



An Introduction to XRI

Working Draft 04, 14 March 2005

Document identifier :

xri-intro-V2.0-wd-04

Location :

<http://docs.oasis-open.org/xri/xri/V2.0>

Editors:

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

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

Contributors:

Fen Labalme, PlaNetwork <fen@idcommons.org>

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

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

Abstract:

This document is a non-normative introduction to the uses and features of XRIs. It is intended to accompany the XRI 2.0 suite of specifications including *Extensible Resource Identifier (XRI) Syntax 2.0 [XRISyntax]*, *Extensible Resource Identifier (XRI) Resolution 2.0 [XRIResolution]*, and *Extensible Resource Identifier (XRI) Metadata 2.0 [XRIMetadata]*.

Status:

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

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

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 Technical Committee web page (<http://www.oasis-open.org/committees/xri/ipr.php>).

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

Table of Contents

38	1.	Introduction	3
39	1.1	What are XRIs?	3
40	1.2	Why a Uniform Abstract Identification Layer?	4
41	2.	The Types of Problems XRIs are Designed to Solve	5
42	2.1	The Broken Links Problem (Persistent Identification)	5
43	2.2	The Multiple Authority Problem (Federated Identification)	5
44	2.3	The N-Squared Mapping Problem (Shared Identification)	6
45	2.4	The Exploding Addresses Problem (Simplified Identification)	6
46	2.5	The Bootstrap Discovery Problem (Metadata Identification)	6
47	2.6	The Public Identifier Problem (Privacy-Protected Identification)	7
48	2.7	The Future-Proofing Problem (Extensible Identification)	7
49	3.	How XRIs Help Solve These Problems: Example Usage Scenarios	8
50	3.1	Persistent Identification	8
51	3.2	Federated Identification	10
52	3.3	Shared Identification	12
53	3.4	Simplified Identification	14
54	3.5	Metadata Identification	15
55	3.6	Privacy-Protected Identification	16
56	3.7	Extensible Identification	17
57	4.	A Brief Guide to the Normative XRI V2.0 Specifications	19
58	4.1	XRI Syntax	19
59	4.2	XRI Resolution	19
60	4.3	XRI Metadata	19
61	5.	References	21
62		Appendix A. Example XRIs	22
63		Appendix B. Acknowledgments	23
64		Appendix C. Revision History	24
65		Appendix D. Notices	25
66			

67 1. Introduction

68 1.1 What are XRIs?

69 XRIs provide a standard syntax and resolution protocol for abstract identifiers—identifiers that are
70 independent of a specific location, domain, application, or protocol. The XRI specifications are
71 built directly on top of the foundation provided by the URI (Uniform Resource Identifier) and IRI
72 (Internationalized Resource Identifier) specifications from IETF and W3C, as shown in Figure 1.
73



74

75

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

76 URIs introduced a standard means of identifying resources across distributed networks that in
77 only 15 years has become the most successful identifier scheme in history. IRIs subsequently
78 extended the generic URI scheme, which supports only the ASCII character set, to include the full
79 UCS (Unicode Character Set).

80 XRIs take a third step by adding additional syntax and resolution features that enable XRIs to
81 solve problems of abstract identification that are not easily addressed by conventional URI or IRI
82 syntax or resolution. One of these problems—persistence—has been addressed by other URI
83 schemes for abstract identifiers, such as the URN (Uniform Resource Name) scheme [RFC2141].
84 While XRIs fulfill the requirements for URNs as specified in [RFC1737], they go on to address a
85 much wider range of issues in abstract identification. As shown in more detail in section 2:

86

87

- *XRIs provide a uniform syntax and resolution protocol for both persistent and reassignable abstract identifiers.*
- *XRIs provide a uniform syntax for delegating abstract identifiers between authorities at any level or context.*
- *XRIs provide a uniform syntax (called “cross-references”) for sharing abstract identifiers across all contexts.*
- *XRIs provide both a generic and a trusted protocol for resolving a single abstract identifier into any number of concrete identifiers.*
- *XRIs provide a simple, standard means of discovering URIs that may be associated with a resource, including those needed for additional metadata discovery.*
- *XRIs provide a standard means of protecting privacy by revealing the least possible information in an identifier and allowing an authority to control further access.*
- *XRIs provide a standard means of extension without sacrificing interoperability.*

88

89

90

91

92

93

94

95

96

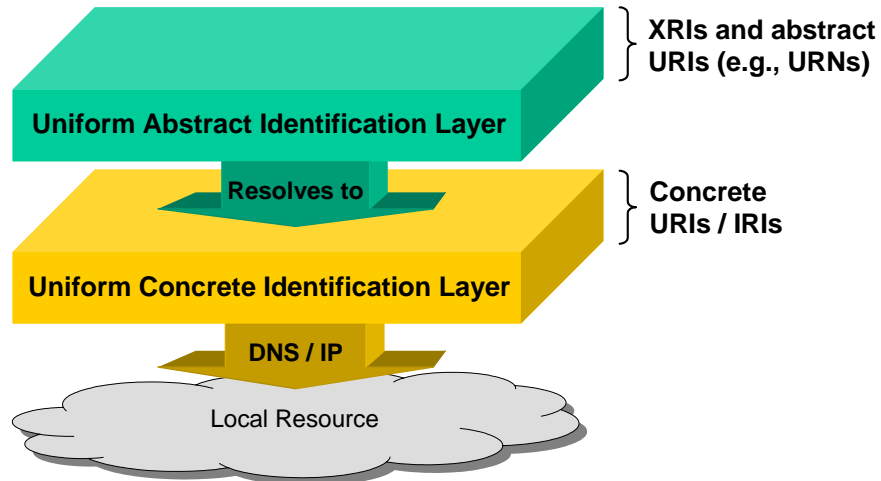
97

98

99

1.2 Why a Uniform Abstract Identification Layer?

100 Not all resources will need the additional layer of indirection that XRI can provide, just as not all
101 resources require URIs or IRIs today. When this additional layer is needed, however, the purpose
102 of XRI is to make it as interoperable across all sites and systems as URIs and IRIs are today. In
103 essence, the goal of XRI architecture is to enable the same *uniform identification layer* for
104 abstract identifiers that URIs and IRIs provide for concrete identifiers today.
105



106

107

Figure 2: The goal of XRI is to provide a uniform abstract identification layer.

108 The key difference between the concrete identification layer and the abstract identification layer is
109 that the latter enables identification of resources independent of a specific context, i.e.:

- 110 • *Independent of a specific network location* where a resource may exist at a particular
111 point in time.
- 112 • *Independent of a specific directory, database, or repository* that might store a resource.
- 113 • *Independent of a specific application* that may be responsible for creating, maintaining, or
114 processing a resource.
- 115 • *Independent of a specific domain, authority, or owner* of a resource at a particular point in
116 time.
- 117 • *Independent of a specific transport or communications protocol* that may be used to
118 interact with a resource.
- 119 • *Independent of a specific semantic label or descriptor* that may be associated with a
120 resource at a particular point in time.

