



Proposal for Extensibility Features in the SAML Schemas

Proposal Draft ~~021~~, ~~119~~ January~~February~~ 2004

Document identifier:

sstc-maler-schema-extension-~~021~~

Location:

http://www.oasis-open.org/committees/documents.php?wg_abbrev=security

Authors:

Scott Cantor, individual (cantor.2@osu.edu)
Eve Maler, Sun Microsystems (eve.maler@sun.com)

Abstract:

This document proposes changes to the various extensibility mechanisms used in SAML's XML schemas. It is hoped that some of the proposal's content will contribute to a technical paper on the SAML TC's extensibility design decisions and/or to the SAML core specification's explanation of extensibility.

Status:

This is the ~~first~~second draft of this proposal, incorporating some feedback and decisions from the 20 January 2004 Security Services TC telecon and additional ideas from the XML community. Comments should be sent to the authors or to the security-services@lists.oasis-open.org mailing list.

22 **Table of Contents**

23 1 Introduction.....3

24 1.1 SAML Customizations to Date.....3

25 1.2 SAML's Extensibility Requirements.....3

26 2 Extensibility Mechanism Choices.....6

27 2.1 Derivable Types.....6

28 2.2 Model Groups and Attribute Groups.....7

29 2.3 Element Substitution.....8

30 2.4 Content Models Containing "Any" Wildcards.....8

31 2.5 Elements Bound to "Any" Types.....9

32 2.6 Global Elements and Attributes.....10

33 2.7 URI-Based Identifiers and NamespacesSAML Namespaces.....10

34 3 Recommendations for SAML Extensibility.....12

35 3.1 Recommendations Related to Native Types.....12

36 3.2 Recommendations Related to Open Content.....12

37 3.3 Recommendations Related to Global Elements.....13

38 3.4 Recommendations Related to SAML Namespaces.....14

39 3.5 Recommendations Related to Extension Understandability.....14

40 4 References.....15

41

1 Introduction

The SAML TC has long had a goal of letting customizers extend SAML in a controlled way in order to maximize interoperability. SAML provides extensibility in three areas:

- **Extensibility of structure:** This includes ways to modify (add to or subtract from) SAML's native XML content models.
- **Extensibility of content:** This includes ways to customize the format and interpretation of the content of SAML's XML elements and attributes. (SAML artifacts could also be considered to come under this category.)
- **Extensibility of "protocol":** This includes ways to define new flows, called profiles, of SAML assertion creation, usage, and exchange. Sometimes these profiles also involve extended XML structures and content, as described above.

SAML was initially developed at a time when the art of W3C XML Schema (XSD) [Schema1] "schemography" was extremely new. Since SAML V1.0 first became stable, a number of users have gained experience with customizing SAML. This experience, along with lessons learned about the abilities and limitations of XSD and its conforming processors, are contributing to a new understanding of how SAML can best achieve its extensibility goal.

This proposal attempts to harness this experience to recommend improvements to SAML's extensibility in the areas of structure and content. (Note that schema features that do not affect extensibility are not discussed here, though this could easily become the subject of another position paper.)

1.1 SAML Customizations ~~to Date~~

The Liberty Alliance schemas [LibProtSchema] are the best-known and most comprehensive customization of SAML. These schemas were developed somewhat in parallel with the original SAML standard and were extended from stable but incomplete drafts of the original specification. This created some divergence as well as some duplication of work as SAML evolved and addressed limitations that Liberty also addressed in the form of extensions; these conflicts have been addressed in part through subsequent revisions of the Liberty schemas, demonstrating that extensions tend to be "evolving" specifications that must adapt as the underlying foundation changes. Understanding these points of contact and minimizing their impact is the essential goal of extensibility.

~~[@@Briefly discuss the nature of Liberty customizations, and add mention of other known customizations]~~

1.2 SAML's Extensibility Requirements

SAML's general requirements for extensibility can be stated as follows.

Note: The SAML Conformance Program Specification [SAMLConform] §2.4 outlines how extensions and their processors must behave in order to conform with SAML. For example, extensions are never allowed to redefine existing semantics nor alter specified behavior.

1. Allow customizers to modify SAML's XML structures in sufficiently powerful, maintainable, and validatable ways-

SAML is generic technology that is designed to be reused, customized, and profiled for a variety of purposes. In general, any customization is likely to be promoted by its creators as a second-order standard, which suggests that customizers' need to create formal schemas is relatively high. This will be particularly true wherever a customizer wishes to make third-order customizations possible.

