



1

2

Security and Privacy Considerations for the OASIS Security Assertion Markup Language (SAML)

3

4

Document identifier: draft-sstc-sec-consider-04

5

6

Location: <http://www.oasis-open.org/committees/security/docs>

7

Publication date: 15 January 2002

8

Maturity Level: Committee Working Draft. This document represents work in progress upon which no reliance should be made.

9

10

Send comments to: security-services-comment@lists.oasis-open.org unless you are subscribed to the security-services list for committee members -- send comments there if so. Note: Before sending messages to the security-services-comment list, you must first subscribe. To subscribe, send an email message to security-services-comment-request@lists.oasis-open.org with the word "subscribe" as the body of the message.

11

12

13

14

15

Contributors:

16

Jeff Hodges, Oblix

17

Chris McLaren, editor (cmclaren@netegrity.com)

18

Prateek Mishra, Netegrity

19

RL "Bob" Morgan, University of Washington

20

Tim Moses, Entrust

21

Evan Prodromou, Securant

22

Marlena Erdos, IBM

23

Rev	Date	Author	What
00	xx-Aug-2001	Jeff Hodges	Created.
01	2001-11-14	Chris McLaren	First substantive draft presented to TC
03	2002-01-09	Chris McLaren	Added comments on KM, filled in additional information, added references to threats and security model in bindings, added privacy section
04	2002-01-15	Chris McLaren	Editorial cleanup

24

24		
25	Security and Privacy Considerations for the OASIS Security Assertion Markup Language	
26	(SAML)	1
27	1. Introduction	4
28	1.1. Background	4
29	1.2. Scope	8
30	1.3. SAML Threat Model	8
31	2. Security Techniques	9
32	2.1. Authentication	9
33	2.1.1. Active Session	9
34	2.1.2. Message-Level	9
35	2.2. Confidentiality	9
36	2.2.1. In Transit	10
37	2.2.2. Message-Level	10
38	2.3. Data Integrity	10
39	2.3.1. In Transit	10
40	2.3.2. Message-Level	10
41	2.4. Notes on Key Management	10
42	2.4.1. Access to the Key	10
43	2.4.2. Binding of Identity to Key	11
44	2.5. TLS/SSL Cipher Suites	11
45	2.5.1. What Is a Cipher Suite?	12
46	2.5.2. Cipher Suite Recommendations	13
47	3. SAML-Specific Security Considerations	13
48	3.1. SAML Assertions	13
49	3.2. SAML Protocol	13
50	3.2.1. Denial of Service	14
51	3.2.1.1. Requiring Client Authentication at a Lower Level	14
52	3.2.1.2. Requiring Signed Requests	14
53	3.2.1.3. Restricting Access to the Interaction URL	14
54	3.3. SAML Protocol Bindings	15
55	3.3.1. SOAP Binding	15
56	3.3.1.1. Eavesdropping	15

57	3.3.1.2.	Replay.....	16
58	3.3.1.3.	Message Insertion.....	16
59	3.3.1.4.	Message Deletion	17
60	3.3.1.5.	Message Modification.....	17
61	3.3.1.6.	Man-in-the-Middle	18
62	3.3.2.	Specifics of SOAP over HTTP.....	18
63	3.4.	Profiles for SAML.....	18
64	3.4.1.	Web Browser-Based Profiles	19
65	3.4.1.1.	Eavesdropping.....	19
66	3.4.1.1.1.	Theft of the User Authentication Information.....	19
67	3.4.1.1.2.	Theft of the Bearer Token	19
68	3.4.1.2.	Replay.....	20
69	3.4.1.3.	Message Insertion.....	20
70	3.4.1.4.	Message Deletion	20
71	3.4.1.5.	Message Modification.....	20
72	3.4.1.6.	Man-in-the-Middle	21
73	3.4.2.	Browser/Artifact Profile.....	21
74	3.4.2.1.	Replay.....	21
75	3.4.3.	Browser/POST Profile.....	22
76	3.4.3.1.	Replay.....	22
77	3.4.4.	SOAP Profile.....	22
78	3.4.4.1.	Holder of Key.....	22
79	3.4.4.1.1.	Eavesdropping.....	22
80	3.4.4.1.2.	Replay.....	23
81	3.4.4.1.3.	Message Insertion.....	23
82	3.4.4.1.4.	Message Deletion	23
83	3.4.4.1.5.	Message Modification.....	23
84	3.4.4.1.6.	Man-in-the-Middle	24
85	3.4.4.2.	Sender Vouches.....	24
86	3.4.4.2.1.	Eavesdropping.....	24
87	3.4.4.2.2.	Replay.....	24
88	3.4.4.2.3.	Message Insertion.....	25
89	3.4.4.2.4.	Message Deletion	25

90	3.4.4.2.5. Message Modification	25
91	3.4.4.2.6. Man-in-the-Middle	25
92	4. References	26
93	Appendix A. Notices	28