121 While it has long been recognized that persistent identification is one of the key problems an
122 abstract identification layer can solve, the recent growth of distributed XML architecture, Web
123 services, and digital identity infrastructure has brought to light a number of other key problems
124 that a uniform abstract identification layer can help address **[XRIReqs]**.

125 The purpose of this document is to first identify these other types of problems (section 2), then
126 show how XRI can be used to solve them through a series of typical usage scenarios (section
127 3). Finally it provides a brief introduction to the normative XRI 2.0 specifications in section 4.

2. The Types of Problems XRIs are Designed to Solve

2.1 The Broken Links Problem (Persistent Identification)

Probably the best known use case for abstract identifiers is the need to maintain a persistent reference to a resource when it moves or changes locations on the network. As explained in the U.S. Government report *Persistent Identification: A Key Component of an E-Government Infrastructure* [CENDI]:

The current addressing structure for the Web is based on the Uniform Resource Locator (URL). This technology uses a physical location (IP address/server/path/file name) to identify and locate digital objects. While the URL provides direct, efficient access, URL-only naming fails whenever the resources are moved or reorganized. In addition, the URL may stay the same, but the object addressed by the URL may change significantly or be replaced with completely different content. Thus, while the URL permits interoperability in assigning an initial address for an object, it offers no assurance that this address will follow the object as it moves among locations. The lack of persistence or “linkrot” leads to 404 errors (file not found), inhibiting access to digital objects and causing problems when archiving material for long-term preservation and permanent access.

This is the fundamental rationale for URNs (Uniform Resource Names), the URI scheme for persistent, location-independent resource identifiers [RFC2141]. From this standpoint, the most significant difference between URN and XRI architecture is that URNs deal with only with persistent identifiers, whereas XRIs deal with all types of abstract identifiers—both persistent and reassignable.

XRIs provide a uniform syntax and resolution protocol for both persistent and reassignable abstract identifiers.

2.2 The Multiple Authority Problem (Federated Identification)

URIs, particularly HTTP URIs (commonly called URLs), achieved their success largely because they solved the problem of how to locate and retrieve a resource anywhere across the global Internet—no matter what domain or authority hosted it. They did this by building on top of two federated addressing systems: DNS naming and IP addressing, both of which allow multiple independent authorities to cooperate in resolving different segments of an identifier.

However with DNS and IP, the “edge” of the federation network is an IP-addressable resource (a server, a device, a network application). There is no standard syntax for continuing delegation between resources beyond this endpoint. So, for example, if a corporate directory system wanted to delegate authority for home contact data to employees, or a government agency wanted to delegate responsibility for portions of a spreadsheet it to different departments (who in turn wanted to delegate parts of it to different employees), XRIs provide the syntax necessary to express this in a uniform manner.

XRIs provide a uniform syntax for delegating abstract identifiers between authorities at any level or context.

168 **2.3 The N-Squared Mapping Problem (Shared Identification)**

169 One of the most pervasive problems in large-scale systems integration is the cross-context
170 mapping of resources and data. As long as a resource is only identified in a local context, then
171 each additional system with which it must be shared requires an additional mapping between the
172 two contexts. This pairwise mapping problem grows exponentially with the number of resources
173 and the number of systems involved.

174 While security and privacy requirements may require pairwise mappings to be used in certain
175 situations, in most other cases establishing a shared abstract identifier collapses the n^2 mapping
176 problem into a much easier $n-1$ mapping problem. This is particularly true of resources whose
177 primary purpose is to establish shared semantics across systems, e.g., data interchange
178 schemas, dictionaries, taxonomies, ontologies, etc.

179 *XRIs provide a uniform syntax (called “cross-references”) for sharing abstract identifiers*
180 *across all contexts.*

181 **2.4 The Exploding Addresses Problem (Simplified** 182 **Identification)**

183 The constant progression of Internet, wireless, and other communications technologies has
184 begun to produce an explosion of addresses. This problem applies both to people—the exploding
185 business card problem—and increasingly to other resources (smart phones, networked
186 applications) that can also be addressable via multiple networks and interfaces.

187 Jamie Lewis, CEO of the Burton Group, summarized this problem succinctly in a recent press
188 article **[Lewis]**:

189 *Today, we use a wide variety of different mechanisms for identification, including e-mail*
190 *addresses, IP addresses, phone numbers, and object identifiers. But most of these are*
191 *specific to one means of interaction. None of them is persistent across the many different*
192 *ways that people, applications, and devices can communicate, and so they don’t function*
193 *well as identifiers in the long run.*

194 A uniform abstract identification layer can solve this problem both for the addressees (who will not
195 longer need to constantly change their published addresses/interfaces), as well as for the
196 addressors, who will no longer need to constantly update their local references.

197 *XRI provide both a standard and a trusted protocol for resolving a single abstract*
198 *identifier into any number of concrete identifiers.*

199 **2.5 The Bootstrap Discovery Problem (Metadata Identification)**

200 Many URI schemes are specific to a particular means of interaction with a resource (http, https,
201 ftp, mailto, telnet, etc.) Binding the identifier for a resource to the protocol used to access that
202 resource makes it easy to interact with a resource via that particular protocol, but it does not
203 facilitate discovery of other means of interaction that might be available. Nor do most of these
204 protocols standardize a means for obtaining additional metadata about a resource, e.g., its
205 version, media type, encoding, author, rights management, etc.

206 In fact, with the rapidly growing number of ways to describe a resource, its attributes, and
207 relationships, there may in fact be multiple sources and formats of this metadata, each with their
208 own interaction-specific URIs identifying them. But how does one discover these? This is the
209 “bootstrap” discovery problem, and one perfectly suited to a uniform abstract identification layer.

210 *XRIs provide a simple, standard means of discovering URIs that may be associated with*
211 *a resource, including those needed for additional metadata discovery.*

212 **2.6 The Public Identifier Problem (Privacy-Protected**
213 **Identification)**

214 Another ramification of identifiers that are specific to a particular means of interaction is the
215 privacy issues they create, particularly those that need to be used publicly. For example, the most
216 convenient way for users to register at many websites today is using an email address. However
217 many Internet users are reluctant to share their email address for fear of spam.

218 As public addresses, abstract identifiers have the virtue of not needing to reveal a specific means
219 of interaction, yet serving being able to serve as a key by which access requests can be made.
220 By applying policy to these access requests (for example, requiring authentication), XRI
221 authorities can increase their privacy and obtain greater control over the sharing of contact and
222 other sensitive information.

