query component is identical to the URI query component as described in Section 3.4 of  
569 **[RFC2396]** with one exception: it may begin with a cross-reference. This permits the  
570 incorporation of metadata in XRI syntax describing the query string syntax. See [ref Appendix B:  
571 Special Identifiers Assigned by the XRI-specification] for more about query syntax identifiers.

572

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

574

575 The characters permitted in a query segment are the full set allowed in a URI path segment.

576

577 `pchar = unreserved / escaped / ";" /`  
578 `":" / "@" / "&" / "=" / "+" / "$" / ","`

## 579 2.1.4 Fragment

580 XRI syntax also supports fragments as described in Section 4.1 of [RFC2396] with the exception  
581 that it may begin with a cross-reference.

582

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

584

585 Fragments are supported primarily for compatibility with generic URI syntax, as XRI syntax can  
586 directly address attributes or secondary representations of a primary resource to any depth. XRIs  
587 can also use cross-references to identify media types or other alternative representations of a  
588 resource.

## 589 2.2 Characters

590 The character set and encoding of an XRI is primarily inherited from generic URI syntax as  
591 defined in [RFC2396] and clarified in [RFC2396bis], however it also includes the expanded  
592 character set defined in [IRI]. XRI characters fall into the same three subsets as URI characters.

593

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

### 595 2.2.1 Reserved Characters

596 XRI reserved characters are used to delimit XRI syntax components and thus are a superset of  
597 the URI reserved character set. Specifically, four characters have been added: opening  
598 parentheses ("("), closing parentheses (")"), dot ("."), and asterisk ("\*").

599

600 `xri-reserved = "/" / "?" / "#" / "[" / "]" / "(" / ")" / ";" / ":" /`  
601 `"," / "." / "&" / "@" / "=" / "+" / "*" / "$"`

602

603 If the use of an unescaped XRI reserved character as a data character would cause the  
604 interpretation of the XRI to be ambiguous, the character MUST be escaped as per the rules in [ref  
605 Escaping section].

### 606 2.2.2 Unreserved Characters

607 With the exception of the expanded UCS character set described in [IRI], the unreserved  
608 character set for XRIs is the same as that of URIs after the subtraction of the four characters  
609 noted above (all of which are in of the "mark" production of [RFC2396] and [RFC2396bis]).

610

611 `xri-unreserved = ALPHA / DIGIT / ucschar / xri-mark`  
612 `xri-mark = "-" / "_" / "!" / "~" / "'"`

613

614 The principle difference between XRI and URI reserved character sets is the inclusion of the UCS  
615 character set.

616

```
617 ucschar = %xA0-D7FF / %xF900-FDCF / %xFDF0-FFEF /  
618 %x10000-1FFFFD / %x20000-2FFFFD / %x30000-3FFFFD /  
619 %x40000-4FFFFD / %x50000-5FFFFD / %x60000-6FFFFD /  
620 %x70000-7FFFFD / %x80000-8FFFFD / %x90000-9FFFFD /  
621 %xA0000-AFFFFD / %xB0000-BFFFFD / %xC0000-CFFFFD /  
622 %xD0000-DFFFFD / %xE1000-EFFFFD
```

623

624 Escaping unreserved characters in an XRI does not change what resource is identified by that  
625 XRI. However, it may change the result of a URI comparison (see [ref Normalization and  
626 Comparison]), so unreserved characters should not be escaped unless necessary.

## 627 2.2.3 Escaped Characters

628 XRIs follow the same rules for escaping characters as URIs, i.e., any data in an XRI MUST be  
629 escaped if: a) it does not have a representation using an unreserved character, and b) using a  
630 reserved character would cause the XRI to be misinterpreted. An XRI thus escaped is said to be  
631 in “escaped normal form”. For consistency, all characters that are not in the ‘xri-unreserved’  
632 production and that are not used as syntactical elements as defined in this specification SHOULD  
633 be escaped. In this context, misinterpretation applies to XRIs used directly (i.e. not as URIs).  
634 Rules for converting an XRI into a legal URI are discussing in section 2.2.3.2. [ This last sentence  
635 is somewhat confusing. Could use some examples here. ]

### 636 2.2.3.1 Escaped Encoding

637 XRIs use the same percent-encoding as URIs as per section 2.4.1 of [RFC2396] and  
638 [RFC2396bis]. An escaped octet is encoded as a character triplet consisting of the percent  
639 character "%" followed by the two hexadecimal digits representing that octet's numeric value.

640

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

642

643 The uppercase hexadecimal digits 'A' through 'F' are equivalent to the lowercase digits 'a' through  
644 'f', respectively. XRIs that differ only in the case of hexadecimal digits used in escaped octets are  
645 equivalent. For consistency, uppercase digits SHOULD be used by XRI generators and  
646 normalizers.

### 647 2.2.3.2 Converting XRIs to URIs

648 Although XRIs can be used directly, there may be times when it is desirable to use an XRI in a  
649 context that expects a URI reference as defined by [RFC2396]. In other cases it may be desirable  
650 to use an XRI in a context that allows an identifier containing characters disallowed by [RFC2396]  
651 but which provides a simple mapping into a legal URI. The anyURI tag in defined in  
652 [XMLSchema2] is an example of the second case, where an escaping procedure is defined for  
653 characters that would otherwise be illegal under [RFC2396]. Additionally, [IRI] is a work-in-  
654 progress that proposes a new protocol element - an Internationalized Resource Identifier, or IRI-  
655 and defines the process for converting an IRI to a URI. IRI to URI conversion differs from the  
656 conversion defined for anyURI in [XMLSchema2] primarily in that it includes an algorithm  
657 appropriate for internationalized domain names. There may be cases in which it is desirable to  
658 use an XRI in a context that expects an IRI.

659 This specification defines the process for transforming an XRI into a legal URI. Depending on the  
660 target application, it may be appropriate to terminate the transformation process before the final  
661 step. If the target application expects an identifier defined as anyURI in [XMLSchema2], for



662 example, the transformation may terminate at the point at which the XRI has reached the  
663 threshold defined for protocol elements allowed under that specification. Where appropriate, the  
664 transformation steps below note such thresholds. Except for transformations specific to XRI  
665 syntax, these steps closely follow the algorithm proposed in [IRI].

666 Applications MUST map XRIs to URIs using the following steps (or any equivalent process that  
667 achieves the same result).

- 668 1. If the XRI is not encoded in UTF-8, convert the XRI to a sequence of characters encoded  
669 in UTF-8, normalized according to Normalization Form C (NFC) as defined in [UTR15].
- 670 2. Optionally add font and language metadata (see note below).
- 671 3. Perform XRI-specific conversion defined in section 2.2.3.3. At this point the identifier may  
672 be used as an IRI.
- 673 4. If the XRI has a 'hostname' component, replace it with the 'hostname' component  
674 converted using the ToASCII operation defined in section 4.1 of [RFC3490], with the  
675 UseSTD3ASCIIRules flag set to true and the AllowUnassigned flag set to false. At this  
676 point the identifier may be used as anyURI defined in [XMLSchema2] or in a comparable  
677 context.
- 678 5. Replace each character that is disallowed in URI references with escaped triplet(s) as  
679 described in section 2.2.3.1, one escaped triplet for each octet in the UTF-8 encoding of  
680 the disallowed character. At this point the identifier may be used as a generic URI.

681 A note on step two above. In some languages, a UTF-8 encoded string (i.e. a sequence of UTF-8  
682 encoded characters) does not contain enough information to determine how to properly render  
683 that string in the intended language. Specifically, to represent the glyph of a UTF-8 encoded  
684 character, language information and font information may be required. On the other hand, local  
685 language encoding always has the language and font information associated with it. To make it  
686 possible to revert back to the local language representation of an XRI, it may be necessary to  
687 record the language and font context of an XRI when converting to UTF-8. If UTF-8 encoding  
688 would lose information required to transform the XRI back into human readable form in the  
689 intended language and font, the transformation MAY include mark up by use of cross references  
690 containing the \$l and/or \$f identifier defined in Appendix B. Once the language and font context is  
691 declared it will be valid until it is reset by another \$l/\$f declaration.

692 The XRI-specific conversion described in step three is not idempotent, i.e. each time this step is  
693 applied it may yield different results. It is very important, therefore, that implementers are careful  
694 not to apply this step more than once since doing so may change the semantics of the identifier.  
695 In general, an application SHOULD use the least escaped version appropriate for the context in  
696 which the identifier appears. If the context, for example, allows an XRI directly, the identifier  
697 SHOULD be in escaped normal form described in section 2.2.3. If the context allows an IRI but  
698 not a XRI, the identifier SHOULD be in the form that results from step three, and so on.

699 The form of the XRI that results from each step in this section is equivalent to the result of any  
700 other step. In other words, applying this conversion does not change the equivalence of the  
701 identifier.

### 702 **2.2.3.3 XRI-specific conversion for use in URIs**

703 This section describes issues that can arise when an XRI is converted to URI. It looks only at  
704 issues specific to XRI syntax and not, for example, at international character issues. It also  
705 defines a conversion operation that performs the XRI-specific conversions required during the  
706 conversion of an XRI into a generic URI. This conversion operation must be done in conjunction  
707 with the steps defined in section 2.2.3.2 in order to effect a complete conversion from an XRI to a  
708 URI. In other words, the conversion in this section has very limited utility on its own. It is intended  
709 to be used as part of the larger conversion process described in section 2.2.3.2.

710 XRIs can contain other URIs as cross-references (see section [\[ref to cross-reference section\]](#)).  
711 These URIs can contain characters that, if unescaped, would cause misinterpretation when the  
712 XRI is converted to a URI. Consider the following XRI.

713

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

715

716 The generic parsing algorithm described in **[RFC2396]** would separate the above XRI into the  
717 following components

718

```
719 scheme = xri  
720 authority = <undefined>  
721 path = @example/(xri:@example2/abc?  
722 query = id=1)
```

723

724 The desired separation is

725

```
726 scheme = xri  
727 authority = <undefined>  
728 path = @example/(xri:@example2?id=1)  
729 query = <undefined>
```

730

731 To avoid this type of misinterpretation, certain characters in a cross-reference must be escaped  
732 when converting an XRI to a URI. In particular, cross-references must be converted such that the  
733 question mark “?” character is escaped as “%3F”, the number sign “#” character is escaped as  
734 “%28”, and the colon “:” character is escaped as “%3A”.