94

95 **1. Introduction**

96 This non-normative document describes and analyzes the security and privacy properties of the
 97 OASIS Security Assertion Markup Language (SAML) defined in the core SAML specification
 98 **[SAMLCore]** and the SAML specification for bindings and profiles **[SAMLBind]**. The intent in
 99 this document is to provide input to the design of SAML, and to provide information to
 100 architects, implementors, and reviewers of SAML-based systems about the following:

- 101 • The threats, and thus security risks, to which a SAML-based system is subject
- 102 • The security risks the SAML architecture addresses, and how it does so
- 103 • The security risks it does not address
- 104 • Recommendations for countermeasures that mitigate those risks

105 Note that terms used in this document are as defined in the SAML glossary **[SAMLGloss]** unless
 106 otherwise noted.

107 The rest of this section describes the background and assumptions underlying the analysis in this
 108 document. Section 4 provides a high-level view of security techniques and technologies that
 109 should be used with SAML. Section 5 analyzes the specific risks inherent in the use of SAML.

110 **2. Privacy**

111 SAML includes the ability to make statements about the attributes and authorizations of
 112 authenticated entities. There are very many common situations in which the information carried
 113 in these statements is something that one or more of the parties to a communication would desire
 114 to keep accessible to as restricted as possible a set of entities. Statements of medical or financial
 115 attributes are simple examples of such cases.

116 Parties making statements, issuing assertions, conveying assertions, and consuming assertions
 117 must be aware of these potential privacy concerns and should attempt to address them in their
 118 implementations of SAML-aware systems.

119 **2.1. Ensuring Confidentiality**

120 Perhaps the most important aspect of ensuring privacy to parties in a SAML-enabled transaction
 121 is the ability to carry out the transaction with a guarantee of confidentiality. In other words, can
 122 the information in an assertion be conveyed from the issuer to the intended audience, and only
 123 the intended audience, without making it accessible to any other parties?

124 It is technically possible to convey information confidentially (a discussion of common methods
125 for providing confidentiality occurs in the Security portion of the document in Section 4.2) and
126 all parties to SAML-enabled transactions should analyze each of their steps in the interaction to
127 ensure that they are taking the appropriate steps to ensure that information that should be kept
128 confidential is actually being kept so.

129 It should also be noted that simply obscuring the contents of assertions may not be adequate
130 protection of privacy. There are many cases where just the availability of the information that a
131 given user (or IP address) was accessing a given service may constitute a breach of privacy (for
132 example, an the information that a user accessed a medical testing facility for an assertion may
133 be enough to breach privacy without knowing the contents of the assertion). Partial solutions to
134 these problems can be provided by various techniques for anonymous interaction, outlined
135 below.

136 **2.2. Notes on Anonymity**

137 **2.2.1. *Definitions that Relate to Anonymity***

138 There are no definitions of anonymity which are satisfying for all cases. Many definitions
139 **[Anonymity]** deal with the simple case of a sender and a message, and discuss “anonymity” in
140 terms of not being able to link a given sender to a sent message, or a message back to a sender.

141 And while that definition is adequate for the “one off” case, it ignores the aggregation of
142 information that is possible over time based on behavior rather than an identifier.

143 Two notions which may be generally useful, and that relate to each other, can help define
144 anonymity.

145 The first notion is to think about anonymity as being “within a set”, as in this comment from
146 “Anonymity, Unobservability, and Pseudonymity” **[Anonymity]**:

147 | “To enable anonymity of a subject, there always has to be an appropriate set of subjects
148 | with potentially the same attributes....

149 | ...Anonymity is the stronger, the larger the respective anonymity set is and the more
150 | evenly distributed the sending or receiving, respectively, of the subjects within that set
151 | is”.

152

153 This notion is relevant to SAML because of the use of authorities. Even if a Subject is
154 “anonymous”, that subject is still identifiable as a member of the set of Subjects within the
155 domain of the relevant authority.

156 In the case where aggregating attributes of the user are provided, the set can become much
157 smaller. For example, if the user is “anonymous” but has the attribute of “student in Course
158 6@mit.edu”. Certainly, the number of Course 6 students is less than the number of MIT-
159 affiliated persons which is less than the number of users everywhere.

160 Why does this matter? It matters because of the second notion. This idea is that non-anonymity
161 leads to the ability of an adversary to harm expressed in Dingledine, Freedman, and Molnar’s
162 Freehaven document [**FreeHaven**]:

163 | “Both anonymity and pseudonymity protect the privacy of the user's location and true
164 | name. Location refers to the actual physical connection to the system. The term “true
165 | name” was introduced by Vinge and popularized by May to refer to the legal identity of
166 | an individual. Knowing someone's true name or location allows you to hurt him or her.”