223 *XRIs provide a standard means of protecting privacy by revealing the least possible*
224 *information in an identifier and allowing an authority to control further access.*

225 **2.7 The Future-Proofing Problem (Extensible Identification)**

226 Many pre-Web identifier schemes have been adapted to the Web through the development of
227 special URI schemes (a good example is the “mailto:” scheme for Internet email addresses).
228 Countless other resources in legacy systems have also been made URI-addressable through
229 authority-specific mapping algorithms.

230 However the introduction of new URI schemes, resolution protocols, and mapping algorithms can
231 require expensive changes to existing infrastructure and can increase, rather than decrease,
232 integration complexity. XRIs help ameliorate this problem both by providing an easy way to “bind”
233 new URI schemes to existing infrastructure. They also provide the alternative simply defining new
234 XRI namespaces that can be cross-referenced by all other XRI namespaces. Either way,
235 providing this “extensibility by design” for identifiers was one of the key motivations for XRI just as
236 it was for XML.

237 *XRIs provide a standard means of extension without sacrificing interoperability.*

238

3. How XRI's Help Solve These Problems: Example Usage Scenarios

239

240

This section will use a set of scenarios to show how XRI architecture can solve the types of problems described in section 2. For continuity, the overall context will be a set of issues a government might face in sharing resources via the Internet. XRI's can be used to identify any type of resource – a person, a network device, a Web page, a web service, an XML namespace, a particular instance of a physical book, etc. This set of examples focuses primarily on documents available via the Web, but the principles apply to any type of XRI-identified resource.

241

242

243

244

245

246

3.1 Persistent Identification

247

XRI's provide a uniform syntax and resolution protocol for both persistent and reassignable abstract identifiers.

248

249

As highlighted in the previous [CENDI] quote, one of the most frustrating aspects of the Internet today is the lack of persistent identification for Web-accessible resources. In most cases the root cause is encoding a resource's location and the access protocol in its URI or IRI. Such an identifier is valid only as long as the resource doesn't change locations or retrieval mechanisms. In practical terms, this means when a web site is reorganized, links break.

250

251

252

253

254

As an example, say a department of a government agency publishes a document named `govdoc.pdf` via the following URI:

255

256

```
ftp://department.agency.example.org/docs/govdoc.pdf
```

257

The document is easily located via DNS and retrieved via FTP. However, when the resource's access protocol changes from FTP to HTTP, the resource's identifier changes to:

258

259

```
http://department.agency.example.org/docs/govdoc.pdf
```

260

Typically, external links to the former address would now be broken. Notice that the resource itself didn't change – `govdoc.pdf` is still exactly the same – but its identifier is different because the protocol used to retrieve the resource changed.

261

262

263

Next, imagine the department undergoes a reorganization. Its current domain name, `department.agency.example.org`, is changed to `newdept.agency.example.org`. Now the resource's identifier changes to:

264

265

266

```
http://newdept.agency.example.org/docs/govdoc.pdf
```

267

Finally, the web site hosting `govdoc.pdf` is reorganized, and `docs` is renamed `documents`. The document's identifier changes yet again, this time to:

268

269

```
http://newdept.agency.example.org/documents/govdoc.pdf
```

270

In each case, the resource stayed the same, but its identifier changed. Because the identifier was intrinsically associated with the resource's location and retrieval mechanism, external links to `govdoc.pdf` are broken with every change.

271

272

273

The first step in persistently identifying a resource, then, is assigning an identifier that doesn't depend on the resource's network location and access mechanism, since both will inevitably change over time. This type of identifier is called "abstract" because it doesn't resolve directly to the resource via a specific protocol, but instead resolves into other concrete identifier(s) and

274

275

276

277 access protocol(s) (or metadata that contains this information.) XRIs fulfill this requirement
278 because they don't encode concrete location or protocol information in the identifier.

279 If the government agency department in the previous example used XRIs, in the first step it might
280 have assigned the resource `govdoc.pdf` an XRI such as:

281 `xri://@example.org*agency*department/docs/govdoc.pdf`

282 Note that stars (“*”) are used instead of dots (“.”) as delegation characters in XRI syntax, so the
283 dot that appears in `@example.org` is simply treated as another data character. Also note the
284 leading character “@” in the XRI authority segment—this is called a *global context symbol*. It is
285 used to indicate the abstract global context of an authority—in this case, an organization.

286 XRI resolution would then resolve this abstract identifier into the concrete URI
287 `ftp://department.agency.example.org/docs/govdoc.pdf`. When the access protocol
288 changed from FTP to HTTP, the XRI would not break. Only the resolution result would change—it
289 would return `http://department.agency.example.org/docs/govdoc.pdf` instead of
290 `ftp://department.agency.example.org/docs/govdoc.pdf`.

291 However we still have a problem when the name of the government department is changed from
292 `department` to `newdept`. In this respect the XRI above has the same problem as the original
293 domain name-based URI—it was constructed from semantic names that are subject to change.
294 We *could* try to solve this by creating a new XRI that reflects the department's new name, i.e.:

295 `xri://@example.org*agency*newdept/docs/govdoc.pdf`

296 We could then use XRI resolution to map this name to the previous XRI—
297 `xri://@example.org*agency*department/docs/govdoc.pdf`. This XRI would finally
298 resolve to the concrete URI
299 `http://department.agency.example.org/docs/govdoc.pdf`. However such an
300 approach would eventually lead to a spaghetti code of new-to-old XRI mappings. It also has the
301 drawback of preventing reassignment of the identifier “department” for another purpose.

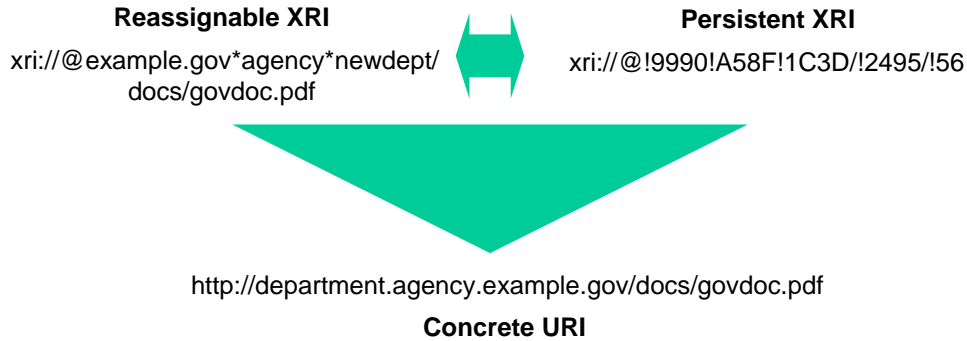
302 A much better solution would be to assign the resource “govdoc.pdf” an identifier that never
303 needs to change or be reassigned. This can be accomplished using a fully persistent XRI such as
304 the following:

305 `xri://@!9990!A58F!1C3D!/!2495`

306 Notice the bang (“!”) symbols that appear throughout this XRI. Bangs are the delimiter used in
307 XRI syntax to indicate that the identifier that follows (called a *sub-segment*) is intended to be
308 persistent (i.e., it fills the role of a URN as defined in **[RFC1737]**).¹ Compare this to the star (“*”)
309 delimiter in the previous example, which indicates that the sub-segment that follows is intended to
310 be *reassignable* (i.e., it fills the same role as a domain name or filename).