735 The example above, then, would be expressed as

736

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

738

739 A slash “/” character in a cross-reference can also be misinterpreted when the XRI is converted  
740 into a URI. Consider

741

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

743

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

746

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

748

749 instead of the intended

750

```
751 xri://example.com/
```

752

753 because the algorithm is defined in terms of the last (right-most) slash character. This problem is  
754 avoided by escaping slashes within cross-references as “%2F”. The above example, then, would  
755 be expressed as

756

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

758

759 Note that ambiguity is possible if an XRI in escaped normal form contains characters that have  
760 been escaped to indicate that they should not be interpreted in their normal syntactical sense. For  
761 example, consider the following XRI in escaped normal form

762

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

764

765 This slash character between 'c' and 'd' is escaped to show that it's not a syntactical element of  
766 the XRI, i.e. that it should be interpreted literally and not as a path separator. To preserve this  
767 type of distinction when converting an XRI to a URI, the percent "%" character must be escaped  
768 as "%25". The above example, fully converted, would be

769

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

771

772 The following, then, are the XRI-specific steps required to convert an XRI into a URI.

- 773 1. Escape all percent "%" characters as "%25" across the entire XRI.
- 774 2. Escape all number sign "#" characters that appear within a cross-reference as "%23".
- 775 3. Escape all question mark "?" characters that appear within a cross-reference as "%3F".
- 776 4. Escape all colon ":" characters that appear within a cross-reference as "%3A".
- 777 5. Escape all slash "/" characters that appear within a cross-reference as "%2F".

778 Note that the XRI must be in escaped normal form and all URIs in cross-references must be in an  
779 escaped form appropriate to their schemes before the above rules are applied.

#### 780 **2.2.3.4 Converting URIs to XRIs**

781 There may be times when it is desirable to convert an XRI in URI escaped form into an XRI in  
782 escaped normal form. This section gives a procedure to do such a conversion. Except for steps  
783 specific to XRIs, this procedure very closely follows the algorithm proposed by [IRI].

784 Conversion from an XRI in URI escaped form into an XRI in escaped normal form MUST use the  
785 following steps (or any equivalent process that achieves the same result).

- 786 1. If the identifier is not encoded in US-ASCII, convert it to a sequence of octets in US-  
787 ASCII.
- 788 2. If the identifier has a 'hostname' component, replace it with the UTF-8 encoded  
789 'hostname' component converted using the ToUnicode operation defined in section 4.2 of  
790 [RFC3490], with the UseSTD3ASCIIRules flag set to true and the AllowedUnassigned  
791 flag set to false.
- 792 3. Convert all escaped characters (as defined in section 2.2.3.1) with their corresponding  
793 octets, except for the percent "%" character, those characters in the 'reserved' production  
794 of [RFC2396] and US-ASCII characters disallowed in URIs by section 2.4.3 of  
795 [RFC2396].
- 796 4. Re-escape any octet produced in step 3 that is not part of a strictly legal UTF-8 octet  
797 sequence. [NOTE: This is verbatim from IRI. Is this ok? Should we elaborate?]
- 798 5. Perform the following XRI-specific conversions
  - 799 a. Convert all escaped slash "/" characters to their corresponding octets.
  - 800 b. Convert all escaped colon ":" characters to their corresponding octets.
  - 801 c. Convert all escaped question mark "?" characters to their corresponding octets.
  - 802 d. Convert all escaped number sign "#" characters to their corresponding octets.
  - 803 e. Convert all escaped percent "%" characters to their corresponding octets.

804 6. Encode the resulting sequence in UTF-8 (except for that portion already converted by  
805 step 3).

## 806 2.2.4 Excluded Characters

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

810

```
811 excluded = invisible / delims / unwise  
812 invisible = CTL / SP / %x80-FF  
813 delims = "<" / ">" / "%" / DQUOTE  
814 unwise = "{" / "}" / "|" / "\" / "^" / "`"
```

## 815 2.2.5 Legal Character Sequence

816 Not all ASCII sequences can be derived from UTF-8 sequences. A valid XRI character sequence  
817 MUST be derivable by escaping an equivalent UTF-8 sequence. [NOTE: This needs  
818 review/expansion.]

## 819 2.3 Character Encoding and Internationalization

820 The basic character encoding of XRI is UTF-8, as recommended by [RFC2718]. When an XRI is  
821 used as a human readable identifier, the representation of the XRI on the underlying document  
822 should use the character encoding of the underlying document. However, this string must be  
823 converted to UTF-8 before any further processing.

## 824 2.4 Relative XRIs

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

828 For a relative XRI reference that does not contain an authority component but whose base XRI  
829 contains an authority component that matches the URI-authority production, the rules for  
830 resolving relative references defined in section 5.2 of [RFC2396] apply.

831 For a relative XRI reference that does not contain an authority component but whose base XRI  
832 contains an authority component that matches the XRI-authority production, the rules defined in  
833 section 5.2 of [RFC2396] need modification because an XRI authority is considered opaque by  
834 generic URI syntax.

835 The following sections, therefore, define the process for resolving a relative XRI reference into a  
836 string that matches the XRI production defined in section 2.1 for all XRIs, including those relative  
837 references that would otherwise be unresolvable because they are considered opaque by  
838 [RFC2396].

### 839 2.4.1 Establishing a Base XRI

840 A base XRI is established according to the rules defined in section 5.1 of [RFC2396]. In other  
841 words, there is no difference between establishing a base XRI and establishing the base of any  
842 generic URI. [ Need to mention that XRIs are not URIs anyway, unless they are converted to URI  
843 form. ]

## 844 2.4.2 Obtaining the Referenced XRI

845 Section 5.2 of [RFC2396] describes rules for resolving relative references to absolute forms of  
846 URIs. For XRIs matching the XRI Authority production in [ref XRI Authority section], these same  
847 rules apply with the following modifications:

- 848 - In step 1, the XRI reference is parsed using an XRI aware parser such that the “authority  
849 component” is interpreted as the "authority-part" production defined in section 2.1.1 of  
850 this specification.
- 851 - Step 4 states, “If the authority component is defined, then the reference is a network-path  
852 and we skip to step 7”. For XRIs, the presence of an authority component does not imply  
853 that the reference is a network-path as defined by [RFC2396] because it may be an XRI-  
854 authority component. However, the instruction to skip to step 7 is still valid for XRIs. In  
855 other words, the processing instruction is correct, but the inference as to the type of  
856 reference is invalid.
- 857 - In step 4, the base XRI is parsed using an XRI aware parser such that the “authority  
858 component” is interpreted as the authority-part production defined in section 2.1.1 of this  
859 specification.

- 860 - In step 7, the block that reads

```
861 if authority is defined then  
862     append "/" to result  
863     append authority to result
```

864 is replaced by

```
865 if authority is defined then  
866     if type-of(authority) == URI-authority  
867         append "/" to result  
868     append authority to result
```

869

870 It is important to note that the algorithm described in section 5.2 of [RFC2396] will generally  
871 produce incorrect results when applied to relative XRI references in which the authority  
872 component matches the XRI-authority production. This type of relative XRI reference, therefore,  
873 should only be used in contexts in which the above algorithm is known to be employed. [ Example  
874 would be useful. ]

## 875 2.5 Normalization and Comparison

876 The scheme component is case-insensitive for comparison for XRIs and all URIs used as cross-  
877 references.

878 Comparison of authority components of two XRIs, as defined in 2.1.1, is case-insensitive for all  
879 characters in the ALPHA production.

880 Two XRI authority components, as defined in 2.1.1, are equivalent if they match using a case-  
881 insensitive comparison after applying steps one and three of the process described in section  
882 2.2.3.2.

883 Two XRIs MUST be equivalent if they are character-for-character equivalent. It follows, then, that  
884 they are equivalent if they are byte-for-byte equivalent when both XRIs use the same character  
885 encoding.

886 All forms of the XRI during the conversion process described in section 2.2.3.2 are equivalent.

887 Two XRIs that differ only in escaped unreserved characters are equivalent.

888 Each application that uses XRIs MAY define additional equivalence rules as appropriate.

889 Section 6 of [RFC2396bis] offers advice on more aggressive strategies for normalization and  
890 comparison as well as best practices for canonicalization of generic URIs. Implementers may find

891 this information useful in developing a strategy for establishing equivalence, particularly with  
892 respect to non-XRI cross-references.

---

## 893 **3 Resolution**

### 894 **3.1 Introduction to Resolution Architecture**

895 Resolution is the process of converting an XRI into data and metadata about the resource  
896 identified by the XRI.

897 Because XRIs will be used in a wide variety of deployments, communities, and applications, no  
898 single resolution mechanism is appropriate for all XRIs. Thus, a resolution framework and  
899 concrete implementations of that framework are defined. This framework MAY be required by  
900 communities to allow resolution of XRIs that they define. Other resolution mechanisms MAY be  
901 defined on a per-community basis.

902 It is important to note that XRIs can be "resolved" in a variety of ways. For example, they may be  
903 used as keys in a database, or used as filenames in a filesystem. The intent of this framework is  
904 to define an interoperable process for discovering and accessing data in an open system such as  
905 the Internet where such data may be distributed across a number of systems.

906 Policies for management of identifiers are defined on a community-by-community basis. Each  
907 community is identified via the authority portion of an XRI (which can be either a URI authority or  
908 an XRI authority as defined in Section 2.1.1). When a community chooses to create a new  
909 identifier authority, it SHOULD define a policy for how identifiers under this authority are assigned  
910 and managed. Furthermore, it SHOULD define what resolution scheme should be used for  
911 resolving those identifiers.

912 Resolution is defined as a set of interactions between a client and a series of "endpoints."  
913 Endpoints are networked systems that participate in XRI resolution using one or more of the  
914 implementations of the framework. These endpoints are usually discovered through processes  
915 defined in the framework, except for those which are defined by XRI communities as being the  
916 "community root" (which is effectively the starting place for the resolution process). These  
917 endpoints are advertised "out of band" to entities wishing to resolve identifiers in the identifier  
918 community.