More controversially, in the past we have had an implicit goal of forcing customizers to define schema customization layers on top of SAML so that we could ensure their correct usage of the native SAML portions. With experience, it seems that this requirement was too nosy and, at times, too impractical.

2. To the extent possible, gracefully tolerate extensions

This is something of an inverse of requirement #1: SAML processors should be able to parse (and

88 | [validate](#)) a document that has an extension in it even if the schema for the extension is not available.
89 | [This probably can't be achieved in all situations, but we can use it to color our choices of](#)
90 | [mechanisms. Note that this requirement might have consequences for the Conformance document.](#)
91 | [David Orchard's recent article on versioning \[Orchard\] discusses some ways to achieve this goal](#)
92 | [more successfully.](#)

93 | **3. Make the extensibility mechanisms practical**

94 | They should not be so cumbersome or costly that customizers avoid them and use cut-and-paste
95 | instead. Support and interoperability of XSD processors and features are key variables here.

96 |
97 |
98 |
99 |
100 |
101 |
102 |
103 |
104 |
105 |

4. **Allow for unambiguous reuse of foreign taxonomies where SAML is just a carrier for them**

In general, SAML tries to avoid doing much more than codifying current usage unless current technologies are wholly inadequate. Thus, if a system of ~~information codes, keywords, or values (such as for resource actions)~~ is already available, SAML needs to allow for a way to indicate interpretation of ~~this~~ information ~~provided this way~~. An example is SAML's action namespace framework, which allows usage of existing sets of actions, such as HTTP verbs and UNIX file permissions. Another example is SAML's attribute namespace framework, which has been less successful because it apparently allows for ambiguous semantics.

2 Extensibility Mechanism Choices

106

107 The following sections describe the major contenders for extensibility mechanisms, SAML V1.1's usage
108 of each, and brief analyses of their strong and weak points. This is not an exhaustive list of the
109 mechanisms made available by XSD. The analysis focuses on these characteristics:

- 110 • **Tools support:** This addresses #1, the practicality goal.
- 111 • **Interactions with other mechanisms:** This also addresses #1, the practicality goal.
- 112 • **Validation power:** This addresses the goal of allowing customizations to be validated in #2. Some
113 checking can be done at “compile time” (with the schema alone as input), and some is done at “run
114 time” (with the schema and the XML instance as input).
- 115 • **Flexibility:** This addresses the goal of allowing sufficiently powerful customizations in #2.
- 116 • **Reuse:** This addresses the goal of allowing sufficiently powerful customizations in #2, and also #3,
117 which deals with SAML's ability to reuse in the other direction.

2.1 Derivable Types

118

119 These are types, abstract and otherwise, that are designed for derivation. XSD allows a type to serve as
120 a base type (a parent) of a specialized type, à la object-oriented systems, unless the original type is set
121 to “final.” This setting can be controlled globally for a whole schema.

122 For the complex types that get bound to XML element content models, you can do type extension, which
123 allows adding to the end of a content model, and type restriction, which allows subtracting any optional
124 element from a content model. For the simple types that get bound to simple element and attribute string
125 content, only restriction through “facets,” extension by allowing a list of multiple same-typed values, and
126 extension by unioning two types are allowed.

127 XSD allows the redefinition (without renaming) of imported types in an importing schema.

128 All of SAML's defined types are non-final and are explicitly documented as being derivable. In several
129 cases, SAML defines “deep” complex type hierarchies (and matching elements) especially for derivation
130 purposes. In the assertion schema, this includes:

- 131 • **ConditionAbstractType > AudienceRestrictionConditionType**
- 132 • **StatementAbstractType > SubjectStatementAbstractType > (AuthenticationStatementType,**
133 **AttributeStatementType, AuthorizationDecisionType)**
- 134 • **AttributeDesignatorType > AttributeType**

135 In the protocol schema, this includes:

- 136 • **RequestAbstractType > RequestType**
- 137 • **QueryAbstractType > SubjectQueryAbstractType > (AuthenticationQueryType,**
138 **AttributeQueryType, AuthorizationDecisionQueryType)**
- 139 • **ResponseAbstractType > ResponseType**

140 The types with “Abstract” in their name are, in fact, set to be abstract, which means they must be derived
141 from in order to be used at all. The SAML core specification [SAMLCore] §6 mentions that derivations of
142 SAML's types may be used either with the XSD `xsi:type` attribute on native elements or through
143 substitution of a foreign element bound to the new type.