311 Once “govdoc.pdf” has been assigned this persistent XRI, the department now has two options
312 for referencing this resource abstractly as shown in Figure 3.

¹ As with URNs, the issue of whether a persistent sub-segment is in fact permanent (never reassigned) is a matter of operational policy for the assigning authority. XRIs can't help with the operational issue, but XRI syntax allows the authority to express its intent.



313

314

Figure 3: The “resolution triangle” of reassignable XRIs, persistent XRIs, and concrete URIs.

315 The two XRIs at the top of the triangle are called *XRI synonyms* because they both resolve to the
 316 same resource. An XRI resolution request for either XRI (see section 3.4) can return the other
 317 XRI synonyms for referencing the resource, as well as the concrete URIs for accessing the
 318 resource, all in one step. So an XRI-consuming application can use human-friendly reassignable
 319 XRIs to transparently discover and “memorize” the persistent XRI synonyms for a resource
 320 without any extra overhead, leading to steadily increasing efficiency and resilience in distributed
 321 systems rather than today’s slowing decaying “link rot”.

322 XRI architecture recognizes that not all identifier authorities may be capable of providing the
 323 infrastructure or management necessary to assign persistent identifiers. It is possible for the
 324 government agency department to have a persistent XRI, for example, but for its internal
 325 document repository not to have this capability. XRI syntax supports this use case by allowing
 326 XRIs to contain a mixture of persistent and reassignable sub-segments, as shown below.²

327

```
xri://@!9990!A58F!1C3D/docs/govdoc.pdf
```

328 This flexibility allows XRIs to be optimized for the context in which they appear: descriptive and
 329 memorable for the side of a bus, for example, or dense and opaque for persistence. XRI
 330 resolution then allows identifier authorities to map reassignable XRIs to persistent XRIs or to
 331 concrete URIs as needed to support both human readability and persistent linking.

322 3.2 Federated Identification

333

XRIs provide a uniform syntax for delegating identifiers between authorities at any level or context.

334

335 As discussed in section 2.2, federated naming and addressing systems are fundamental to
 336 Internet and Web architecture. By allowing multiple independent authorities to cooperate in
 337 naming and resolving different segments of an identifier, they provide both scalability and
 338 distributed management.

339 The best known example of a federated naming system is DNS. Let’s look again at the original
 340 URI for govdoc.pdf:

341

```
ftp://department.agency.example.org/docs/govdoc.pdf
```

² To our knowledge, this feature is unique to XRIs. URNs [RFC2141] and other location-independent identifier systems support only persistent identifiers.

342 The first portion of this URI between the “//” and the first “/” is called the authority segment. This
343 URI uses the DNS name `department.agency.example.org` to specify the authority. DNS
344 resolves each dot-delimited sub-segment right-to-left, i.e., the `gov` nameserver is queried for
345 `example`, the `example` nameserver is queried for `agency`, and the `agency` nameserver is
346 queried for `department`.

347 XRIs work the same way, except federation works from left-to-right (the same direction as the rest
348 of the path). For example, in our original XRI:

```
349 xri://@example.org*agency*department/docs/govdoc.pdf
```

350 The authority segment is `@example.org*agency*department`. To resolve this using XRI
351 resolution, the `@` authority is queried for `*example.org`, the `*example.org` authority is queried
352 for `*agency`, and the `*agency` authority is queried for `*department`.³

353 Notice that each of the sub-segments includes the preceding star (“*”).⁴ This allows an XRI
354 authority to distinguish between reassignable and persistent sub-segments. For example, say the
355 government agency department also assigned `govdoc.pdf` the following persistent XRI:

```
356 xri://@!9990!A58F!1C3D/!2495
```

357 The authority segment here is `@!9990!A58F!1C3D`. To resolve this using XRI resolution, the `@`
358 authority is queried for `!9990`, the `!9990` authority is queried for `!A58F`, and the `!A58F` authority
359 is queried for `!1C3D`.

360 However from the standpoint of federation the most significant difference between DNS and XRI
361 infrastructure is not just that XRIs support both reassignable and persistent identifiers or that
362 delegation is processed from left to right, but that XRIs provide consistent federation syntax
363 throughout the identifier, not just in the XRI Authority portion. Consider the following XRI:

```
364 xri://@example.org*agency*department/offices*east*library/govdoc.pdf
```

365 In this case, delegation continues in the path portion of the XRI. Once we have reached the public
366 XRI authority `@example.org*agency*department`, it can switch to internal delegation (using
367 local XRI resolution or another resolution protocol). The `/offices` authority can delegate to a
368 particular office, `*east`, which in turn can delegate to its internal library, `*library`. Each of
369 these authorities can also maintain corresponding persistent XRIs to support persistent linking if
370 required.

371 A publication could continue delegation internally. For example, `govdoc.pdf` could itself be a
372 compilation of subdocuments that are delegated by the `govdoc.pdf` author. This ability to
373 delegate and federate both reassignable and persistent identifiers at any level is another unique
374 feature of XRIs that has many applications in referencing, linking, and sharing resources across
375 distributed, multi-authority systems.

³ It should also be noted that XRI resolution supports standard URI or IRI authorities. In other words, the XRI `xri://department.agency.example.org/docs/govdoc.pdf` can be resolved using DNS and HTTP content negotiation to an XRI Descriptor. See section 2.3 of **[XRIResolution]**.

⁴ Note that no star appears between “@” and “example.org” in the XRI itself. To make XRI syntax as simple and human-friendly as possible, a star is optional when a reassignable sub-segment follows a global context symbol or a forward slash. However the star is always made explicit when performing XRI resolution. By contrast a bang is never optional—it is always required as the prefix for a persistent XRI sub-segment.