167

168 This leads to a unification of the notion of anonymity within a set and ability to harm, from the
169 same source [**FreeHaven**]:

170 | “We might say that a system is partially anonymous if an adversary can only narrow
171 | down a search for a user to one of a ‘set of suspects.’ If the set is large enough, then it is
172 | impractical for an adversary to act as if any single suspect were guilty. On the other hand,
173 | when the set of suspects is small, mere suspicion may cause an adversary to take action
174 | against all of them.”

175

176 SAML-enabled systems are limited to "partial anonymity" at best because of the use of
177 authorities. An entity about whom an assertion is made is already identifiable as one of the pool
178 of entities in a relationship with the issuing authority.

179 The limitations on anonymity can be a lot worse than simple authority association, depending on
180 how identifiers are employed, as reuse of pseudonymous identifiers allows accretion of
181 potentially identifying information (see Section 2.2.2). Additionally, users of SAML-enabled
182 systems can also make the breach of anonymity worse by their actions (see Section 2.2.3).

183 **2.2.2. Pseudonymity & Anonymity**

184 Apart from legal identity, any identifier for a Subject can be considered a pseudonym. And even
185 notions like “holder of key” can be considered as serving as the equivalent of a pseudonym in
186 linking an action (or set of actions) to a Subject. Even a description such as “the user that just
187 requested access to object XYZ at time 23:34” can serve as an equivalent of a pseudonym.

188 The point is, that with respect to “ability to harm” it makes no difference whether the user is
189 described with an identifier or described by behavior (i.e. use of a key, or performance of an
190 action).

191 What does make a difference is how often the particular equivalent of a pseudonym is used.

192 [3] gives a taxonomy of pseudonyms starting from personal pseudonyms (like nicknames) that
193 are used all the time, through various types of role pseudonyms (e.g. Secretary of Defense), on to
194 “one time use” pseudonyms.

195 Only one time use pseudonyms can give you anonymity (within SAML, consider this as
196 "anonymity within a set").

197 The more often you use a given pseudonym, the more you reduce your anonymity and the more
198 likely it is that you can be harmed. In other words re-use of a pseudonym allows additional

199 potentially identifying information to be associated with the pseudonym. Over time this will lead
200 to an accretion that can uniquely identify the identity associated with a pseudonym.

201 **2.2.3. *Behavior and Anonymity***

202 As Joe Klein can attest, anonymity isn't all it is cracked up to be.

203 Klein is the "Anonymous" who authored Primary Colors. Despite his denials he was unmasked
204 as the author by Don Foster, a Vassar professor who did a forensic analysis of the text of Primary
205 Colors. Foster compared that text with texts from a list of suspects that he devised based on their
206 knowledge bases and writing proclivities.

207 It was Klein's idiosyncratic usages that did him in (though apparently all authors have them).

208 The relevant point for SAML is that an "anonymous" user (even one that is never named) can be
209 identified enough to be harmed by repeated unusual behavior. Here are some examples:

- 210 • A user who each Tuesday at 21:00 access a database that correlates finger lengths and life
211 span starts to be non-anonymous. Depending on that user's other behavior, she or he may
212 become "traceable" [**Pooling**] in that other "identifying" information may be able to be
213 collected.
- 214 • A user who routinely buys an usual set of products from a networked vending machine,
215 certainly opens themselves to harm (by virtue of booby-trapping the products).

216

217 **2.2.4. *Implications For Privacy***

218 Origin site authorities (i.e. Authentication Authorities and Attribute Authorities) can provide a
219 degree of "partial anonymity" by employing one-time-use identifiers or keys (for the "holder of
220 Key" case).

221 This anonymity is "partial" at best because the Subject is necessarily confined to the set of
222 Subjects in a relationship with the Authority.

223 This set may be further reduced (thus further reducing anonymity) when aggregating attributes
224 are used that further subset the user community at the origin site.

225 Users who truly care about anonymity must take care to disguise or avoid unusual patterns of
226 behavior that could serve to "de-anonymize" them over time.

227 **3. Security**

228 **3.1. Background**

229 Communication between computer-based systems is subject to a variety of threats, and these
230 threats carry some level of associated risk. The nature of the risk depends on a host of factors,
231 including the nature of the communications, the nature of the communicating systems, the
232 communication mediums, the communication environment, the end-system environments, and so

233 on. Section 3 of the IETF guidelines on writing security considerations for RFCs [Rescorla-Sec]
234 provides an overview of threats inherent in the Internet (and, by implication, intranets).
235 SAML is intended to aid deployers in establishing security contexts for application-level
236 computer-based communications within or between security domains. By serving in this role,
237 SAML addresses the “endpoint authentication” aspect (in part, at least) of communications
238 security, and also the “unauthorized usage” aspect of systems security. Communications security
239 is directly applicable to the design of SAML. Systems security is of interest mostly in the
240 context of SAML’s threat models. Section 2 of the IETF guidelines gives an overview of
241 communications security and systems security.