919 The resolution framework here expects the XRI being resolved to have been converted into a  
920 URI-compatible form, following the rules in Section 2.2.3.2, "Converting XRIs to URIs."

#### 921 **3.1.1 Assumptions**

922 XRI resolution makes several minimal assumptions about XRIs:

- 923 • The endpoints representing the top-level authority for any globally unique XRI are  
924 identified with the "uri-authority" or "xri-authority" part of the XRI.
- 925 • Data corresponding to a single XRI may be retrieved or manipulated by multiple protocols  
926 at multiple endpoints.
- 927 • Each endpoint and protocol may present a different subset, type, or representation of the  
928 data and metadata associated with the identified resource.

#### 929 **3.1.2 Phases of Resolution**

930 The XRI resolution framework is designed to be as flexible as possible given the assumptions  
931 described above and the wide number of anticipated uses for XRIs. The framework reflects the  
932 structure of XRIs, and consists of two phases:

- 933 • Authority Resolution
- 934 • Local Access

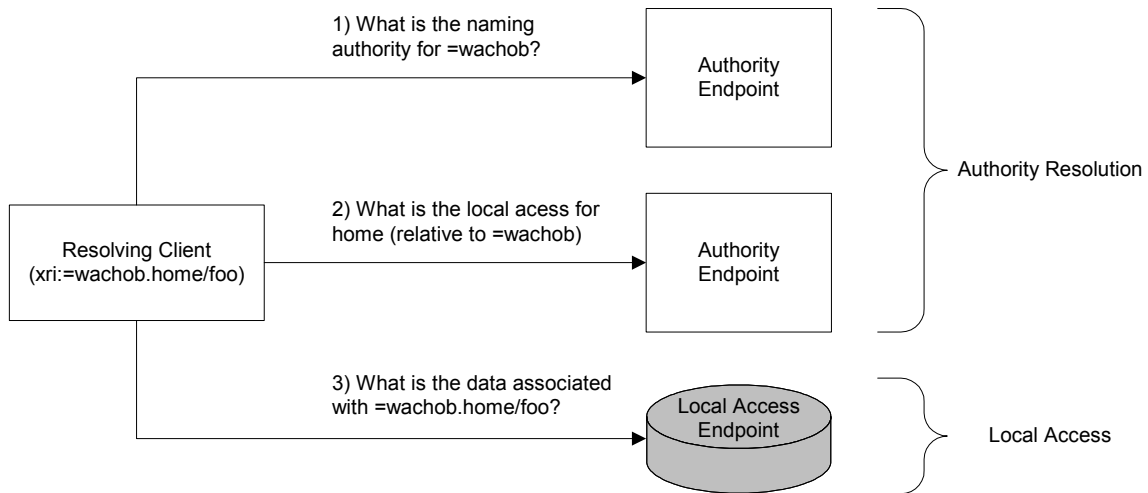
935 Authority Resolution is the process of finding the endpoint or endpoints representing the authority  
936 that controls the community of identifiers in which the XRI is defined. Authority Resolution results

937 in a list of local access descriptors, each describing an endpoint providing “local access” service.  
 938 These Local Access Descriptors are defined in section 3.3.1. Resolving clients choose an  
 939 endpoint described by one of these descriptors and select a local access protocol with which to  
 940 access that endpoint.

941 Figure 1 demonstrates the phases of XRI resolution:

942

943



944

945

Figure 1: Phases of Resolution

## 946 3.2 Phase 1: Authority Resolution

### 947 3.2.1 General Description

948 Authority Resolution is the process of finding a system representing an authority identified by the  
 949 “authority” component of the XRI. That component of the XRI can either be in the form of a DNS  
 950 name, an IP address, a GCS identifier, or a cross-reference. Each type of authority has a  
 951 separate method for resolution. DNS-specified authorities are resolved using DNS-based  
 952 Authority Resolution (DNAR) described in section 3.2.2 below. Authorities identified by GCS  
 953 identifiers or cross-references are resolved using the XRI-Authority Resolution Framework  
 954 (XARF) defined in section 3.2.3 below. Finally, authorities identified by an IP address are  
 955 resolved according to the process in section 3.2.4 below.

956

Type of Authority	XRI BNF Production	Description	Example	Method of Resolution
DNS- or IP-Address Specified	URI-authority	An authority identified by a DNS name or IP address.	<b>xri://example.com/foo.bar</b>	DNAR
Abstract	XRI-authority	An authority identified by a resolvable GCS	<b>xri:+example/foo.bar</b>	XARF



		identifier or a cross-reference.	
--	--	----------------------------------	--

957

Table 1: Types of Authority Identifiers

958

Whether DNAR or XARF is used, the result is a list of descriptors of local access endpoints.

959

These local access descriptors are defined in section 3.3.1 below. The resolving client can then

960

choose which endpoint it wishes to use to access data, attributes, or services associated with the

961

XRI.

962

### 3.2.2 DNS-Specified Authority Resolution (DAR)

963

The process for resolving DNS-specified authorities is a DDDS “application” which takes

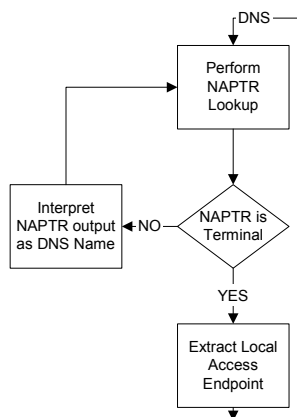
964

advantage of NAPTR and SRV records associated with DNS names.

965

DAR resolution is formally defined as a DDDS application with the following attributes:

<b>Application Unique String:</b>	The XRI being resolved
<b>First Well Know Rule:</b>	The DNS name specified in the Authority Name segment
<b>Flags:</b>	<p>“s” – a terminal flag which indicates that the result of the application of the regular expression (or replacement) is a DNS name pointing to an SRV record</p> <p>“u” – a terminal flag which indicates that the result of the application of the regular expression (or replacement) is a URI</p>
<b>Service Parameters:</b>	<p>This field is empty unless a terminal flag is present.</p> <p>This field contains a local access protocol descriptor as described in section 3.3.2 (“Format of Local Access Protocol Descriptors”). Initially, the only protocol descriptors using the ‘http’ local access service descriptor (defined in section 3.3.5.1 below) may be used. Local access service descriptors are unique to each use of a network protocol, and are described more fully in section 3.3.3 below.</p>
<b>Valid Databases:</b>	DNS from RFC 3403



966

967

Figure 2: DAR Algorithm

968

The DAR algorithm is as follows:

- 969 1) The DNS name that is the authority in the XRI is the “initial key” for the DDDS algorithm.  
970 The full XRI being resolved is the application unique identifier.  
971 2) The DNS name is resolved using a query type of “NAPTR”.  
972 3) The identifier is applied to the NAPTR record (either the regular expression or the  
973 replacement, following RFC 3402 rules).  
974 4) If the NAPTR record does not contain a terminal flag, then the result of step 3 is  
975 interpreted as a DNS name. The algorithm loops back to step 2 with the result of step 3  
976 as the DNS name to be resolved.  
977 5) If the NAPTR record contains a “u” flag, then the result of step 3 is interpreted as a URI.  
978 This URI is the network location of the local access service, and the “service” field is the  
979 local access protocol descriptor used for local access. DAR terminates.  
980 6) If the NAPTR record contains a “s” flag, then the result of step 3 is interpreted as a DNS  
981 name. This DNS name is resolved into a set of SRV records, each of which describes a  
982 network location of the local access service using a hostname and port number. The local  
983 access protocol descriptor is copied from the service field. These host and IP pairs  
984 corresponding to each SRV record are then reconstructed back into a URI based on the  
985 local access protocol descriptor and the SRV hostname. For example, if the DNS name  
986 that is the result of step 3 is “\_thttp.\_tls.\_tcp” and the local access protocol descriptor is  
987 “thttps+I2R”, and an SRV record contains a target of “thttp.xri.example.com” and a port of  
988 443, then the resulting network location would be https://thttp.xri.example.com/ as  
989 defined in section 3.3.5.1. After this network location is extracted, DAR terminates.  
990 [ The construction of a network endpoint from an SRV record probably needs more  
991 explanation and/or examples ]

## 992 3.2.3 XRI Authority Resolution Framework (XARF)

### 993 3.2.3.1 Introduction

994 The process for resolving abstractly-identified authorities (corresponding to the XRI-Authority  
995 BNF production) is comprised of an initialization step and an iterative series of operations. One of  
996 these iterations is performed for each node in the Authority component of the XRI. Each iteration  
997 consists of an invocation of a “relative lookup mechanism”. The term “relative” emphasizes the  
998 fact that only one node of the Authority component is being resolved at a time, from left to right.

999 This specification defines two of these relative lookup mechanisms: an HTTP-based mechanism  
1000 (“XRI-HTTP Relative Lookup Mechanism (NA1)”) and one derived from DDDS (“RDDDS Relative  
1001 Lookup Mechanism (NA2)”).

1002 The XARF algorithm results in a set of one or more Local Access Descriptors, the form of which  
1003 is defined in section 3.3.1.

### 1004 3.2.3.2 User Relative XRIs

1005 XRIs beginning with the user-relative community symbol (“\*”) are a special case for resolution.  
1006 The authority for these identifiers is defined by the user of the XRI, and not uniquely specified in  
1007 the XRI itself. Thus, these XRIs are not resolvable without the establishment of an authority for  
1008 the XRI from some source other than the characters in the XRI.

1009 XRIs beginning with the User-Relative community identifiers MUST be transformed into XRIs with  
1010 an explicit Authority identifier (other than one based on the user relative community) before they  
1011 can be resolved using the resolution mechanisms defined in this specification.

1012 Note that in most cases, this transformation is simply the replacement of the “\*” character with a  
1013 prefix corresponding to an Authority identifier. For example, if a client is configured with a default  
1014 community of “@employer”, then the xri “xri:\*workstation/identifier” would be converted into  
1015 “xri:@employer.workstation/identifier”.