376 3.3 Shared Identification

377 *XRIs provide a uniform syntax (called “cross-references”) for sharing identifiers across all*
378 *contexts.*

379 Imagine that `govdoc.pdf` goes into wide circulation and the department responsible for it
380 obtains an International Standard Book Number, or ISBN, of 0-395-36341-1. This can be
381 represented as a URN as:

```
382 urn:ISBN:0-395-36341-1
```

383 This URN represents a logical resource, not a physical resource. In other words, it identifies the
384 registered identity of the book, not a particular physical copy of the book. If the URN were
385 resolvable, it might reference some information about the book maintained by an ISBN registry,
386 but we wouldn't expect it to resolve to a digital version of the book itself.

387 On the other hand, a reference to a physical resource *can* be constructed by using this ISBN in a
388 specific context. For example, we might refer to, “The National Library's copy of the book with the
389 ISBN 0-395-36341-1.” Notice that the ISBN number *in the context* of the National Library
390 represents a different resource than the ISBN number by itself – a particular copy of the book as
391 opposed to the notion of the book. XRIs provide an elegant way to express the concept of
392 context-based identification. Let's say the National Library has an XRI of:

```
393 xri://@example.org*national.library
```

394 We can now create an XRI that expresses the concept of “The National Library's copy of the book
395 with the ISBN 0-395-36341-1”:

```
396 xri://@example.org*national.library/!(urn:ISBN:0-395-36341-1)
```

397 We can also create both reassignable and persistent way to express the concept of, “The
398 government agency department's representation of the book with the ISBN 0-395-36341-1”:

```
399 xri://@example.org*agency*department/!(urn:ISBN:0-395-36341-1)  
400 xri://@!9990!A58F!1C3D/!(urn:ISBN:0-395-36341-1)
```

401 The same ISBN number in the context of a book store would represent yet another resource:

```
402 xri://@example.bookstore/!(urn:ISBN:0-395-36341-1)
```

403 XRIs allowed us to reuse the identifier `urn:ISBN:0-395-36341-1` consistently across each of
404 these contexts by expressing it as a cross-reference, syntactically distinguished by parentheses.
405 In essence, all of the authorities represented above are able to share the same identifier for the
406 same logical resource (the publication we know as `govdoc.pdf`) in a manner that is
407 understandable to both people and machines, yet allows of them to designate the “meaning” (i.e.,
408 resolution) in their own context.

409 The ability to share identifiers across contexts using compound identifiers is very powerful. It
410 allows us to express ideas like, “The agency's employee with the email address
411 ‘bob@example.com’”:

```
412 xri://@!9990!A58F!1C3D/employees*(mailto:bob@agency.example.org)
```

413 Of course XRIs themselves can also be used in the context of other XRIs. For example, say Bob
414 had registered the following personal XRI under the XRI global context symbol, “=” (used to
415 represent individual persons as authorities.)

416 `xri://=example.bob`

417 Now Bob has the option to use this personal identifier in the various contexts in which he might
418 be known or have a relationship. For example, if Bob was an employee of the government
419 agency, both of the following XRIs could express this relationship:

420 `xri://@example.org*agency*department/employees*(=example.bob)`
421 `xri://@!9990!A58F!1C3D/!24*(=example.bob)`

422 Another use of cross-references for shared identification is identifier metadata, such as date/time
423 stamps and version identifiers. The global context symbol "\$" is reserved for XRIs specified by
424 standards bodies for this purpose. The first two such specifications are **[XRIResolution]**, which
425 establishes the top-level namespace `$res` for identifiers needed in XRI resolution, and
426 **[XRIMetadata]**, which establishes \$ namespaces for XRI language tags, date/time stamps,
427 versioning, and annotations. Following are several examples from **[XRIMetadata]**; see the
428 specification for more.

429 `xri://@France*(($l/fr)/pays)*place`
430 `xri://@example/resource*($d/1994-11-05T08:15:30-05:00Z)`
431 `xri://@example/resource*($v/3.04.2)`
432 `xri://!!1000*($-Example%20Corp.)!5678*($-West%20Coast)/resource`

433 Cross-references are also the key to XRI extensibility, further discussed in section 3.7.

434 3.4 Simplified Identification

435 *XRI provide both a standard and a trusted protocol for resolving a single abstract*
436 *identifier into any number of concrete identifiers.*

437 XRIs permit the abstract identification of a resource independent of location and interaction
438 protocol. That means a single XRI can be used to represent the same resource stored in multiple
439 locations and accessed through multiple protocols. Discovering those concrete endpoints and
440 protocols is the job of XRI resolution, which leverages both the broadly deployed HTTP and
441 HTTPS infrastructure and the flexibility of XML.

442 XRI resolution is a two phase process. The first phase, *authority resolution*, resolves to the XRI
443 authority responsible for the resource. The second phase, *local access*, uses URIs and metadata
444 from the authority to interact with the identified resource. During authority resolution each sub-
445 segment in the authority segment is resolved to an XML document that lists concrete locators
446 (such as HTTP URIs) for the named resource.

447 Let's look again at

448 `xri://@example.org*agency*department/docs/govdoc.pdf`

449 The first phase of resolution operates on the XRI Authority portion of the XRI, in this case,
450 `@example.org*agency*department`. The mechanics of XRI resolution are covered in
451 **[XRIResolution]**, but the final result is an XML document called an XRI Descriptor (often
452 abbreviated as XRID) that describes the XRI Authority. The XRID for
453 `@example.org*agency*department` might look something like this:

```
454 <xrid:XRIDDescriptor xmlns="...">  
455   ...  
456   <xrid:Service>  
457     <xrid:Type>  
458       xri://$res*local.access/X2R  
459     </xrid:Type>  
460     <xrid:URI>  
461       http://department.agency.example.org/  
462     </xrid:URI>  
463     <xrid:URI>  
464       https://department.agency.example.org/  
465     </xrid:URI>  
466   </xrid:Service>  
467   ...  
468 </xrid:XRIDDescriptor>
```

469 The `xrid:Service` element describes a service offering of type
470 `xri://$res*local.access/X2R`. This is a very simple service defined by the XRI resolution
471 protocol that allows a resource to be accessed using the semantics defined by HTTP. In this
472 case, `govdoc.pdf` may be retrieved by choosing one of the listed URIs, appending the path
473 portion of the XRI and performing an HTTP GET on the result. The path portion of original XRI
474 was:

475 `/docs/govdoc.pdf`

476 Appending this to one of the listed URIs in the `xrid:Service` element above would produce:

477 `http://department.agency.example.com/docs/govdoc.pdf`

478 This is one URI that can be used to perform HTTP access on `govdoc.pdf`. Since the XRID also
479 indicates that HTTPS access is supported, a client application can calculate a second URI that
480 can be used for secure access:

481 `https://department.agency.example.com/docs/govdoc.pdf`

482 When the location(s) or retrieval protocol(s) supported by this government agency department
483 change, only its XRI Descriptor needs to change. In addition, the XRID can list multiple locations,
484 access protocols, media types etc. supported by this XRI authority.