242 **3.2. Scope**

243 Some areas that impact broadly on the overall security of a system that uses SAML are explicitly
244 outside the scope of SAML. While this document does not address these areas, they should
245 always be considered when reviewing the security of a system. In particular, these issues are
246 important, but beyond the scope of SAML:

- 247 • Initial authentication: SAML allows statements to be made about acts of authentication
248 that have occurred, but includes no requirements or specifications for these acts of
249 authentication. Consumers of authentication assertions should be wary of blindly trusting
250 these assertions unless and until they know the basis on which they were made.
251 Confidence in the assertions must never exceed the confidence that the asserting party
252 has correctly arrived at the conclusions asserted.
- 253 • Trust Model: In many cases, the security of a SAML conversation will depend on the
254 underlying trust model, which is typically based on a key management infrastructure
255 (e.g., PKI, secret key). For example, SOAP messages secured by means of XML
256 Signature [XMLSig] are secured only insofar as the keys used in the exchange can be
257 trusted. Undetected compromised keys or revoked certificates, for example, could allow a
258 breach of security. Even failure to require a certificate opens the door for impersonation
259 attacks. PKI setup is not trivial and must be implemented correctly in order for layers
260 built on top of it (such as parts of SAML) to be secure.

261 **3.3. SAML Threat Model**

262 The general Internet threat model described in the IETF guidelines for security considerations
263 [Rescorla-Sec] is the basis for the SAML threat model. We assume here that the two or more
264 endpoints of a SAML transaction are uncompromised, but that the attacker has complete control
265 over the communications channel.

266 Additionally, due to the nature of SAML as multi-party authentication and authorization
267 statement protocol, cases must be considered where one or more of the parties in a legitimate
268 SAML transaction—who operate legitimately within their role for that transaction—attempt to
269 use information gained from a previous transaction maliciously in a subsequent transaction.

270 In all cases, the local mechanisms that systems will use to decide whether or not to generate
271 assertions are out of scope. Thus, threats arising from the details of the original login at an
272 authentication authority, for example, are out of scope as well. If an authority issues a false

273 assertion, then the threats arising from the consumption of that assertion by downstream systems
274 are explicitly out of scope.

275 The direct consequence of such a scoping is that the security of a system based on assertions as
276 inputs is only as good as the security of the system used to generate those assertions. When
277 determining what issuers to trust, particularly in cases where the assertions will be used as
278 inputs to authentication or authorization decisions, the risk of security compromises arising from
279 the consumption of false but validly issued assertions is a large one. Trust policies between
280 asserting and relying parties should always be written to include significant consideration of
281 liability and implementations must be provide an audit trail. .

282 **4. Security Techniques**

283 The following sections describe security techniques and various stock technologies available for
284 their implementation in SAML deployments.

285 **4.1. Authentication**

286 Authentication here means the ability of a party to a transaction to determine the identity of the
287 other party in the transaction. This authentication may be in one direction or it may be bilateral.

288 **4.1.1. *Active Session***

289 Non-persistent authentication is provided by the communications channel used to transport a
290 SAML message. This authentication may be unilateral—from the session initiator to the
291 receiver—or bilateral. The specific method will be determined by the communications protocol
292 used. For instance, the use of a secure network protocol, such as RFC 2246 [**RFC2246**] or the IP
293 Security Protocol [**IPsec**], provides the SAML message sender with the ability to authenticate the
294 destination for the TCP/IP environment.

295 **4.1.2. *Message-Level***

296 XML Signature [**XMLSig**] provides a method of creating a persistent “authentication” that is
297 tightly coupled to a document. This method does not independently guarantee that the sender of
298 the message is in fact that signer (and indeed, in many cases where intermediaries are involved,
299 this is explicitly not the case).

300 Any method that allows the persistent confirmation of the involvement of a uniquely resolvable
301 entity with a given subset of an XML message is sufficient to meet this requirement.

302 **4.2. Confidentiality**

303 Confidentiality means that the contents of a message can be read only by the desired recipients
304 and not anyone else who encounters the message while it is in transit.

305 **4.2.1. *In Transit***

306 Use of a secure network protocol such as RFC 2246 [RFC2246] or the IP Security Protocol
307 [IPsec] provides transient confidentiality of a message as it is transferred between two nodes.

308 **4.2.2. *Message-Level***

309 XML Encryption [XMLEnc] is a draft specification for the selective encryption of XML
310 documents. This encryption method provides persistent, selective confidentiality of elements
311 within an XML message.

312 Until XML Encryption is an accepted standard, confidentiality may be implemented in transit
313 (and not end-to-end) by reliance on transports that provide in-transit confidentiality (as described
314 in Section 4.2.1 above).

315 **4.3. Data Integrity**

316 Data integrity is the ability to confirm that a given message as received is unaltered from the
317 version of the message that was sent.

318 **4.3.1. *In Transit***

319 Use of a secure network protocol such as RFC 2246 [RFC2246] or the IP Security Protocol
320 [IPsec] may be configured so as to provide for integrity check CRCs of the packets transmitted
321 via the network connection.