Tools support

144

145 Extension tends to be supported. No problems have been reported with SAML customizations of this sort
146 so far. Support for restriction of complex types may be a concern, though trivial restriction from **anyType**
147 or **anySimpleType** is likely not to pose a problem; XSD processors tend to be sloppy about checking
148 compatibility of the restricted type at compile time. At least one expert is suspicious of the overall
149 interoperability of derivation from complex types [Kawaguchi]. Element substitution of restricted types
150 appears to be a problem for some XSD processors.

151 We don't know of any cases of restriction from the SAML schemas.

152 **Interactions with other mechanisms**

153 Derivation gains significant new abilities in combination with substitution and redefinition, but need not
154 be used with them (and in this way avoids their negatives).

155 Use of the `<xs:any>` wildcard at the end of a sequence causes problems when extending the original
156 sequence because of ambiguities between optional elements and wildcard matching.

157 **Validation power**

158 Derivations of SAML types are validatable in XSD processors, though it should be noted that at least one
159 expert is concerned about run-time misinterpretations [Obasanjo1]. Any other mechanism that allows for
160 or requires derivation will gain its validation power. When abstract types are used (as in some cases in
161 SAML), derivation (and validation of it) is forced.

162 **Flexibility**

163 Based on experience to date, the ability to extend types by adding to the end of the content model is
164 sufficient for most customizations, except when the original model terminates in the `<xs:any>` wildcard.

165 **Reuse**

166 Customizers' reuse of native constructs: Non-final types are all available to be used in the construction of
167 an entirely new schema that does not use the native elements. It is not possible to forbid this.

168 Native reuse of others' constructs: Our schemas of interest (SAML) could reuse the types defined by
169 others' schemas. (It currently does not; the only "foreign" types it are ones built into the XSD
170 specification.)

171 **2.2 Model Groups and Attribute Groups**

172 Model groups are like programming macros that allow for usage of a previously defined XML content
173 model particle or list of attributes. They can be used to build new content models without the usual
174 extension/restriction constraints, and more importantly for the purposes this paper, an importing schema
175 can redefine a model group (without changing its name) to have different content.

176 SAML does not currently use model groups, so schema customizations can't redefine them.

177 **Tools support**

178 We don't know if tools support is an issue here. At a guess, proper support for redefinition is less
179 universal than support for basic usage of model and attribute groups.

180 **Interactions with other mechanisms**

181 Model and attribute groups don't depend on other mechanisms, though they can be used in combination
182 with type derivation.

183 **Validation power**

184 Model groups form content models, against which XML instances are validated by conforming XSD
185 processors. Attribute groups contribute to the list of attributes on an element, against which instances
186 are likewise validated. Any redefinitions of groups also get validated. There is no way to turn this
187 validation off.

188 **Flexibility**

189 Model and attribute groups allow for a greater degree of flexibility in content model/attribute list
190 modification than type derivation does. For example, you can use a model group for the middle of a
191 content model, which then allows for redefinition of only that portion of the content model; type extension
192 can only add to the end of a content model.

193 **Reuse**

194 Customizers' reuse of native constructs: Model and attribute groups ~~[@@?test this]~~ can, we believe, be
195 reused by any schemas that import the schema in which the groups are originally defined.

196 Native reuse of others' constructs: Our schemas of interest (SAML) could reuse model and attribute
197 groups defined by others' schemas. (It currently contains no model or attribute groups, native or
198 otherwise.)

199 **2.3 Element Substitution**

200 In XSD, unless an element has been declared as “blocked,” other elements with a compatible type (the
201 same or a derived type) can appear in place of it. This setting can be controlled globally throughout a
202 schema. Substitution can be thought of as a type-controlled, late-binding version of an `<xs:choice>`
203 group, which allows a mutually exclusive choice among several elements.

204 SAML currently allows all elements to be heads of substitution groups.

205 ***Tools support***

206 Some tools do not seem to fully support this mechanism, especially when combined with extension by
207 restriction.

208 ***Interactions with other mechanisms***

209 Substitution is typically used with derivable types, in order to substitute derived-type elements for base-
210 type elements. It's possible to use substitution without derivation by substituting an element with an
211 identical type, rather than a derived one, for the original element.