485 This XML-mediated abstraction layer can significantly simplify the process of identifying and
486 interacting with any resource maintained by an XRI authority, whether this authority is an
487 organization, a person, a community, etc.

488 3.5 Metadata Identification

489 *XRIs provide a simple, standard means of discovering URIs that may be associated with*
490 *a resource, including those needed for additional metadata discovery.*

491 XRI Descriptors are not limited simply to providing URIs. They can also publish additional
492 metadata about the types of resources available from an authority. To facilitate metadata
493 discovery, XRIDs natively support a `MediaType` element as shown in the following example:

```
494 <xrid:XRIDDescriptor xmlns="...">  
495   ...  
496   <xrid:Service>  
497     <xrid:Type>  
498       xri://$res*local.access/X2R  
499     </xrid:Type>  
500     <xrid:URI>  
501       http://department.agency.example.org/  
502     </xrid:URI>  
503     <xrid:URI>  
504       https://department.agency.example.org/  
505     </xrid:URI>  
506     <MediaType>  
507       application/pdf  
508     </MediaType>  
509     <MediaType>  
510       application/rdf+xml  
511     </MediaType>  
512   </xrid:Service>  
513   ...  
514 </xrid:XRIDDescriptor>
```

515 This element would, for example, allow a client application to discover that an RDF description of
516 `govdoc.pdf` may be available.

517 XRIDs are also not limited to descriptions of local access services—the XML schema is
518 extensible and can include elements from other namespaces. For example, in XRI trusted
519 resolution, where the goal of resolution is to obtain a response with some assurance of
520 authenticity, XRIDs contain SAML assertions issued by the XRI authority producing the response.
521 See section 3.3.1 of **[XRIResolution]** for more details.

522 Following this extension model, XRIDs can also include any other metadata the authority chooses
523 to publish. For example, the XRI Descriptor above might have included the following:

524

525

```

526 <xrid:XRIDescriptor xmlns="...">
527   ...
528   <other:AuthorityDetails>
529     <other:TechnicalContact>
530       mailto:bob@agency.example.org
531     </other:TechnicalContact>
532   ...
533   </other:AuthorityDetails>
534   ...
535 </xrid:XRIDescriptor>

```

536 The AuthorityDetails element is defined in an external schema identified by the other
537 namespace. Resolvers that understand this element can process it, while those that don't are free
538 to ignore it.

539 It is important to note that generic XRI resolution, like DNS resolution, is public, i.e. it does not
540 require client authentication and it does not return different results for different requestors. When
541 using generic XRI resolution, XRI Descriptors should contain only information that is appropriate
542 to share with the public.

543 3.6 Privacy-Protected Identification

544 *XRIs provide a standard means of protecting privacy by revealing the least possible*
545 *information in an identifier and allowing an authority to control further access.*

546 The proceeding example of including an email address as a public identifier highlights another
547 benefit of abstract identifiers: they can be constructed to not reveal private information while still
548 providing a means to access such information with the proper authorization.

549 For instance, the XRI Descriptor in the previous example might have contained the following XRI
550 rather than an explicit email address:

```

551 <xrid:XRIDescriptor xmlns="...">
552   ...
553   <other:AuthorityDetails>
554     <other:TechnicalContact>
555       xri://@example.org*agency*department*contact/(=example.bob)
556     </other:TechnicalContact>
557   ...
558   </other:AuthorityDetails>
559   ...
560 </xrid:XRIDescriptor>

```

561 This XRI will resolve to an XRID describing how the individual represented by =example.bob
562 might be contacted. For example, it could contain an xrid:Service element that specified an
563 HTTP or HTTPS endpoint for submitting a contact request that can be validated according to a
564 local authentication policy.

565 XRI syntax can also be used to express the complex identifiers that often emerge in federated
566 identity scenarios such as those in SAML, Liberty Alliance, WS-Federation, etc. For example, it's
567 possible to express:

568 The identity A calls Bob.

569 `xri://@a/Bob`

570 The identity in B's namespace that A calls Bob.

571 `Xri://@b/(@a/Bob)`

572 The identity in B's namespace known by the public identifier `mailto:bob@example.com`.

```
573 xri://@b/(mailto:bob@example.com)
```

574 The identity in B's namespace known by the public identifier `=bob`.

```
575 xri://@b/(=bob)
```

576 The identity in B's namespace that A calls 123. In this case, 123 might be a pseudonymous
577 identifier, i.e. an identifier used only when A talks to B about that identity.

```
578 xri://@b/(@a/123)
```

579 The identity in B's namespace that A calls 456 when A talks to C about that identity. This might
580 be appropriate when C is facilitating an introduction between A and B. In that case, 456 might be
581 used only once, during the introduction, and then discarded.

```
582 xri://@b/(@c/(@a/456))
```

583 XRI authority and cross-reference syntax is flexible enough to accommodate all these use
584 cases in a manner that is interoperable across all these protocols and frameworks.

585 3.7 Extensible Identification

586 *XRIs provide a standard means of extension without sacrificing interoperability.*

587 Because there will always be new and better ways to identify resources, the “X” in “XRI” reflects
588 the same design principle as it does in “XML”: extensibility-by-design. This is accomplished using
589 XRI cross-reference syntax in two ways.

590 The first is using cross-references to incorporate identifiers created in new URI schemes. For
591 example, say our fictitious government agency wanted to create a database of all the country's
592 jurisdictions for referencing in government legislation. It might decide the best way to make these
593 jurisdictional identifiers Web-compatible is to create a new URI scheme. This new `juris:`
594 scheme might be partitioned into cities, counties, and provinces as shown below:

```
595 juris:city:example.city  
596 juris:county:example.county  
597 juris:province:example.province
```

598 Despite the fact that this is a brand new URI scheme, perhaps with its own resolution protocol, it
599 is immediately interoperable with XRI syntax because a `juris:` reference can be incorporated
600 into an XRI as a cross-reference. For example, our department could specify one of the provinces
601 over which it has authority as:

```
602 xri://@example.org*agency*department/(juris:province:example.province)
```

603 However, if the government is already using XRIs for abstract identification, it also has the option
604 of simply creating a new XRI namespace. One XRI global context symbol, “+”, is specifically
605 designated for generic identifiers—the online equivalent of generic nouns. Any XRI authority can
606 create identifiers in this space that can be cross-referenced by any other XRI authority.