1016 **3.2.3.3 Authority Descriptors**

1017 Endpoints that act as Authorities are described with Authority Descriptors, as defined below:

<b>Data Item</b>	<b>Description</b>	<b>Example</b>
Location	The network address where this endpoint can be reached. This is in the form of a URI, for “xri-http” authorities and in the form of a DNS URI for “rddds” authorities.	dns:authority.example.com
Protocol Descriptor	The protocol to use when communicating with this authority endpoint. See the definition of Authority Protocol Descriptor at section 3.2.3.4	na2

1018 Table 2: Authority Descriptor

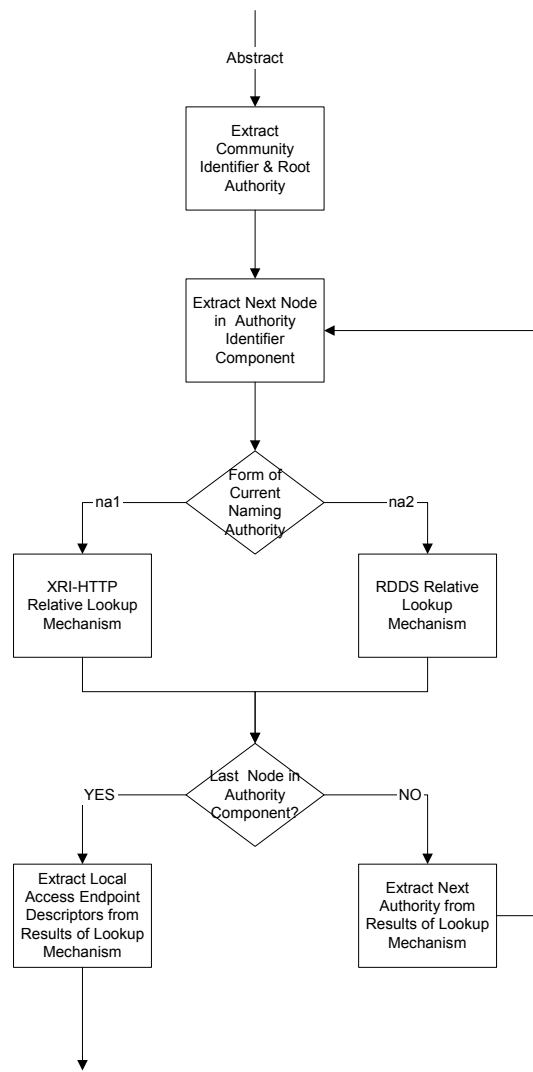
1019 **3.2.3.4 Authority Protocol Descriptors**

1020 Each Authority Relative Lookup Mechanism defines a short string to uniquely identify the protocol  
1021 used to perform a lookup at the Authority. There are currently two Authority Protocol Descriptors  
1022 defined:

<b>Protocol</b>	<b>Authority Protocol Descriptor</b>
XRI-HTTP	na1
RDDDS	na2

1023 Table 3: Authority Protocol Descriptors

### 3.2.3.5 Algorithm



1025

1026

Figure 3: XARF Algorithm

1027 The XARF algorithm consists of an initialization step and repeated invocations of a relative lookup  
 1028 mechanism. Each invocation of the relative lookup mechanism operates on a specific node (the  
 1029 “current node”) of the XRI Name Authority component, progressing from left to right.

1030 The initialization step is the extraction of the community identifier from the Authority component of  
 1031 the XRI. This first node may specify a global community (using a global community identifier) or a  
 1032 privately defined community (using a cross reference). In the case of a global community  
 1033 identifier, the global community identifier is treated as a separate node of the authority identifier.  
 1034 Also, every global community identifier has one or more associated Authority Descriptors. In the case  
 1035 of a cross reference, there must be one or more Authority Descriptor associated with that  
 1036 cross reference. The discovery of these associated Authority Descriptors is an out of band  
 1037 process that is assumed to have taken place when the resolving framework is deployed.

1038 After the initial Authority Descriptor is selected, the next node of the Authority becomes the initial  
 1039 value for the “current node” in the iteration described below.

1040

<b>XRI</b>	xri:@example.internal/foo
<b>Authority</b>	@example.internal
<b>Community Identifier</b>	@
<b>First Node Resolved</b>	.example

1041

Table 4: Global Community Identifiers

1042

<b>XRI</b>	xri:(http://www.example.com).internal/foo
<b>Authority</b>	(http://www.example.com).internal/foo
<b>Community Identifier</b>	(http://www.example.com) [Would this be escaped in URI form?]
<b>First Node Resolved</b>	.internal

1043

Table 5: Cross-Reference Community Identifiers

1044

1045 Each XARF iteration begins with an Authority Descriptor, and results in a set of Authority  
 1046 Descriptors or Local Access Descriptors.

1047 If the node being operated on is the last node in the XRI Name Authority component, then the  
 1048 results should be a list of Local Access Descriptors.


1049 If the node being operated on is not the last node in the XRI Name Authority component, then the  
 1050 results of the iteration should be a list of Authority Descriptors. The next iteration will use one of  
 1051 these Authority Descriptors from the pervious iteration to perform the lookup on the next node in  
 1052 the XRI Authority.

1053 Each iteration consists of the following steps:

- 1054 1) Select a Authority Descriptor which is available from the previous iteration, or, if this is the  
 1055 first iteration, select a Authority Descriptor from those configured to correspond to the first  
 1056 node of the XRI Authority component as described above.
- 1057 2) Perform the relative lookup mechanism specified in the Authority Descriptor using the  
 1058 current node as the query for the lookup. Currently, there are two of these relative lookup  
 1059 mechanisms: XRI-HTTP and RDDDS.
- 1060 3) If the current node is not the last node in the XRI Authority component, then the results of  
 1061 step 2 are a set of Authority Descriptors. The loop repeats at step 1, using the next node of  
 1062 the XRI Authority component, and using the Authority Descriptors just retrieved from  
 1063 step 2.
- 1064 4) If the current node is the last node in the XRI Authority component, then the results of  
 1065 step 2 must be a set of Local Access Descriptors. XARF terminates with these Local  
 1066 Access Descriptors as the result.

1067 If step 2 results in an empty list of Descriptors, this is equivalent to an unresolvable XRI Authority  
 1068 component, and SHOULD reported to the user of the resolver as an error.

1069 **3.2.3.6 XRI-HTTP Relative Lookup Mechanism (NA1)**



1070  
1071

Table 6: XRI-HTTP Algorithm

1072 The XRI-HTTP relative lookup mechanism performs a simple HTTP GET to resolve a node of an  
1073 XRI Authority component. This relative lookup mechanism has the Authority Protocol Descriptor  
1074 “na1”. Any Authority Descriptor with the Authority Protocol Descriptor “na1” MUST contain a  
1075 network location of the form of a HTTP or HTTPS URI.

1076 A resolver performing the XRI-HTTP relative lookup mechanism constructs a request URL from  
1077 the Authority Descriptor location field and the current node (from the XARF iteration). This URL is  
1078 a concatenation of the location field and the current node as described below:

1079  
1080  
1081

```
nal-url = authority-location ?(“/”) url-escape(current-node)
```

1082 The separator “/” is inserted between the authority-location in the descriptor and the current node  
1083 only if the authority location does not end in a “/”. The current node must be “URL-escaped”  
1084 [\[reference\]](#) before being inserted into the request URL.

1085 The resolving client MUST use the HTTP/1.1 protocol. All HTTP semantics are available to the  
1086 resolving client and Authority endpoint. Specifically, redirects, security, caching and other HTTP-  
1087 defined semantics should be employed where necessary. However, for the purposes of  
1088 interoperability and ease of implementation, use of such features should be minimized to the  
1089 extent possible.

1090 The content of the result of the HTTP request is a document that contains a list of newline-  
1091 separated [\[better way to say this?\]](#) lines of plain text. The document is of content-type “text/plain”.  
1092 Each line is of the following form:

1093  
1094  
1095

```
nal-result = protocol-descriptor space location newline
```

1096 The value of every service-descriptor field is MUST be a legal value defined for Authority Protocol  
1097 Descriptors (section 3.2.3.4) or Local Access Protocol Descriptors (section 3.3.2). The value of  
1098 every location field is constrained by the definition of the service descriptor. For example, if the  
1099 service descriptor is “na2”, then the location must be a DNS URI.

1100 Each line is parsed and transformed into an Authority Descriptor or Local Access Descriptor  
1101 depending on whether the current node is the last node of the XRI Authority component. The  
1102 result of each XRI-HTTP invocation is a list of descriptors.

1103  
1104  
1105

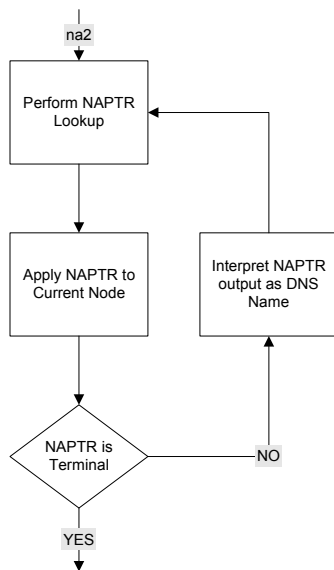
```
Non-Terminal Results Example:
```

1106 The name authority is available at a HTTP na1 endpoint, an HTTPS na2  
 1107 endpoint, and at a domain name for use with na2 (DDDS).  
 1108  
 1109 na1 http://next.step.example.com/resolve/b/  
 1110 na1 https://alternate.step.example.com/resolve/c/  
 1111 na2 dns:my.nazone.example.com

1112  
 1113 Terminal Results Example:  
 1114  
 1115 There is a thhttp-based I2R service for accessing wsd1 associated with  
 1116 the XRI and a ldap-based I2R service for accessing wsil associated with  
 1117 the XRI.  
 1118  
 1119 thhttp+I2R/wsd1 http://authority.example.com/uri-res/  
 1120 ldap+I2R/wsil ldap://user@foo.ldap.example.com/

1121  
 1122 Complete HTTP Request-Response Example:  
 1123 The resolving client is performing an "na1" request on the URL  
 1124 <http://xrib.example.com/naresolve> with the relative identifier "c"  
 1125  
 1126 Client to Server:  
 1127 GET /naresolve/c HTTP/1.1  
 1128 Host: xrib.example.com  
 1129  
 1130 Response to Client from Server:  
 1131 HTTP/1.1 200 OK  
 1132 Content-Type: text/plain  
 1133  
 1134 na1 http://next.step.example.com/resolve/b/  
 1135 na1 https://alternate.step.example.com/resolve/c/  
 1136 na2 dns:my.nazone.example.com

1137 **3.2.3.7 RDDDS Relative Lookup Mechanism (NA2)**



1138  
 1139 Figure 4: RDDDS Algorithm

1140 **3.2.3.7.1 Introduction**

1141 This lookup mechanism is a DDDS application as described in RFC 3401, and has the Authority  
 1142 Protocol Descriptor “na2”. Each invocation of RDDDS on a node of the XRI Authority component  
 1143 is a complete run of the DDDS algorithm as defined in RFC 3402. This algorithm takes an  
 1144 identifier and finds an Authority for that identifier.