212 ***Validation power***

213 The type compability of the substituted element for the head element is validated at compile time by XSD
214 processors. This can't be turned off.

215 ***Flexibility***

216 Some have commented that substitution is more limited than `<xs:choice>` groups [Kawaguchi]
217 because it can only occur with compatible types rather than arbitrary types. However, substitution
218 groups allow for customizers to extend “choices” after the original SAML schema creation, which
219 `<xs:choice>` can't do. Note that type derivation used alone also allows for this possibility with the use
220 of the `xsi:type` attribute.

221 ***Reuse***

222 Does not apply.

223 **2.4 Content Models Containing “Any” Wildcards**

224 These are content models that contain the XSD `<xs:any>` and `<xs:anyAttribute>` particles. They
225 create partially or fully “open” portions of a content model, where a variety of specific elements not
226 foreseen by the original schema may appear.

227 The `<xs:anyAttribute>` wildcard is currently not used anywhere in SAML. The `<xs:any>` wildcard
228 is used in content models as follows:

- 229 • `<Advice>` in the assertion schema contains `<Assertion>`, `<AssertionIDReference>`, and
230 `##other` laxly validated elements
- 231 • `<StatusDetail>` in the protocol schema contains `##any` laxly validated elements

232 ***Tools support***

233 There seems to be reasonable support for these features in tools.

234 ***Interactions with other mechanisms***

235 When `<xs:any>` is used with the `##any` namespace setting in a content model mixed with native
236 elements bound to derivable types, and a customizer derives a new type from one of these types and
237 creates XML instances where the element bound to this derived type is supplied in an `<xs:any>` spot
238 and uses `xsi:type`, the content model becomes “ambiguous.” The workaround used in the Liberty
239 Alliance customization was to invent an `<Extension>` element to hold the `<xs:any>` particle

240 exclusively, avoiding mixing it together with native elements in a content model.
241 `<xs:any>` also causes ambiguity if a new type extends a sequence containing `<xs:any>` and adds new
242 elements to the sequence. The new elements are not distinguishable from the original wildcard. This
243 suggests that `<xs:any>` is not effective in conjunction with type derivation and may be best used in
244 types that block [\(or lack the need for\)](#) derivation.

245 The effect of the `<xs:anyAttribute>` wildcard can be achieved implicitly when the `xs:anyType`
246 ~~[`@@` or `xs:anySimpleType`? it appears so, but test]~~ is bound to an element.

247 An element must be global to satisfy an `<xs:any>` particle.

248 **Validation power**

249 The `<xs:any>` [wildcard](#) allows the schemographer to set the desired level of validation when XML
250 content fills the “any” slot: `strict` (there must be a customized schema present that provides a
251 definition for this subtree), `lax` (validate if a customized schema is present covering this subtree), or
252 `none` (no validation will be done on this subtree). Note that use of `xsi:type` supersedes this optionality
253 and forces validation.

254 **Flexibility**

255 The universe of allowed elements in an “any” slot is controllable by namespace (a list of namespaces to
256 allow or bar, or arbitrarily namespaced elements with `##any`, or just elements from foreign namespaces
257 with `##other`). The `##any` setting gives maximum choice to the customizer (but causes problems, as
258 noted above under Interactions).

259 Also, as noted just above under Validation, the flexibility around validation power is greatest in this
260 mechanism.

261 **Reuse**

262 Does not apply.

263 **2.5 Elements Bound to “Any” Types**

264 These are elements that are assigned the `xs:anyType` and `xs:anySimpleType` types. These are types
265 that allow for “open” content but that also may be extended or restricted in order to close down some
266 options.

267 The `xs:anySimpleType` type is not currently used anywhere in SAML. Two elements in the assertion
268 schema are bound to `xs:anyType`:

- 269 • `<SubjectConfirmationData>`
- 270 • `<AttributeValue>`

271 **Tools support**

272 SAML has used this mechanism without any [significant](#) tools problems surfacing, [though some older](#)
273 [parsers do occasionally exhibit problems validating mixed and attribute content](#).

274 **Interactions with other mechanisms**

275 Using `xs:anyType` ~~[`@@` and also `xs:anySimpleType`?]~~ gives the effect of `<xs:anyAttribute>` for
276 free on that element.