607 So the government could create a new XRI namespace for jurisdictional references such as the
608 following:

```
609 xri://+juris
```

610 This new XRI namespace immediately enjoys all the other standard benefits of XRI syntax and
611 resolution.⁵ For example, the government could partition this namespace (and even delegate
612 these partitions to separate authorities) the same way it could with a new URI scheme, but
613 without the need to create a new resolution protocol.

```
614 xri://+juris*city/example.city  
615 xri://+juris*county/example.county  
616 xri://+juris*province/example.province
```

617 Now other XRI authorities, such as our example agency department, can reference this
618 namespace using cross-references. For example, the department could specify one of the
619 provinces over which it has authority as follows.

```
620 xri://@example.org*agency*department/(+juris*province/example.province)
```

621 Again, the ability of cross-references to share identifiers across contexts, when combined the
622 ability for global context symbols to establish the abstract global context of an identifier, enables
623 XRIs to provide the same interoperable richness for identifiers that XML provides for data.

⁵ Note that XRIs in the + namespace may only locally resolvable using XRI “dictionary” authorities. Multiple dictionary authorities may then cooperate to establish common synonyms, eventually evolving broadly adopted generic XRIs the same way generic nouns evolve in human language.

624 4. A Brief Guide to the Normative XRI V2.0 625 Specifications

626 4.1 XRI Syntax

627 **[XRISyntax]** serves the same role for XRIs that RFC 3986 **[URI]** serves for URIs and RFC 3987
628 **[IRI]** serves for IRIs. It includes the following major sections:

- 629 • *Syntax Components* defines the normative ABNF grammar.
- 630 • *Transformations* defines the rules for transforming a valid XRI into a valid IRI and a valid
631 URI and back again (using these transformations, an XRI may be used any place an IRI
632 or URI may be used).
- 633 • *Relative References* covers the rules for dealing with relative XRIs.
- 634 • *Normalization and Comparison* covers XRI equivalence.
- 635 • *Security and Data Protection* discusses the ways in which compound identifiers may be
636 subject to attack or information leakage.

637 It also includes an appendix on transforming HTTP URIs into XRIs and a Glossary that serves all
638 three normative XRI V2.0 specifications.

639 4.2 XRI Resolution

640 **[XRIResolution]** defines an optional HTTP/HTTPS-based resolution protocol for XRIs. It includes
641 the following major sections:

- 642 • *Generic Resolution* defines the basic two-phase resolution protocol, including the
643 structure of XRI Descriptors (XRIDs), the construction of HTTP(S) URIs for XRI
644 authorities, the use of XRI synonyms for redirects, IRI authorities, local access service,
645 HTTP headers, and HTTP caching.
- 646 • *Trusted Resolution* specifies how generic resolution is extended to create a chain of trust
647 between the participating authorities using SAML assertions. It defines the additional
648 XRID elements needed, the client and server behavior, and the special requirements of
649 authorities offering trusted resolution.
- 650 • *Extensibility and Versioning* explains the extensibility points built into the protocol and
651 how future versioning of the specification and XML schema will be handled.
- 652 • *Security and Data Protection* covers special considerations for the use of XRI resolution.

653 Appendix A includes the normative XML schema for XRIDs, and appendix B includes the non-
654 normative RelaxNG compact syntax schema.

655 4.3 XRI Metadata

656 **[XRIMetadata]** defines four namespaces under the XRI global context symbol “\$” for common
657 types of identifier metadata:

- 658 • *Language Metadata* (\$l) permits expressing the human language in which an XRI, or an
659 XRI sub-segment, is intended to be understood, interpreted, or pronounced.
- 660 • *Date/time Metadata* (\$d) provides a standard means of representing the date and time
661 that an XRI was assigned to a resource.
- 662 • *Version Metadata* (\$v) defines a standard means of representing the version of an XRI or
663 XRI sub-segment.

664 • *Annotation Metadata* (\$) enables XRI producers to provide human-readable annotations
665 in an XRI, or an XRI sub-segment, without affecting resolution or comparison.

666 The specification also covers normalization and comparison rules for XRI metadata and explains
667 how these and other forms of XRI metadata can be extended by any XRI authority using cross-
668 references.

669

5. References

670

671

672 [CENDI]

Persistent Identification: A Key Component of an E-Government Infrastructure. http://cendi.dtic.mil/publications/04-2persist_id.pdf, March, 2004.

673

674

675 [IRI]

M. Duerst, M. Suignard, *Internationalized Resource Identifiers (IRIs)*, <http://www.ietf.org/rfc/rfc3987.txt>, RFC 3987, January 2005.

676

677 [Lewis]

BetaNews, <http://www.betanews.com/article/1075980052>, February 5, 2004

678

679 [XRI Syntax]

D. Reed, D. McAlpin, *Extensible Resource Identifier (XRI) Syntax V2.0*, <http://docs.oasis-open.org/xri/xri/V2.0/xri-syntax-V2.0-cd-01.pdf>, March 2005.

680

681

682 [XRI Resolution]

G. Wachob, *Extensible Resource Identifier (XRI) Resolution V2.0*, <http://docs.oasis-open.org/xri/xri/V2.0/xri-resolution-V2.0-cd-01.pdf>, March 2005.

683

684

685 [XRI Metadata]

D. Reed, *Extensible Resource Identifier (XRI) Metadata V2.0*, <http://docs.oasis-open.org/xri/xri/V2.0/xri-metadata-V2.0-cd-01.pdf>, March 2005.

686

687

688 [RFC1737]

K. Sollins, L. Masinter, *Functional Requirements for Uniform Resource Names*, <http://www.ietf.org/rfc/rfc1737.txt>, RFC 1737, December 1994.

689

690 [RFC2119]

S. Bradner, *Key words for use in RFCs to Indicate Requirement Levels*, <http://www.ietf.org/rfc/rfc2119.txt>, RFC 2119, March 1997.

691

692 [RFC2141]

R. Moats, *URN Syntax*, <http://www.ietf.org/rfc/rfc2141.txt>, IETF RFC 2141, May 1997.

693

694 [RFC2234]

D. H. Crocker and P. Overell, *Augmented BNF for Syntax Specifications: ABNF*, <http://www.ietf.org/rfc/rfc2234.txt>, RFC 2234, November 1997.

695

696 [RFC3066]

H. Alvestrand, *Tags for the Identification of Languages*, <http://www.ietf.org/rfc/rfc3066.txt>, RFC 3066, January, 2001.

697

698 [SAML]

S. Cantor, J. Kemp, R. Philpott, E. Maler, *Assertions and Protocols for the OASIS Security Assertion Markup Language (SAML) V2.0*, <http://www.oasis-open.org/committees/security>, March 2005.