1145 Note that while RDDDS is technically compliant with RFC 3402, it is not in conformance to the  
 1146 original intent of the DDDS specification because it performs resolution on only part of the XRI  
 1147 (i.e. a single node in the XRI Authority component), and does not apply the entire XRI to the  
 1148 NAPTR records retrieved in the DDDS algorithm.

1149 RDDDS uses the DDDS algorithm straightforwardly, but there are some concepts that must be  
 1150 mapped from this specification to the DDDS suite of specifications. The following table describes  
 1151 the conceptual mapping:

DDDS Concept	XRI Concept	Description
Application-unique string	Current node	Each invocation of RDDDS corresponds to a single node in the XRI Authority component, and thus appears to the DDDS algorithm as the entire application unique string.
First Well-Known Rule	Location from the current Authority Descriptor	Each invocation of RDDDS begins with an Authority Descriptor that contains a DNS name. This DNS name is how the DDDS algorithm initially queries DNS for the first iteration of the DDDS algorithm.

1152 Table 7: Mapping DDDS and XRI Concepts for RDDDS

1153 The RDDDS relative lookup mechanism is formally defined here as an DDDS application:

<b>Application Unique String:</b>	The current node from the XRI Authority Resolution Framework
<b>First Well Know Rule:</b>	The first key is AuthorityLocation
<b>Flags:</b>	<p>“s” (a terminal flag) which signifies that the result of the NAPTR regex/replacement is a domain name which has one or more SRV records associated with it. <b>The use of SRV records needs further specification – also see SRV records in section 3.2.2.</b></p> <p>“u” (a terminal flag) which signifies that the result of the NAPTR regex/replacement is a URI (see RFC 3404). <i>This needs further fleshing out – see description above.</i></p> <p>The absence of a terminal flag means that the DDDS resolution continues with the domain name that is a result of applying the regex/replacement to the current node. In this case, DDDS continues with the key equal to this domain name.</p>
<b>Service Parameters:</b>	<p>If this DDDS resolution is NOT the last step in the XARF protocol, then the service must indicate a XARF resolution mechanism (e.g. ‘na1’ or ‘na2’ that are defined in this document). This indicates that the client resolver may use the indicated resolution proctol and endpoint for the next iteration in the XARF protocol.</p> <p>If this is the last iteration in the XARF protocol, the service types must indicate a local access protocol. In this case, see Section</p>



	3.3.2, "Format of Local Access Protocol Descriptors" for the format of the service parameters.
<b>Valid Databases:</b>	DNS from RFC 3403

1154

Table 8: Formal RDDDS Definition

1155

### 3.2.3.7.2 Algorithm

1156

As an implementation of the DDDS algorithm, each invocation of RDDDS consists of the execution of an iterative step one or more times. Each iteration of the loop begins with a current DNS name, along with the current node of the XRI Authority component that is being resolved. The "current node" does not change across the entire invocation of RDDDS; this algorithm operates on one node at a time. The initial DNS name is extracted from the current Authority Descriptor. The result of the RDDDS invocation is a set of Local Access Descriptors or Authority Descriptors as described in Section 3.2.3, "XRI Authority Resolution Framework (XARF)".

1162

The steps for the loop inside the RDDDS lookup mechanism are:

1164

- 1) A NAPTR DNS query is performed using the current DNS name, resulting in a set of NAPTR records.
- 2) The set of NAPTR records is iterated through, according to the algorithm in RFC 3402, section 3.3 (??) until a NAPTR record that matches the current node (using the regex or replacement fields) and matches any requirements for service type defined by the resolver. *(Needs more formalization, but the idea is that the first matching NAPTR is used, based on the order and preference fields of the NAPTR records)*
- 3) The current node is applied to the regex or replacement field (as the application unique string) to the first NAPTR record that matches from step 3.
- 4) If the matched NAPTR record from step 3 does NOT contain any flags, then the result of the regex substitution (or replacement) from step 3 is a DNS name. The current DNS name is set to the result of the regex substitution (or replacement) from step 3. The algorithm jumps back to step 1.
- 5) If the NAPTR record has the "u" flag, then the result of applying the CurrentName to the regex in the NAPTR is a URI describing a network location. For each service listed in this NAPTR record, a new descriptor is created. This descriptor may be a Local Access Descriptor (if the current node is the last node of the XRI Authority component), or a Authority Descriptor (if the current node is not the last node of the XRI Authority component). Both Authority Descriptors and Local Access Descriptors have Location and Protocol Descriptor fields. The Protocol Descriptor field is set to the content of the NAPTR record's service field. The Location field is set to the URI value that results from the regex substitution (or replacement). RDDDS then terminates with this new descriptor as a result.
- 6) If the NAPTR record contains a "s" flag, then the result of step 3 is interpreted as a DNS name. This DNS name is resolved into a set of SRV records, each of which describes a network location using a hostname and port number. As with step 5, a new descriptor is created from each SRV record. The protocol descriptor is copied from the service field of the matching NAPTR record. The host and IP pairs corresponding to each SRV record are then reconstructed back into a URI based on the protocol descriptor and the SRV hostname. For example, if the DNS name that is the result of step 3 is "\_tthttp.\_tls.\_tcp" and the local access protocol descriptor is "thttps+I2R", and an SRV record contains a target of "thttp.xri.example.com" and a port of 443, then the resulting network location would be https://thttp.xri.example.com/ as defined in section 3.3.5.1. RDDDS then terminates with this set of descriptors as a result. *The use of SRV records needs further specification – also see SRV records in section 3.2.2.*

1198

1199 **3.2.4 IP-Address Authority Resolution (IAR)**

1200 [Its not clear what the use case is here, but for consistency this section defines a way of  
1201 resolving identifiers at a "IP-address specified" authority.]

- 1202 • IP address defines the local access endpoint
- 1203 • Use Local Access protocol "thttp" binding below

1204 **3.3 Phase 2: Local Access**

1205 Local Access is the process of asking an authoritative endpoint to do something with the  
1206 identifier. Local Access protocols typically are instances of data access or directory lookup  
1207 protocols.

1208 After performing the Authority Resolution step, a resolving client will choose which of the Local  
1209 Access protocols and endpoints it wishes to use for the Local Access phase of resolution. This  
1210 decision will be based on several factors:

- 1211 • The type of data the client is looking for. This is akin to query types in DNS. A client is  
1212 assumed to be looking for a particular type of information about the identifier. The data  
1213 type associated with an endpoint is sometimes available as part of the Local Access  
1214 Protocol Descriptor as described in section 3.3.2.
- 1215 • The protocol through which each endpoint can be accessed. Clients may only implement  
1216 a subset of Local Access protocols, or have preferences for certain Local Access  
1217 protocols. Clients SHOULD implement at least the protocols described in this document.
- 1218 • The identity of the network endpoint. Clients may choose different network endpoints  
1219 because they have other knowledge about those endpoints, such as previous failed  
1220 attempts to access the endpoint, or the security features associated with a particular  
1221 endpoint (ie HTTPS vs. HTTP).

1222

1223 Example of Choosing an Endpoint:

- 1224 1. A resolving client is given the option of accessing a Local  
1225 Access endpoint using LDAP or HTTP. Because the client has not  
1226 implemented LDAP, it chooses the HTTP endpoint
- 1227 2. A resolving client is given the option of accessing a Local  
1228 Access endpoint that provides WSDL data and one that provides  
1229 RDDL data. Because the client is interested in discovering  
1230 SOAP messaging endpoints, it chooses the Local Access endpoint  
1231 that provides WSDL.

1232 **3.3.1 Format of Local Access Descriptors**

1233 All Local Access endpoints must be described with sufficient detail to allow resolving clients to  
1234 make the sort of decisions described above. To do that, all Local Access endpoints are described  
1235 with Local Access Descriptors.

1236 A **Local Access Descriptor** is a pair of data items as described in the following table:

Data Item	Description	Example
Location	The network address where this endpoint can be reached. This is in the form of a URI	http://xri.example.com/I2R
Protocol Descriptor	The protocol to communicate with. This field optionally includes the type of data available at this endpoint. See the definition of "Local Access Protocol Descriptor" at section 3.3.2	thttp+i2r/wsdl

1237

Table 9: Local Access Descriptor

### 1238 3.3.2 Format of Local Access Protocol Descriptors

1239 The format for a local access protocol descriptor follows that from RFC 3404, but with a  
1240 modification to include the data type associated with the service (if needed).

1241

```
1242 service_field = protocol * ("+" rs)  
1243 protocol      = ALPHA *31ALPHANUM  
1244 rs           = ALPHA *31ALPHANUM * ("/" type)  
1245 type        = ALPHA *31ALPHANUM
```

1246

1247 The protocol and type fields are limited to 32 characters. The protocol, rs, and type fields must  
1248 start with an alphabetic character.

1249 Note that the protocol element is always required for local access protocol descriptors. This  
1250 specification does not enumerate legal data “types”. Communities wishing to use XRI identifiers  
1251 SHOULD enumerate which, if any, data “types” are legal for that community.