277 More tightly specified content models can be freely defined by extension of such an element, allowing
278 second order extensions to take an open content model and precisely define its use in that context.

279 **Validation power**

280 Since essentially any content is compatible with such an element, an XSD processor will permit anything
281 to appear without complaining. Thus, validation is essentially disabled for the element, unless `xsi:type`
282 or element substitution is used.

283 **Flexibility**

284 An element bound to `xs:anyType` essentially stands in the original schema as a placeholder for

285 whatever the same schema or an extension might define as a specific instance of that element. It thus
286 provides a "conceptual" placeholder in a content model in which any XML can be used to represent the
287 concept.

288 It also enables a combination of validating and non-validating behavior, which is useful when
289 interpretation of the non-validated content is strictly optional or governed by other agreement
290 mechanisms. SAML attribute values are an example of this.

291 **Reuse**

292 Does not apply.

293 **2.6 Global Elements and Attributes**

294 Elements and attributes that are global are allowed to be reused independently in the creation of a new
295 foreign schema. Global elements can appear in `<xs:any>` wildcard content, and can serve as the root
296 element of an XML document.

297 All of SAML's elements are global (and none of its attributes are). SAML has not consciously considered
298 this to be an extensibility mechanism, but rather a general schema style that we hoped would be well
299 supported by tools. However, it would be wise to examine the potential reuse of SAML elements in this
300 light.

301 **Tools support**

302 There seems to be very good tools support for both native global elements and reuse of them in a
303 foreign schema, as this was the first use case XSD needed to solve.

304 **Interactions with other mechanisms**

305 See [Maler] for a comparative analysis of global elements, local elements, named types, and anonymous
306 types.

307 **Validation power**

308 Global elements (as well as local ones) have declarations and are bound to types, and are thus validated
309 against these types. Unless an element is bound to something like `xs:anyType`, the validation can't
310 effectively be turned off.

311 **Flexibility**

312 Global elements are more flexible than local ones, but are also more complex because of the usage (and
313 reuse) choices they offer.

314 **Reuse**

315 The main benefit of global elements is their reuse – appearing in more than one content model of other
316 elements – in the same or importing schemas.

317 **2.7 URI-Based Identifiers and NamespacesSAML Namespaces**

318 ~~SAML defines what it calls "namespaces," after the fashion of XML namespaces, uses URI-based~~
319 ~~identifiers~~ for interpreting ~~selected SAML~~ element and attribute content correctly. They are indicated
320 through an attribute that contains a URI reference, which is designed to disambiguate which of the many
321 possible meanings for a particular ~~subtree or keyword~~~~SAML construct~~ is meant. It is also possible ~~in~~
322 ~~some cases~~ for customizers to assign a validatable type to the XML construct ~~containing the subtree or~~
323 ~~keyword~~, so that it can be syntactically validated.

324 ~~In some of the cases where URI-based identifiers are used, SAML terms the mechanism "SAML~~
325 ~~namespaces."~~ This name is typically used when the structure to be interpreted is inside the element on
326 ~~which the URI appears and when the content is keyword-like.~~

327 This mechanism differs from the other mechanisms discussed above; in that it is a ~~SAML-specific~~
328 technique ~~specific to the SAML vocabulary and not global to XSD~~. The SAML assertion schema uses ~~the~~
329 ~~namespace-this~~ mechanism in ~~two the following~~ cases:

- 330 • ~~Names of Interpreting the name of an **action** to be performed on a resource by referring to a URI-~~
331 ~~based identifier representing a set of names of actions to be performed on resources("action~~
332 ~~namespace")~~
- 333 • ~~Interpreting the name of an **attribute** by referring to a URI-based identifier representing a set of~~
334 ~~attribute names Names of attributes("attribute namespace")~~
- 335 • ~~Conveying the **authentication method** that was used for a subject in an authentication statement, or~~
336 ~~the authentication method being provided to confirm the binding of a subject to the bearer in any kind~~
337 ~~of assertion, by referring to a URI-based identifier representing an authentication method~~
- 338 • ~~Interpreting the format of a **name identifier** by referring to a URI-based identifier representing the~~
339 ~~format~~

340 **Tools support**

341 This mechanism seems well supported, though it should be noted that the “semantics” gained from it is
342 outside the knowledge of a generic XSD processor; it is specific to the SAML schemas.