322 **4.3.2. *Message-Level***

323 XML Signature [XMLSig] provides a method of creating a persistent guarantee of the unaltered
324 nature of a message that is tightly coupled to that message.

325 Any method that allows the persistent confirmation of the unaltered nature of a given subset of
326 an XML message is sufficient to meet this requirement.

327 **4.4. Notes on Key Management**

328 Many points in this document will refer to the ability of systems to provide authentication, data
329 integrity, and non-repudiation via various schemes involving digital signature and encryption.
330 For all these schemes the security provided by the scheme is limited based on the key
331 management systems that are in place. Some specific limitations are detailed below:

332 **4.4.1. *Access to the Key***

333 It is assumed that if key-based systems are going to be used for authentication, data integrity, and
334 non-repudiation, that security is in place to guarantee that access to the key is not available to
335 inappropriate parties. For example, a digital signature created with Bob's private key is only
336 proof of Bob's involvement to the extent that Bob is the only one with access to the key.

337 In general, access to keys should be kept to the minimum set of entities possible (particularly
338 important for corporate or organizational keys, etc.) and should be protected with pass phrases
339 and other means. Standard security precautions (don't write down the passphrase, don't leave a
340 window with the key accessed open when you're away from a computer, etc.) apply.

341 **4.4.2. *Binding of Identity to Key***

342 For a key-based system to be used for authentication there must be some trusted binding of
343 identity to key. Verifying a digital signature on a document can determine if the document is
344 unaltered since its signature, and that it was actually signed by a given key. However, this is no
345 way confirms that the key used is actually the key of a specific individual.

346 This key-to-individual binding must be established. Common solutions include local directories
347 that store both identifiers and key—which is simple to understand but difficult to maintain—or
348 the use of certificates.

349 Certificates, which are in essence signed bindings of identity-to-key are a particularly powerful
350 solution to the problem, but come with their own considerations. A set of trusted root Certifying
351 Authorities (CAs) must be identified for each consumer of signatures—i.e. “Who do I trust to
352 make statements of identity-to-key binding”. Verification of a signature then becomes a process
353 of verifying first the signature (to determine that the signature was done by the key in question
354 and that the message has not changed) and then verification of the certificate chain (to determine
355 that the key is bound to the right identity).

356 Additionally, with certificates steps must be taken to ensure that the binding is currently valid—a
357 certificate typically has a “lifetime” built into it, but if a key is compromised during the life of
358 the certificate then the key-to-identity binding contained in the certificate becomes invalid while
359 the certificate is still valid on its face. Also certificates often depend on associations that may end
360 before their lifetime expires (for example certificates that should become invalid when someone
361 changes employers, etc.) This problem is solved by Certificate Revocation Lists (CRLs) which
362 are lists of certificates from a given CA that have been revoked since their issue. Another
363 solution is the Online Certificate Status Protocol (OCSP) which defines a method for calling
364 servers to ask about the current validity of a given certificate. Some of this same functionality is
365 incorporated into the higher levels of the XML Key Management Specification (XKMS) which
366 allows requests to be made for “valid” keys.

367 A proper key management system is thus quite strong but very complex. Verifying a signature
368 ends up being a three-stage process of verifying the document-to-key binding, then verifying the
369 key-to-identity binding, then verifying the current validity of the key-to-document binding.

370 **4.5. TLS/SSL Cipher Suites**

371 The use of SSL 3.0 or TLS 1.0 (RFC 2246) [**RFC2246**] over HTTP is recommended at many
372 places in this document. However TLS/SSL can be configured to use many different cipher
373 suites, not all of which are adequate to provide “best practices” security. The following sections
374 provide a brief description cipher suites and recommendations for cipher suite selection.

375 **4.5.1. What Is a Cipher Suite?**

376 **Note:** *While references to the US Export restrictions are now obsolete, the constants naming the*
377 *cipher suites have not changed. Thus, `SSL_DHE_DSS_EXPORT_WITH_DES40_CBC_SHA` is still*
378 *a valid cipher suite identifier, and the explanation of the historical reasons for the inclusion of*
379 *“EXPORT” has been left in place in the following summary.*

380 A cipher suite combines four kinds of security features, and is given a name in the SSL protocol
381 specification. Before data flows over a SSL connection, both ends attempt to negotiate a cipher
382 suite. This lets them establish an appropriate quality of protection for their communications,
383 within the constraints of the particular mechanism combinations which are available. The
384 features associated with a cipher suite are:

- 385 1. The type of key exchange algorithm used. SSL defines many; the ones that provide server
386 authentication are the most important ones, but anonymous key exchange is supported.
387 (Note that anonymous key exchange algorithms are subject to “man in the middle”
388 attacks, and are **not recommended** in the SAML context.) The “RSA” authenticated key
389 exchange algorithm is currently the most interoperable algorithm. Another important key
390 exchange algorithm is the authenticated Diffie-Hellman “DHE_DSS” key exchange,
391 which has no patent-related implementation constraints.¹
- 392 2. Whether the key exchange algorithm is freely exportable from the United States of
393 America. Exportable algorithms must use short (512-bit) public keys for key exchange
394 and short (40-bit) symmetric keys for encryption. These keys are currently subject to
395 breaking in an afternoon by a moderately well-equipped adversary.
- 396 3. The encryption algorithm used. The fastest option is the RC4 stream cipher; DES and
397 variants (DES40, 3DES-EDE) are also supported in "cipher block chaining" (CBC)
398 mode, as is null encryption (in some suites). (Null encryption does nothing; in such cases
399 SSL is used only to authenticate and provide integrity protection. Cipher suites with null
400 encryption do not provide confidentiality, and **should not be used** in cases where
401 confidentiality is a requirement.)
- 402 4. The digest algorithm used for the Message Authentication Code. The choices are MD5
403 and SHA1.

404 For example, the cipher suite named `SSL_DHE_DSS_EXPORT_WITH_DES40_CBC_SHA`
405 uses SSL, an authenticated Diffie-Hellman key exchange (DHE_DSS), is export grade
406 (EXPORT), uses an exportable variant of the DES cipher (DES40_CBC), and uses the SHA1
407 digest algorithm in its MAC (SHA).

408 A given implementation of SSL will support a particular set of cipher suites, and some subset of
409 those will be enabled by default. Applications have a limited degree of control over the cipher
410 suites that are used on their connections; they can enable or disable any of the supported cipher
411 suites, but cannot change the cipher suites which are available.

¹ RSA patents have all expired, hence this issue is mostly historical.

412 **4.5.2. *Cipher Suite Recommendations***

413 The following cipher suites adequately meet requirements for confidentiality and message
414 integrity, and can be configured to meet the authentication requirement as well (by forcing the
415 presence of X.509v3 certificates). They are also well supported in many client applications.
416 Support of these suites is recommended:

- 417 • TLS_RSA_WITH_3DES_EDE_CBC_SHA (when using TLS)
- 418 • SSL_RSA_WITH_3DES_EDE_CBC_SHA (when using SSL)

419 However, the IETF is moving rapidly towards mandating the use of AES, which has both speed
420 and strength advantages. Forward-looking systems would be wise as well to implement support
421 for the AES cipher suites, such as:

- 422 • TLS_RSA_WITH_AES_128_CBC_SHA

423 **5. SAML-Specific Security Considerations**

424 The following sections analyze the security risks in using and implementing SAML and describe
425 countermeasures to mitigate the risks.

426 **5.1. SAML Assertions**

427 At the level of the SAML assertion itself, there is little to be said about security concerns—most
428 concerns arise during communications in the request/response protocol, or during the attempt to
429 use SAML by means of one of the bindings. However, one issue at the assertion level bears
430 analysis: An assertion, once issued, is out of the control of the issuer.

431 This fact has a number of ramifications. For example, the issuer has no control over how long the
432 assertion will be persisted in the systems of the consumer; nor does the issuer have control over
433 the parties with whom the consumer will share the assertion information. These concerns are
434 over and above concerns about a malicious attacker who can see the contents of assertions that
435 pass over the wire unencrypted (or insufficiently encrypted).

436 While efforts have been made to address many of these issues within the SAML specification,
437 nothing contained in the specification will erase the requirement for careful consideration of
438 what to put in an assertion. At all times, issuers should consider the possible consequences if the
439 information in the assertion is stored on a remote site, where it can be directly misused, or
440 exposed to potential hackers, or possibly stored for more creatively fraudulent uses. Issuers
441 should also consider the possibility that the information in the assertion could be shared with
442 other parties, or even made public, either intentionally or inadvertently.

443 **5.2. SAML Protocol**

444 The following sections describe security considerations for the SAML request-response protocol
445 itself, apart from any threats arising from use of a particular protocol binding.

446 **5.2.1. Denial of Service**

447 The SAML protocol is susceptible to a denial of service (DOS) attack. Handling a SAML request
448 is potentially a very expensive operation, including parsing the request message (typically
449 involving construction of a DOM tree), database/assertion store lookup (potentially on an
450 unindexed key), construction of a response message, and potentially one or more digital
451 signature operations. Thus, the effort required by an attacker generating requests is much lower
452 than the effort needed to handle those requests.

453 **5.2.1.1. Requiring Client Authentication at a Lower Level**

454 Requiring clients to authenticate at some level below the SAML protocol level (for example,
455 using the SOAP over HTTP binding, with HTTP over TLS/SSL, and with a requirement for
456 client-side certificates that have a trusted Certificate Authority at their root) will provide
457 traceability in the case of a DOS attack.

458 If the authentication is used only to provide traceability then this does not in itself prevent the
459 attack from occurring, but does function as a deterrent.