1252

1253 Examples (the type descriptors here are hypothetical and not yet  
1254 defined):

1255

1256 To describe a thttp-based service to return a wsdl document  
1257 corresponding to an XRI:

1258

```
1259 thttp+I2R/wsdl
```

1260

1261 To describe a thttp-based service to return a canonical XRI for a XRI:

1262

```
1263 thttp+I2I
```

1264

1265 To describe a ldap-based service to return a WSIL document  
1266 corresponding to an XRI:

1267

```
1268 ldap+I2R/wsdl
```

### 1269 3.3.3 Local Access Service Descriptors

1270 Local access service descriptors are short strings that unambiguously identify the use of a  
1271 network protocol as a XRI Local Access mechanism, including the selection of any options that  
1272 the protocol may otherwise provide. This use of a protocol is called a “Local Access binding”.  
1273 Usually, there is one Local Access Service Descriptor for each Local Access binding (as  
1274 described in section 3.3.4 below). Local Access bindings may define multiple Local Access  
1275 Service Descriptors to provide different options on using the protocol described in the Local  
1276 Access binding. For example, if a local access service provides an unauthenticated version and  
1277 an authenticated version, there should be separate service descriptors for each. Local Access  
1278 service descriptors make up part of a Local Access protocol descriptor, as described in section  
1279 3.3.2

### 1280 3.3.4 Requirements for Local Access Bindings

1281 Local Access bindings are required to unambiguously describe the use of a network protocol for  
1282 local access. It is expected that most Local Access bindings will refer largely to underlying  
1283 network protocols such as HTTP, LDAP, or SOAP. However, in all cases, there are aspects of  
1284 using such an underlying network protocol that must be explicitly specified for use with XRIs.

1285 Thus, an XRI local access binding specifies:

- 1286 • The Local Access Service Descriptors that identify use of this binding. For example, the  
1287 THTTP Local Access binding defines “thttp” to refer to the THTTP binding.

- 1288
- 1289
- 1290
- 1291
- 1292
- 1293
- 1294
- 1295
- 1296
- 1297
- 1298
- The underlying network protocol used to access the Authority. Examples would be HTTP and LDAP.
  - How the XRI is mapped to protocol-specific fields. For example, if using LDAP as the underlying network protocol, the binding must describe how is the XRI mapped to a LDAP query.
  - What sorts of interaction are possible with the Authority using this binding. Some Local Access protocols will be "read-only", while others will be "read/write".
  - The range of data types the Local Access protocol can accommodate. Some bindings may deal only with a certain type of data, but usually the type of data a particular binding supports is unlimited. For example, if a Local Access binding uses LDAP, then potentially any type of data can be accessed via that binding.

1299 **3.3.5 Local Access Bindings**

1300 This specification defines one useful Local Access binding using RFC2438. It is expected that  
1301 other Local Access bindings will be defined in separate specifications.

1302 **3.3.5.1 THTTP Local Access Binding**

1303 The functionality and content of the response of this local access protocol is defined in RFC 2438.  
1304 Generally, this protocol defines a method for getting a variety of types of data corresponding to an  
1305 URI. In its use here, it can be used to retrieve data about an XRI.

1306 THTTP is based on a simple HTTP 1.1 GET request. THTTP defines the Local Access Service  
1307 Descriptors "thttp" for use with HTTPS URLs and "thttps" for use with HTTPS URLs in this binding  
1308 specification.

1309 The actual GET request is quite simple, and is constructed from several pieces of data. All  
1310 semantics of the GET are inherited from RFC 2169 and RFC 2438, except as described below.

1311 The HTTP URL is constructed from the Location field of the Local Access Protocol Descriptor.  
1312 The Local Access Protocol Descriptor field is also used in constructing the request URL. The  
1313 Local Access Protocol Descriptor is broken up into "protocol", and a series of "rs", and "type"  
1314 fields (as described in section 3.3.2). A single pair of rs and type fields are chosen which specify  
1315 the THTTP service type and the data type the client wishes to perform.

1316

1317 `thttp-url = location ?("/"/) rs "/" ?(type "/") url-encode(xri-local-  
1318 path)`

1319

1320 The separator "/" is inserted between the authority-location in the descriptor and the current node  
1321 only if the authority location does not end in a "/". The type field (and trailing slash) are inserted  
1322 only if there is a type that corresponds to the selected rs in the Local Access Protocol Descriptor  
1323 field. The local path node must be "URL-escaped" (reference) before being inserted into the  
1324 request URL.

1325 The result of the THTTP request is defined by RFC 2438 (RFC 2169?)

1326

1327 **Example:**

1328

1329 Suppose the local access phase is begun with a Local Access Descriptor  
1330 the xri "xri://a.b.c/1.2.3" containing the following fields:

1331

1332 Location: <http://xriaccess.example.com>

1333

1334 ProtocolDescriptor: thttpd+I2R/wsdl+I2R/wsdl

1335

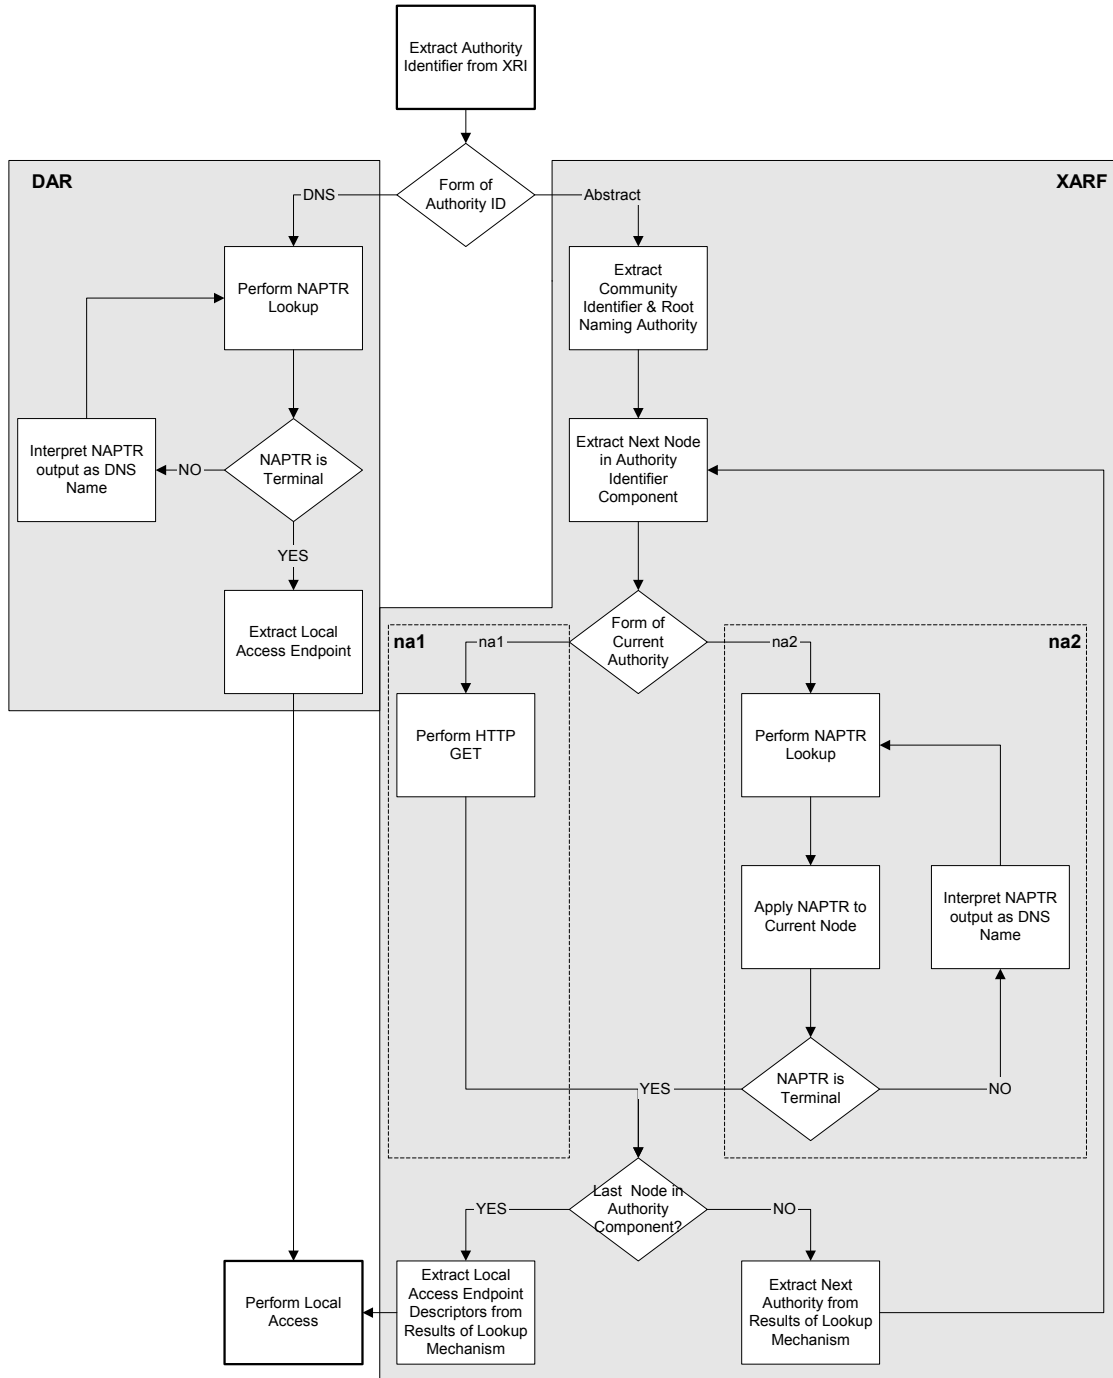
1336 Then a local access request for wsdl associated with the XRI using the  
1337 thttpd I2R service would be invoked by performing a HTTP GET to the  
following URI:

1338  
1339  
1340  
1341

http://xriaccess.example.com/I2R/wsd1/1.2.3  
The result would be a WSDL document associated with xri://a.b.c/1.2.3

1342  
1343

### 3.4 Flowchart of Authority Resolution



1344

---

## 1345 4 Security and Data Protection

### 1346 4.1 XRI Usage in Legacy Infrastructure

1347 Where XRIs are used within the legacy (pre-XRI) Internet and computing infrastructure, the  
1348 security and data protection considerations relating to XRIs are similar to those of other URI  
1349 schemes. In this context the material in section 7, *Security Considerations*, of [RFC2396bis] is  
1350 informative. It include a discussion of the following topics:

- 1351 • Reliability and Consistency
- 1352 • Malicious Construction
- 1353 • Rare IP Address Formats
- 1354 • Sensitive Information
- 1355 • Semantic Attacks

1356 This material notes that “a URI does not in itself pose a direct security threat.” This statement  
1357 remains true only for the use of XRIs in legacy environments, and may not be accurate as new  
1358 infrastructure evolves that takes full advantage of the extensibility of XRI architecture.