343 **Interactions with other mechanisms**

344 The way SAML has defined this mechanism, it typically works hand in hand with assignment of a type for
345 syntactic checking of the content, though this validation is not required.

346 **Validation power**

347 If the namespaced content is assigned a type as discussed above, it can be validated through normal
348 XSD means. Otherwise, any validation must be application-specific. (See [Maler] for a description of an
349 XML standard that uses a different mechanism to encourage syntactic validation much more strongly.)

350 **Flexibility**

351 Because this mechanism can be entirely divorced from XSD validation, and because anyone can invent
352 a new namespace, it is very flexible. However, it may be that our underspecification of the attribute
353 namespace case may have allowed too much flexibility, causing some divergence in usage.

354 **Reuse**

355 Namespaces defined by the SAML specifications and by customizers are available to be reused by
356 others.

3 Recommendations for SAML Extensibility

357

358 The following sections contain our recommendations for modifications to the SAML schemas and
359 specifications in order to support our extensibility goals more closely.

3.1 Recommendations Related to Native Types

360

Finality ✓

361

362 Keep all types non-final.

362

363 **Rationale:** The derivation mechanism is the most well-supported for allowing SAML constructs to be
364 extended, and is known to interact successfully with other mechanisms SAML uses (or should use). The
365 extreme power of model groups for extension has not proven to be needed, may not be fully supported in
366 tools, and is an unfamiliar design pattern for most schemographers.

367 **Action:** None.

367

Abstraction ✓

368

369 Continue to make selected base types abstract and provide matching elements for them in an
370 `<xs:choice>` group with non-abstract versions. The types that should be subjected to this treatment
371 are the ones that have no useful semantics in SAML all by themselves, but have partial semantics that
372 we feel may be useful to customizers are building blocks for their own SAML-derived work.

373 **Rationale:** This seems to be the only way to make a “deep” type hierarchy available, and experience
374 with existing customizations seems to suggest it’s working successfully.

375 **Issues:** Do we have a problem with the fact that this forces customizers to define derived types off of the
376 abstract ones? We don’t necessarily care if they do this, but the mechanism gives no choice.

377 **Action:** None.

377

Substitution ✓

378

379 **Option 1:** Block all substitution because its support, particularly in combination with other mechanisms,
380 is iffy. The action in this case would be to set `blockDefault` attribute at the top level of the schema.

381 **Option 2:** Allow it, recognizing its limitations, since customizers have the right to choose it if they wish.
382 No action would be required in this case.

383 **Decision:** The TC decided on 20 January 2004 to accept Option 1.

3.2 Recommendations Related to Open Content

384

xs:anyType vs. <xs:any>

385

386 **Option 1:** Migrate to using `xs:anyType` exclusively in order to stay away from known and potential
387 problems with `<xs:any>`, allow for schema-less processing, and become consistent across the SAML
388 schemas. The action in this case would be to work on the `<Advice>` declaration and **AdviceType**
389 definition to add a new subelement that would be bound to the `xs:anyType` type, and to work on the
390 `<StatusDetail>` element declaration to bind it to `xs:anyType` instead of **StatusDetailType**, which
391 could be removed.

392 **Option 2:** Migrate to using `<xs:any>` exclusively in order to give additional validation-level flexibility to
393 customizers, allow for schema-less processing, and become consistent across the SAML schemas. Note
394 that the substitution-blocking decision made earlier mitigates the major problems with `<xs:any>`;
395 however, it should be used as the lone content of whatever content model it appears in. The action in
396 this case would be to change `<SubjectConfirmationData>` and `<AttributeValue>` to contain
397 `<xs:any namespace="##any" processContents="lax">` and to change `<Advice>` and
398 **AdviceType** to add a new subelement that would contain only `<xs:any>`.