460 If the authentication is coupled with some access control system, then DOS attacks from non-
461 insiders is effectively blocked. (Note that it is possible that overloading the client-authentication
462 scheme could still function as a denial-of-service attack on the SAML service, but that this attack
463 needs to be dealt with in the context of the client authentication scheme chosen.)

464 Whatever system of client authentication is used, it should provide the ability to resolve a unique
465 originator for each request, and should not be subject to forgery. (For example, in the
466 traceability-only case, logging the IP address is insufficient since this information can easily be
467 spoofed.)

468 **5.2.1.2. Requiring Signed Requests**

469 In addition to the benefits gained from client authentication discussed in Section 5.2.1.1,
470 requiring a signed request also lessens the order of the asymmetry between the work done by
471 requester and responder. The additional work required of the responder to verify the signature is
472 a relatively small percentage of the total work required of the responder, while the process of
473 calculating the digital signature represents a relatively large amount of work for the requester.
474 Narrowing this asymmetry decreases the risk associated with a DOS attack.

475 Note however that an attacker can theoretically capture a signed message and then replay it
476 continually, getting around this requirement. This situation can be avoided by requiring the use
477 of the XML Signature element `<ds:SignatureProperties>` containing a timestamp; the
478 timestamp can then be used to determine if the signature is recent. In this case, the narrower the
479 window of time after issue that a signature is treated as valid, the higher security you have
480 against replay denial of service attacks.

481 **5.2.1.3. Restricting Access to the Interaction URL**

482 Limiting the ability to issue a request to a SAML service at a very low level to a set of known
483 parties drastically reduces the risk of a DOS attack. In this case, only attacks originating from

484 within the finite set of known parties are possible, greatly decreasing exposure both to potentially
485 malicious clients and to DOS attacks using compromised machines as zombies.

486 There are many possible methods of limiting access, including placing the SAML responder
487 inside a secured intranet, implementing access rules at the router level, etc.

488 **5.3. SAML Protocol Bindings**

489 The security considerations in the design of the SAML request-response protocol depend to a
490 large extent on the particular protocol binding (as defined in the SAML bindings specification
491 **[SAMLBind]**) that is used. Currently the only binding sanctioned by the OASIS SAML
492 Committee is the SOAP binding.

493 **5.3.1. *SOAP Binding***

494 Since the SAML SOAP binding requires no authentication and has no requirements for either in-
495 transit confidentiality or message integrity, it is open to a wide variety of common attacks, which
496 are detailed in the following sections. General considerations are discussed separately from
497 considerations related to the SOAP-over-HTTP case.

498 **5.3.1.1. *Eavesdropping***

499 Since there is no in-transit confidentiality requirement, it is possible that an eavesdropping party
500 could acquire both the SOAP message containing a request and the SOAP message containing
501 the corresponding response. This acquisition exposes both the nature of the request and the
502 details of the response, possibly including one or more assertions.

503 Exposure of the details of the request will in some cases weaken the security of the requesting
504 party by revealing details of what kinds of assertions it requires, or from whom those assertions
505 are requested. For example, if an eavesdropper can determine that site *X* is frequently requesting
506 authentication assertions with a given confirmation method from site *Y*, he may be able to use
507 this information to aid in the compromise of site *X*.

508 Similarly, eavesdropping on a series of authorization queries could create a “map” of resources
509 that are under the control of a given authorization authority.

510 Additionally, in some cases exposure of the request itself could constitute a violation of privacy.
511 For example, eavesdropping on a query and its response may expose that a given user is active
512 on the querying site, which could be information that should not be divulged in cases such as
513 medical information sites, political sites, and so on. Also the details of any assertions carried in
514 the response may be information that should be kept confidential. This is particularly true for
515 responses containing attribute assertions; if these attributes represent information that should not
516 be available to entities not party to the transaction (credit ratings, medical attributes, and so on),
517 then the risk from eavesdropping is high.

518 In cases where any of these risks is a concern, the countermeasure for eavesdropping attacks is to
519 provide some form of in-transit message confidentiality. For SOAP messages, this
520 confidentiality can be enforced either at the SOAP level or at the SOAP transport level (or some
521 level below it).

522 Adding in-transit confidentiality at the SOAP level means constructing the SOAP message such
523 that, regardless of SOAP transport, no one but the intended party will be able to access the
524 message. The general solution to this problem is likely to be XML Encryption [XMLEnc]. This
525 draft specification allows encryption of the SOAP message itself, which eliminates the risk of
526 eavesdropping unless the key used in the encryption has been compromised. Alternatively, until
527 XML Encryption is widely supported, deployers will need to depend on the SOAP transport
528 layer, or a layer beneath it, to provide in-transit confidentiality.

529 The details of how to provide this confidentiality depend on the specific SOAP transport chosen.
530 Using HTTP over TLS/SSL (described further in Section 5.3.2) is one method. Other transports
531 will necessitate other in-transit confidentiality techniques; for example, an SMTP transport might
532 use S/MIME.