### 1359 4.2 Secure Resolution

1360 The resolution mechanisms described in section 3 are not intrinsically trustworthy. It is expected  
1361 that, in practice, some combination of DNSSEC, SSL and other existing technologies will be  
1362 employed to increase the security of the resolution process. Such considerations are outside the  
1363 scope of this document, although follow-on work may be done to define best practices and  
1364 facilitate inoperability.

### 1365 4.3 XRI Usage in Evolving Infrastructure

1366 As XRIs are adopted as abstract identifiers, it is anticipated that new services will be developed  
1367 that take advantage of their extensibility. In particular, XRIs may enable new solutions to security  
1368 and data protection problems that are not possible using existing URI schemes.

1369 For example, XRI cross-reference syntax permits the inclusion of identifier metadata such as an  
1370 encrypted or integrity-checked path, query, or fragment. Cross-references can also be used to  
1371 indicate methods of obfuscating, proxying, or redirecting resolution to prevent the exposure of  
1372 private or sensitive data. These capabilities may enable new security and data protection features  
1373 at the fundamental level of resource identifiers.

1374 A complete discussion of this topic is out of scope for this document. However, as a consequence  
1375 of the extensibility of XRIs, it is not possible to make definitive statements regarding all security  
1376 and data protection considerations relating to XRIs.

---

## 5 References

### 5.1 Normative

- 1379 [RFC2396] T. Berners-Lee, R. Fielding, L. Masinter, *Uniform Resource Identifiers (URI):*  
1380 *Generic Syntax*, <http://www.ietf.org/rfc/rfc2396.txt>, RFC 2396, August 1998.
- 1381 [XMLSchema2] P. Biron, A. Malhotra, *XML Schema Part 2: Datatypes W3C Recommendation*,  
1382 <http://www.w3.org/TR/xmlschema-2/>, May 2001.
- 1383 [RFC2119] S. Bradner, *Key words for use in RFCs to Indicate Requirement Levels*,  
1384 <http://www.ietf.org/rfc/rfc2119.txt>, RFC 2119, March 1997.
- 1385 [XML] T. Bray, J. Paoli, C.M. Sperberg-McQueen, E. Maler, *Extensible Markup*  
1386 *Language (XML) 1.0 (Second Edition) W3C Recommendation*, <http://www.w3.org/TR/REC-xml>,  
1387 October 2000.
- 1388 [RFC2234] Crocker, D.H. and Overell, P., *Augmented BNF for Syntax Specifications: ABNF*,  
1389 <http://www.ietf.org/rfc/rfc2234.txt>, RFC 2234, November 1997.
- 1390 [UTR15] M. Davis, M. Duerst, *Unicode Normalization Forms*,  
1391 <http://www.unicode.org/unicode/reports/tr15/tr15-23.html>, April 17, 2003.
- 1392 [RFC3490] P. Faltstrom, P. Hoffman, A. Costello, *Internationalizing Domain Names in*  
1393 *Applications (IDNA)*, <http://www.ietf.org/rfc/rfc3490>, RFC 3490, March 2003.
- 1394 [RFC2732] R. Hinden, B. Carpenter, L. Masinter, *Format for Literal IPv6 Addresses in URL's*,  
1395 <http://www.ietf.org/rfc/rfc2732.txt>, RFC 2732, December, 1999.
- 1396 [RFC2718] L. Masinter, H. Alvestrand, D. Zigmond, R. Petke, *Guidelines for New URL*  
1397 *Schemes*, <http://www.ietf.org/rfc/rfc2718.txt>, RFC 2718, November 1999.
- 1398 [RFC3305] M. Mealing, R. Denenberg, *Uniform Resource Identifiers (URIs), URLs, and*  
1399 *Uniform Resource Names (URNs): Clarifications and Recommendations*,  
1400 <http://www.apps.ietf.org/rfc/rfc3305.html>, RFC 3305, August 2002.
- 1401 [RFC2141] R. Moats, *URN Syntax*, <http://www.ietf.org/rfc/rfc2141.txt>, IETF RFC 2141, May  
1402 1997.
- 1403 [UML] Object Management Group, *Unified Modeling Language (UML) Version 1.5*,  
1404 <http://www.omg.org/technology/documents/formal/uml.htm>, March 1, 2003.
- 1405 [RFC1737] K. Sollins, L. Masinter, *Functional Requirements for Uniform Resource Names*,  
1406 <http://www.ietf.org/rfc/rfc1737.txt>, RFC 1737, December 1994.
- 1407 [Unicode] The Unicode Consortium, *The Unicode Standard, Version v3.0*, Addison-Wesley  
1408 Pub Co; ISBN: 0201616335, February, 2000.

### 5.2 Informative

- 1410 [IRI] M. Duerst, M. Suignard, *Internationalized Resource Identifiers (IRIs)*,  
1411 <http://www.ietf.org/internet-drafts/draft-duerst-iri-04.txt>, Work-In-Progress, June 2003.
- 1412 [RFC2396bis] R. Fielding, *Uniform Resource Identifiers (URI): Generic Syntax*, Internet Draft  
1413 draft-fielding-uri-rfc2396bis-03, <http://www.apache.org/~fielding/uri/rev-2002/rfc2396bis.html>,  
1414 Work-In-Progress, June 2003.
- 1415 [XRIReqs] G. Wachob, D. Reed, M. Le Maitre, D. McAlpin, D. McPherson, *Extensible*  
1416 *Resource Identifier (XRI) Requirements and Glossary v1.0*, [http://www.oasis-](http://www.oasis-open.org/apps/org/workgroup/xri/download.php/2523/xri-requirements-and-glossary-v1.0.doc)  
1417 [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),  
1418 June 2003.

1419

## Appendix A. Collected ABNF for XRI

1420 This section contains the complete ABNF for XRI, which includes the complete ABNF for URI  
 1421 from **[RFC2396bis]** since XRI syntax is a superset. XRI productions use green shading and URI  
 1422 productions yellow shading. A valid XRI MUST conform to this ABNF.

1423

```

1424 abs-path      = "/" path-segments
1425
1426 alphanum     = ALPHA / DIGIT
1427
1428 authority    = [ userinfo "@" ] host [ ":" port ]
1429
1430 authority-part = URI-authority / XRI-authority
1431
1432 dec-octet    = DIGIT                               ; 0-9
1433              / %x31-39 DIGIT                       ; 10-99
1434              / "1" 2DIGIT                          ; 100-199
1435              / "2" %x30-34 DIGIT                   ; 200-249
1436              / "25" %x30-35                        ; 250-255
1437
1438 delims       = "<" / ">" / "%" / DQUOTE
1439
1440 domainlabel  = alphanum [ 0*61( alphanum / "-" ) alphanum ]
1441
1442 escaped       = "%" HEXDIG HEXDIG
1443
1444 excluded     = invisible / delims / unwise
1445
1446 fragment    = *( pchar / "/" / "?" )
1447
1448 gcs-char     = "+" / "=" / "@" / "$" / "*"
1449
1450 global-path  = [ "!" ] authority-part [ local-path ]
1451
1452 global-xri   = global-path [ "?" xri-query ] [ "#" xri-fragment ]
1453
1454 h4           = 1*4HEXDIG
1455
1456 hier-part    = net-path / abs-path / rel-path
1457
1458 host         = [ hostname / IPv4address / IPv6reference ]
1459
1460 hostname     = idomainlabel qualified
1461
1462 idomainlabel = 1*ucschar
1463
1464 invisible    = CTL / SP / %x80-FF
1465
1466 IPv4address  = dec-octet "." dec-octet "." dec-octet "." dec-octet
1467
1468 IPv6address  = 6( h4 ":" ) 1s32
1469              / "::" 5( h4 ":" ) 1s32
1470              / [ h4 ":" 4( h4 ":" ) 1s32
1471              / [ *1( h4 ":" ) h4 ] "::" 3( h4 ":" ) 1s32
1472              / [ *2( h4 ":" ) h4 ] "::" 2( h4 ":" ) 1s32
1473              / [ *3( h4 ":" ) h4 ] "::" h4 ":" 1s32
1474              / [ *4( h4 ":" ) h4 ] "::" 1s32
1475              / [ *5( h4 ":" ) h4 ] "::" h4

```



```

1476           / [ *6( h4 ":" ) h4 ] "::-"
1477
1478 IPv6reference = "[" IPv6address "]"
1479
1480 local-path   = "/" relative-path
1481
1482 ls32        = ( h4 ":" h4 ) / IPv4address
1483             ; least-significant 32 bits of address
1484
1485 mark        = "-" / "_" / "." / "!" / "~" / "*" / "'" / "(" / ")"
1486
1487 net-path    = "//" authority [ abs-path ]
1488
1489 path-segments = segment *( "/" segment )
1490
1491 pchar       = unreserved / escaped / ";" /
1492             ":" / "@" / "&" / "=" / "+" / "$" / ",",
1493
1494 port        = *DIGIT
1495
1496 qualified   = *( "." idomainlabel ) [ "." ]
1497
1498 query       = *( pchar / "/" / "?" )
1499
1500 relative-path = *( [ "." ] "/" ) xri-segments
1501
1502 rel-path    = path-segments
1503
1504 reserved    = "/" / "?" / "#" / "[" / "]" / ";" /
1505             ":" / "@" / "&" / "=" / "+" / "$" / ",",
1506
1507 scheme      = ALPHA *( ALPHA / DIGIT / "+" / "-" / "." )
1508
1509 segment     = *pchar
1510
1511 sub-segment = *xri-pchar / xref
1512
1513 uchar       = %xA0-D7FF / %xF900-FDCF / %xFDF0-FFEF /
1514             %x10000-1FFFD / %x20000-2FFFD / %x30000-3FFFD /
1515             %x40000-4FFFD / %x50000-5FFFD / %x60000-6FFFD /
1516             %x70000-7FFFD / %x80000-8FFFD / %x90000-9FFFD /
1517             %xA0000-AFFFD / %xB0000-BFFFD / %xC0000-CFFFD /
1518             %xD0000-DFFFD / %xE1000-EFFFD
1519
1520 unreserved  = ALPHA / DIGIT / mark
1521
1522 unwise      = "{" / "}" / "|" / "\" / "^" / "`"
1523
1524 URI         = scheme ":" hier-part [ "?" query ] [ "#" fragment ]
1525
1526 URI-authority = "//" [ userinfo "@" ] host [ ":" port ]
1527
1528 uric        = reserved / unreserved / escaped
1529
1530 userinfo    = *( unreserved / escaped / ";" /
1531             ":" / "&" / "=" / "+" / "$" / ",", )
1532
1533 xref        = "(" ( global-xri / URI ) ")"
1534
1535 xref-authority = xref ( "." sub-segment / ":" sub-segment ) *( "."
1536             sub-segment / ":" sub-segment )
1537
1538 XRI         = "xri:" xri-value