699

700

701 [URI]

T. Berners-Lee, R. Fielding, L. Masinter, *Uniform Resource Identifier (URI): Generic Syntax*, <http://www.ietf.org/rfc/rfc3986.txt>, STD 66, RFC 3986, January 2005.

702

703

704 [XMLSchema2]

P. Biron, A. Malhotra, *XML Schema Part 2: Datatypes Second Edition W3C Recommendation*, <http://www.w3.org/TR/xmlschema-2/>, October 2004.

705

706

707 [RFC3066bis]

A. Phillips, M. Davis, *Tags for Identifying Languages*, <http://www.ietf.org/internet-drafts/draft-phillips-langtags-09.txt>, Work-In-Progress, January, 2005.

708

709

710 [XRI Reqs]

G. Wachob, D. Reed, M. Le Maitre, D. McAlpin, D. McPherson, *Extensible Resource Identifier (XRI) Requirements and Glossary v1.0*, <http://www.oasis-open.org/apps/org/workgroup/xri/download.php/2523/xri-requirements-and-glossary-v1.0.doc>, June 2003.

711

712

713

714

715

716

Appendix A. Example XRIs

717 The purpose of this appendix is to provide a consolidated set of XRIs that help illustrate key
718 concepts of XRI syntax. (Note the heavy use of “example” and “!1000” for authority names in
719 order to prevent the use of real authorities.)

720 Fully Persistent XRIs

721 `xri://@!1000!123!0-395-36341-1`

722 `xri://@!1000!1234/!rfc!2141`

723 This XRI includes a persistent sub-segment that is “semantically reflective”, i.e. it has
724 meaning that is apparent to the reader. This is supported in XRI syntax, but should
725 generally be avoided.

726 `xri://!!1000!34F2!A98E!B8FC/!3283`

727 This XRI is rooted on the bang (“!”) global context symbol, the context for persistent
728 identification of physical resources, independent of a person or organization.

729 Fully Reassignable XRIs

730 `xri://@example.gov*agency*department/docs/govdoc.pdf`

731 `xri://@example/john.doe`

732 `xri://=example.john.doe`

733 This XRI is rooted on the equals (“=”) global context symbol, the context for identification
734 of individual persons.

735 `xri://=example.john.doe*home`

736 `xri://=example.john.doe*work`

737 `xri://=example.john.doe*work/(+phone)`

738 This XRI contains a cross-reference rooted on the plus (“+”) global context symbol, the
739 context for generic identifiers. The interpretation of these identifiers is established by
740 convention and/or shared XRI dictionaries.

741 `xri://xri.example.com/john.doe`

742 This XRI is rooted on a DNS name.

743 `xri://(http://example.com)/john.doe`

744 This XRI is rooted on a cross-reference, which allows any local community to establish a
745 root authority without registering with a central authority.

746 XRIs with both persistent and reassignable sub-segments

747 `xri://@example.bookstore/!1E2F!34A2`

748 `xri://@!1000!3432!2397/documents/file.txt`

749 `xri://@!1000!3432!2397/documents/example.xml#34234`

750 XRIs support the same query and fragment syntax as URIs and IRIs

751 `xri://@example.gov*national.library/!(urn:ISBN:0-395-36341-1)`

752 This XRI has a persistent cross-reference (URNs satisfy the requirements for persistent
753 XRI segments).

754

755

Appendix B. Acknowledgments

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

- 758 • Geoffrey Strongin, Advanced Micro Devices
- 759 • Ajay Madhok, AmSoft Systems
- 760 • Jean-Jacques Dubray, Attachmate
- 761 • William Barnhill, Booz Allen and Hamilton
- 762 • Drummond Reed, Cordance Corporation
- 763 • Marc Le Maitre, Cordance Corporation
- 764 • Dave McAlpin, Epok
- 765 • Loren West, Epok
- 766 • Peter Davis, NeuStar
- 767 • Masaki Nishitani, Nomura Research
- 768 • Nat Sakimura, Nomura Research
- 769 • Tetsu Watanabe, Nomura Research
- 770 • Owen Davis, PlaNetwork Inc
- 771 • Victor Grey, PlaNetwork Inc
- 772 • Fen Labalme, PlaNetwork Inc
- 773 • Mike Lindelsee, Visa International
- 774 • Gabriel Wachob, Visa International
- 775 • Dave Wentker, Visa International
- 776 • Bill Washburn, XDI.ORG

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

779 Thomas Bikeev, EAN International; Krishna Sankar, Cisco; Winston Bumpus, Dell; Joseph
780 Moeller, EDS; Steve Green, Epok; Lance Hood, Epok; Adarbad Master, Epok; Davis McPherson,
781 Epok; Chetan Sabnis, Epok; Phillipe LeBlanc, GemPlus; Jim Schreckengast, Gemplus; Xavier
782 Serret, Gemplus; John McGarvey, IBM; Reva Modi, Infosys; Krishnan Rajagopalan, Novell;
783 Tomonori Seki, NRI; James Bryce Clark, OASIS; Marc Stephenson, TSO; Mike Mealling,
784 Verisign; Rajeev Maria, Visa International; Terence Spielman, Visa International; John Veizades,
785 Visa International; Lark Allen, Wave Systems; Michael Willett, Wave Systems; Matthew Dovey;
786 Eamonn Neylon; Mary Nishikawa; Lars Marius Garshol; Norman Paskin; Bernard Vatan.

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

788

Appendix C. Revision History

Rev	Date	By Whom	What
01	2/24/2005	Drummond Reed	Initial document.
02	3/3/2005	Drummond Reed, Dave McAlpin	Revised sections 1 and 2, scenarios developed for section 3
03	3/14/2005	Dave McAlpin, Drummond Reed	Editorial revisions to sections 1 and 2, rough draft of almost all of section 3
04	3/14//2005	Dave McAlpin, Drummond Reed	Complete draft including appendix A

789

790

Appendix D. Notices

791 OASIS takes no position regarding the validity or scope of any intellectual property or other rights
792 that might be claimed to pertain to the implementation or use of the technology described in this
793 document or the extent to which any license under such rights might or might not be available;
794 neither does it represent that it has made any effort to identify any such rights.

795 Information on OASIS's procedures with respect to rights in OASIS specifications can be found at
796 the OASIS website. Copies of claims of rights made available for publication and any assurances
797 of licenses to be made available, or the result of an attempt made to obtain a general license or
798 permission for the use of such proprietary rights by implementors or users of this specification,
799 can be obtained from the OASIS President.

800 OASIS invites any interested party to bring to its attention any copyrights, patents or patent
801 applications, or other proprietary rights which may cover technology that may be required to
802 implement this specification. Please address the information to the OASIS President.

803 **Copyright © OASIS Open 2005. All Rights Reserved.**

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

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

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