399 **Option 3:** Develop a set of criteria that provide guidance on when to use each mechanism, measure
400 each current usage against the criteria, and synchronize the usage as necessary. NB: This set of criteria

401 ~~does not exist yet and would need to be developed! I think~~ One of the likely criteria amounts to a guess
402 ~~as to whether the element containing the <xs:any> would be likely to need derivation. Metadata is an~~
403 ~~example. I think~~ An element that amounts to an "element bag" seems like a poor candidate for
404 ~~extension and using <xs:any> seems like a way of allowing new optional "bag contents" from extension~~
405 ~~schemas to be defined while being ignored by older implementations.~~

406 **Option 4:** ~~Do nothing and keep the current ad hoc split between the two mechanisms. This requires no~~
407 ~~action, but also provides no guidance if we add future elements that need flexible/foreign content.~~

408 ~~Wherever open content is desired, bind the element to xs:anyType rather than using the~~
409 ~~<xs:any> content model particle.~~

410 **Rationale:** ~~The use of xs:anyType has proven to be more robust, flexible, and well-supported in~~
411 ~~tools than <xs:any>, and it doesn't have as many weird interactions with other mechanisms. It~~
412 ~~doesn't require customizers to build a formal schema, but this is of no concern to SAML. And it~~
413 ~~forces the open content to be sequestered in its own element, but this seems like good practice~~
414 ~~anyway for reasons of clear specification.~~

415 **Action:** ~~This would require work on the <Advice> declaration and AdviceType definition to add a~~
416 ~~new subelement that would be bound to the xs:anyType type. It would also require work on the~~
417 ~~<StatusDetail> element declaration to bind it to xs:anyType instead of StatusDetailType, which~~
418 ~~could be removed.~~

419 **<xs:anyAttribute>** ✓

420 **Option 1:** Consider adding <xs:anyAttribute> to all elements that aren't bound to **xs:anyType**
421 (which implies <xs:anyAttribute>), because allowing for open attribute content is a harmless but
422 powerful kind of extensibility that will allow SAML users to pick up and use interesting new global
423 attributes produced in other venues. The action in this case would be to add the <xs:anyAttribute>
424 particle to all complex type definitions except for those bound to xs:anyType.

425 **Option 2:** Don't do this because new attributes can always be added through type extension. No action
426 would be required in this case.

427 **Decision:** The TC decided on 20 January 2004 to take no action now, but to consider adding
428 <xs:anyAttribute> on a case-by-case basis as needed. (Note: David Orchard's article [Orchard]
429 recommending wide application of <xs:anyAttribute> has come to the authors' attention since the
430 decision was made.)

431 **3.3 Recommendations Related to Global Elements**

432 **Restrained Globalness**

433 Keep all elements global, but add prose saying which elements are allowed to be root elements, roots of
434 SOAP payloads, and so on. Keep all attributes local, unless we see fit in the future to define any
435 "common attributes" for use in SAML that we think have utility beyond SAML.

436 **Rationale:** SAML has succeeded with its global elements so far, and its vocabularies are small enough
437 that there seems to be no pressure to have multiple elements with the same name but different content
438 (the main value of local elements). However, without further qualification, global elements are more
439 powerful than we intend most of our elements to be. Adding specification prose to prohibit usage of
440 elements that are "subelements by nature" will avoid odd conformance scenarios. If we were to switch to
441 using local qualified elements to enforce the "root element" constraint in a machine-readable way, we
442 would probably cause any XSLT-based SAML processing (at the least) to have to change, since local
443 element names can be globally non-unique and thus need to be additionally anchored (e.g., simple child
444 ladders would have to be anchored all the way up to a global ancestor).

445 **Action:** This would require adding prose to mention that only certain elements – we propose to start with
446 assertions, requests, responses, name identifiers, and attributes – MAY be reused in importing schemas
447 and appear as root elements in an XML document. I think this is potentially a bit strict in a few cases. I
448 think <NameIdentifier> will be self-describing in most contexts, for one. Possibly <Attribute> would be
449 another example?

450 3.4 Recommendations Related to SAML Namespaces

451 *Attribute Clarity*

452 **Note:** Be prepared to recognize the “XML attribute” vs. “attribute of a SAML subject”
453 distinction when reading this recommendation!

454 Add a new XML attribute on the <AttributeDesignator> and <Attribute> elements to capture the
455 source of a SAML attribute.

456 **Rationale:** Rather than continuing to allow users to overload `AttributeNamespace` with multiple
457 meanings, it's better to settle on exactly what “metadata” is needed to fully disambiguate an attribute
458 name and allow SAML users to provide that metadata.

459 **Action:** This would require adding an attribute to **AttributeDesignatorType**. Note that any action here
460 may have dependencies on the resolutions of the W-28* (attribute-related) work items.

461 *Semantic vs. Syntactic Clarity*

462 For SAML attributes and actions (and future any uses of the namespace mechanism), document more
463 thoroughly the relationship between a namespace (identifying the “semantic” basis for the value) and the
464 ability to assign a type to the element that holds the value of the attribute or action (identifying the
465 syntactic constraints to which the value should be held during validation).