```

```

1539
1540 XRI-authority = ( gcs-char xri-segment ) / xref-segment
1541
1542 xri-characters = xri-reserved / xri-unreserved / escaped
1543
1544 xri-fragment = [ xref ] * ( pchar / "/" / "?" )
1545
1546 xri-mark = "-" / "_" / "~" / "'"
1547
1548 xri-path = global-path / local-path / relative-path
1549
1550 xri-pchar = xri-unreserved / escaped / ";" / "!" / "*"
1551 "@ " / "&" / "=" / "+" / "$" / ","
1552
1553 xri-query = [ xref ] * ( pchar / "/" / "?" )
1554
1555 xri-reserved = "/" / "?" / "#" / "[" / "]" / "(" / ")" / ";" / ":" /
1556 ", " / "." / "&" / "@" / "=" / "+" / "*" / "$" / "!"
1557
1558 xri-segment = ( [ "." ] sub-segment / ":" sub-segment )
1559 *( "." sub-segment / ":" sub-segment )
1560
1561 xri-segments = xri-segment *( "/" xri-segment )
1562
1563 xri-unreserved = ALPHA / DIGIT / ucschar / xri-mark
1564
1565 xri-value = [ xri-path ] [ "?" xri-query ] [ "#" xri-fragment ]
1566

```

---

## Appendix B. Special Identifiers Assigned by the XRI Specification

1567  
1568

1569 As defined in Section 2.1.1.2.1, Global Context Symbols (GCS), the GCS character "\$" is  
1570 reserved for identifiers for which the XRI specification is the authority. The purpose of this special  
1571 set is to define metadata that is specific to identifiers and the act of identification (resolution).  
1572 Establishing these identifiers at the level of the XRI specification enables interoperability of this  
1573 metadata among XRI implementations. Specifically this includes:

- 1574 • *Human-readable metadata* that allows free text comments to be embedded in an XRI.
- 1575 • *Versioning metadata* that identifies the syntax of a version identifier.
- 1576 • *Linguistic metadata* that identifies the language or font of an internationalized identifier.
- 1577 • *Internationalization encoding metadata* that identifies the level of encoding of an XRI. (See  
1578 section [ref I18N section].
- 1579 • *Query metadata* that identifies the syntax of a query string.

1580

1581 [DSR: I ran out of time to complete this section by turning the portion below into a 3-column table:  
1582 Identifier, Identifier Purpose, Comments and Requirements. I also intend to prefix this with a set of  
1583 requirements for the \$ namespace as a whole, include terseness, the use of URI-legal chars, and  
1584 the use of all-lowercase.]

1585

1586	\$!	= non-resolvable free text human comment
1587	\$v	= version (default is standard xri-segment syntax)
1588	\$v.d	= version in XML datetime format
1589	\$l	= language (when necessary for disambiguation of internationalized XRIs)
1590	\$f	= font (when necessary for disambiguation of internationalized XRIs)
1591	\$i	= internationalization encoding (when necessary for equivalence)
1592	\$q	= query (default is standard xri-segment syntax)
1593	\$q.xpath	= query in XPath syntax

1594

1595 Note that like all authority segments, a slash delimits the end of the segment.

1596 [DSR note: should also discuss the use of cross-references using "+" syntax for common names.]

---

1597 **Appendix C. Transforming HTTP URIs to XRI**

1598 [This section should discuss:

1599 a) relationship of HTTP URIs and XRIs (e.g., answer the questions brought up on the list), and

1600 b) specify the non-normative rules required to transform an HTTP URI into a legal XRI.]

---

1601 **Appendix D. Acknowledgments**

1602 The following individuals were members of the committee during the development of this  
1603 specification:

- 1604 • Numerous people

1605 In addition, the following people made contributions to this specification:

- 1606 • Other people

1607

## Appendix E. Revision History

1608

*[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.

1609

---

1610

## Appendix F. Notices

1611 OASIS takes no position regarding the validity or scope of any intellectual property or other rights  
1612 that might be claimed to pertain to the implementation or use of the technology described in this  
1613 document or the extent to which any license under such rights might or might not be available;  
1614 neither does it represent that it has made any effort to identify any such rights. Information on  
1615 OASIS's procedures with respect to rights in OASIS specifications can be found at the OASIS  
1616 website. Copies of claims of rights made available for publication and any assurances of licenses  
1617 to be made available, or the result of an attempt made to obtain a general license or permission  
1618 for the use of such proprietary rights by implementors or users of this specification, can be  
1619 obtained from the OASIS Executive Director.

1620 OASIS invites any interested party to bring to its attention any copyrights, patents or patent  
1621 applications, or other proprietary rights which may cover technology that may be required to  
1622 implement this specification. Please address the information to the OASIS Executive Director.

1623 **Copyright © OASIS Open 2003. All Rights Reserved.**

1624 This document and translations of it may be copied and furnished to others, and derivative works  
1625 that comment on or otherwise explain it or assist in its implementation may be prepared, copied,  
1626 published and distributed, in whole or in part, without restriction of any kind, provided that the  
1627 above copyright notice and this paragraph are included on all such copies and derivative works.  
1628 However, this document itself does not be modified in any way, such as by removing the  
1629 copyright notice or references to OASIS, except as needed for the purpose of developing OASIS  
1630 specifications, in which case the procedures for copyrights defined in the OASIS Intellectual  
1631 Property Rights document must be followed, or as required to translate it into languages other  
1632 than English.

1633 The limited permissions granted above are perpetual and will not be revoked by OASIS or its  
1634 successors or assigns.

1635 This document and the information contained herein is provided on an "AS IS" basis and OASIS  
1636 DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO  
1637 ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE  
1638 ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A  
1639 PARTICULAR PURPOSE.

## Appendix G. Issues

Issue	Section	Status
Need to remove this section.	Appendix G	Open
Need to add link to XRI Primer in Abstract when it exists	<a href="#">Abstract</a>	Open
Need to add link to errata page in the Status section	<a href="#">Status</a>	Open
Make sure internationalization text satisfies the HFI internationalization requirement or note that it's not supported in the current spec.	<a href="#">1.2.3</a>	Addressed
Definition for concrete identifier is unclear. An HTTP URI resolves to an IP address, an IP address resolves to a MAC address, etc. Why aren't they abstract by this definition?	<a href="#">1.3.3</a>	Addressed
Definition of abstract identifier may need to be revised to be consistent with clarified definition of concrete identifier.	<a href="#">1.3.3</a>	Addressed
Definition of non-resolvable identifier raised this question, "In my mind, xri:! <a href="#">@IETF/rfc.2396</a> is non-resolvable not because there's no data and/or metadata about it but because it represents the abstract notion of the RFC rather than a particular digital representation of the text. Does this idea match the definition of non-resolvable identifier?"	<a href="#">1.3.3</a>	Addressed
Need text for Character Encoding and Internationalization	<a href="#">2.3</a>	Addressed
Relative resolution is broken by the ! (non-resolvable) symbol. Need to figure out how broken it is and how to fix.	<a href="#">2.4</a>	Addressed
Need text for Internationalized XRI Equivalence	<a href="#">2.5.3</a>	Addressed
Probably need to rename "THTTP Local Access Binding" from thttp to something XRI specific, since its not really RFC2169 compliant	<a href="#">3.3.5.1</a>	Open
Need text for Privacy Considerations.	<a href="#">4.3</a>	Addressed
References to RFC2277 and Unicode aren't used. If they aren't needed by internationalization text, they should be removed.	<a href="#">5.1</a>	Addressed
Need to discuss the vocabulary of the \$ namespace in appendix B. The list there is just the current candidates. They should be approved or removed from the spec.	Appendix B	Open
Appendix C "Transforming HTTP URIs to XRIs" needs text	Appendix C	Open
Appendix D "Acknowledgements" needs to be filled out with the current membership list.	Appendix D	Open
Resolution section needs thorough review	<a href="#">3</a>	Open
Need to review BNF for completeness and correctness (i.e. need to prove the grammar)	Appendix A	Open



Need a section similar to Appendix B of RFC2396 where we provide tools and guidance for parsing XRIs	None	Open
Security section should comment on lack of secure resolution.	4	Addressed
Hyperlinks in doc aren't all enabled. Need to make a pass through the doc and correct links to references and other sections of this document	All	Open
Need to review use of normative keywords ("MUST", "SHOULD", etc) for consistency and correctness.	All	Open
There's a possible terminology issue with section 4.3 "Privacy Considerations". In Europe, "data protection" is the code-word for "privacy". Since we already have the section title as Security and Data Protection, a separate section on Privacy Considerations appears redundant.	4.3	Addressed
Examples and tables don't have or have lost captions	3	Open
Clear up the intent of having multiple resolution mechanisms	3	Open
Mention the fact that resolution is currently only defined on URI-legal character strings, and confirm that this is a reasonable approach.	3	Open
Conversion of SRV records to Authority Descriptors needs fleshing out	3	Open
Use of IP addresses as Authority identifiers needs fleshing out	3	Open
Step 2 of section 2.2.3.2, combined with the last paragraph of that section implies that font and language tags are irrelevant for establishing equivalence. Are we ok with this?	2.2.3.2	Open

1641