466 **Rationale:** The spec says little on this point today, and the mechanism is likely to be better understood
467 and more widely and correctly used in future if we explain it better.

468 **Action:** This would require some writing only: no schema changes. Note that any action here may have
469 dependencies on the resolutions of the W-28* (attribute-related) work items.

470 *Namespace Terminology*

471 Currently there's the namespace camp and the (unnamed) camp, with no solid connection between them.

472 **Option 1:** Unify the descriptions of all the parts of SAML that use the URI-based identifier mechanism,
473 focusing on the “identifier” language rather than the “namespace” language because it's more broad.
474 This would require changing the names of `AttributeNamespace` and `ActionNamespace` and
475 modifying any associated prose.

476 **Option 2:** Unify the descriptions of all the parts of SAML that use the URI-based identifier mechanism,
477 focusing on the “namespace” language rather than the “identifier” language because it's more evocative.
478 This would require modifying prose that discusses URI-based identifiers and calling them “SAML
479 namespaces.”

480 **Option 3:** Do nothing because the distinction hasn't caused any problems to date.

481 3.5 Recommendations Related to Extension Understandability

482 *Must Ignore*

483 David Orchard's article [Orchard] recommends that vocabularies adopt an explicit model stating how to
484 process extensions, with the default recommendation being the “Must Ignore” rule (“Document receivers
485 MUST ignore any XML attributes or elements in a valid XML document that they do not recognize”) in
486 combination with extensibility techniques of the sort described in this document. Another choice might be
487 a fallback model. Do we want to adopt this posture in some fashion? It might require some expression in
488 the metadata, so that entities exchanging SAML can agree on which extensions they produce and
489 expect.

4 References

490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514

- [Kawaguchi]** K. Kawaguchi. "W3C XML Schema Made Simple", XML.com, 6 June 2001. <http://www.xml.com/pub/a/2001/06/06/schemasimple.html>.
- [LibProtSchema]** S. Cantor et al., *Liberty ID-FF Protocols and Schema Specification, Version 1.2*, Liberty Alliance Project, 12 November 2003. <http://www.projectliberty.org/specs/liberty-idff-protocols-schema-v1.2.pdf>.
- [Maler]** E. Maler. "Schema Design Rules for UBL...and Maybe for You", XML 2002 paper, 8-13 December 2002. http://www.idealliance.org/papers/xml02/dx_xml02/papers/05-01-02/05-01-02.html. Note particularly Section 4. Also note that the schema examples in this section share a common error: `<xs:sequence>` is missing from around the content of `<xs:complexType>`.
- [Obasanjo1]** D. Obasanjo. "XML Schema Design Patterns: Is Complex Type Derivation Unnecessary?", XML.com, 29 October 2003. <http://www.xml.com/pub/a/2003/10/29/derivation.html>.
- [Orchard]** D. Orchard. "Versioning XML Vocabularies", XML.com, 3 December 2003. <http://www.xml.com/pub/a/2003/12/03/versioning.html>.
- [SAMLConform]** E. Maler et al. *Conformance Program Specification for the OASIS Security Assertion Markup Language (SAML)*. OASIS, September 2003. Document ID oasis-sstc-saml-conform-1.1. <http://www.oasis-open.org/committees/security/>.
- [SAMLCore]** E. Maler et al. *Assertions and Protocol for the OASIS Security Assertion Markup Language (SAML)*. OASIS, September 2003. Document ID oasis-sstc-saml-core-1.1. <http://www.oasis-open.org/committees/security/>.
- [Schema1]** H. S. Thompson et al. *XML Schema Part 1: Structures*. World Wide Web Consortium Recommendation, May 2001. <http://www.w3.org/TR/xmlschema-1/>.

A. Revision History

Rev	Date	By Whom	What
01	19 Jan 2003	Eve Maler, Scott Cantor	Initial draft.
<u>02</u>	<u>28 Jan 2003</u>	<u>Eve Maler, Scott Cantor</u>	<u>Takes into account the feedback and decisions from the 20 January 2004 SSTC telecon, additional discussions between the authors, and an interesting article from David Orchard.</u>

517

B. Notices

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

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

529 **Copyright © OASIS Open 2004. All Rights Reserved.**

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

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

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