533 In some cases, a layer beneath the SOAP transport might provide the required in-transit
534 confidentiality. For example, if the request-response interaction is carried out over an IPsec
535 tunnel, then adequate in-transit confidentiality may be provided by the tunnel itself.

536 **5.3.1.2. Replay**

537 There is little vulnerability to replay attacks at the level of the SOAP binding. Replay is more of
538 an issue in the various profiles. The primary concern about replay at the SOAP binding level is
539 the potential for use of replay as a denial-of-service attack method.

540 In general, the best way to prevent replay attacks is to prevent the message capture in the first
541 place. Some of the transport-level schemes used to provide in-transit confidentiality will
542 accomplish this goal. For example, if the SAML request-response conversation occurs over
543 SOAP on HTTP/TLS, third parties are prevented from capturing the messages.

544 Note that since the potential replayer does not need to understand the message to replay it,
545 schemes such as XML Encryption do not provide protection against replay. If an attacker can
546 capture a SAML request that has been signed by the requester and encrypted to the responder,
547 then the attacker can replay that request at any time without needing to be able to undo the
548 encryption. This is a particular issue since the SAML request does not include information about
549 the issue time of the request, thus making it difficult to determine if replay is occurring. The only
550 recourse is to design systems that use the unique key of the request (its RequestID) to determine
551 if this is a replay request or not.

552 Additional threats from the replay attack include cases where a “charge per request” model is in
553 place. Replay could be used to run up large charges on a given account.

554 Similarly models where a client is allocated (or purchases) a fixed number of interactions with a
555 system, the replay attack could exhaust these uses unless the issuer is careful to keep track of the
556 unique key of each Request.

557 **5.3.1.3. Message Insertion**

558 The message insertion attack for the SOAP binding amounts to the creation of a request. The
559 ability to make a request is not a threat at the SOAP binding level.

560 **5.3.1.4. Message Deletion**

561 The message deletion attack would either prevent a request from reaching a responder, or would
562 prevent the response from reaching the requestor.

563 In either case, the SOAP binding does not address this threat. The SOAP protocol itself, and the
564 transports beneath it, may provide some information depending on how the message deletion is
565 accomplished.

566 Examples of reliable messaging systems that attenuate this risk include reliable HTTP (HTTPR)
567 [HTTPR] at the transport layer and the use of reliable messaging extensions in SOAP such as
568 Microsoft's SRMP for MSMQ [SRMPPres].

569 **5.3.1.5. Message Modification**

570 Message modification is a threat to the SOAP binding in both directions.

571 Modification of the request to alter the details of the request can result in significantly different
572 results being returned, which in turn can be used by a clever attacker to compromise systems
573 depending on the assertions returned. For example, altering the list of requested attributes in the
574 <AttributeDesignator> elements could produce results leading to compromise or rejection of
575 the request by the responder.

576 Modification of the request to alter the apparent issuer of the request could result in denial of
577 service or incorrect routing of the response. This alteration would need to occur below the
578 SAML level and is thus out of scope.

579 Modification of the response to alter the details of the assertions therein could result in vast
580 degrees of compromise. The simple examples of altering details of an authentication or an
581 authorization decision could lead to very serious security breaches.

582 In order to address these potential threats, a system that guarantees in-transit message integrity
583 must be used. The SAML protocol and the SOAP binding neither require nor forbid the
584 deployment of systems that guarantee in-transit message integrity, but due to this large threat, it
585 is **highly recommended** that such a system be used. At the SOAP binding level, this can be
586 accomplished by digitally signing requests and responses with a system such as XML Signature
587 [XMLSig]. The SAML specification allows for such signatures see the SAML Core
588 Specification [SAMLCore] Section 5 for further information.

589 If messages are digitally signed (with a sensible key management infrastructure, see Section 4.4)
590 then the recipient has a guarantee that the message has not be altered in transit, unless the key
591 used has been compromised.

592 The goal of in-transit message integrity can also be accomplished at a lower level by using a
593 SOAP transport that provides the property of guaranteed integrity, or is based on a protocol that
594 provides such a property. SOAP over HTTP over TLS/SSL is a transport that would provide
595 such a guarantee.

596 Encryption alone does not provide this protection, as even if the intercepted message could not
597 be altered per se, it could be replaced with a newly created one.

598 **5.3.1.6. Man-in-the-Middle**

599 The SOAP binding is susceptible to man-in-the-middle (MITM) attacks. In order to prevent
600 malicious entities from operating as a man in the middle (with all the perils discussed in both the
601 eavesdropping and message modification), some sort of bilateral authentication is required.

602 A bilateral authentication system would allow both parties to determine that what they are seeing
603 in a conversation actually came from the other party to the conversation.

604 At the SOAP binding level, this goal could also be accomplished by digitally signing both
605 requests and responses (with all the caveats discussed in Section 5.3.1.5 above). This method
606 does not prevent an eavesdropper from sitting in the middle and forwarding both ways, but he is
607 prevented from altering the conversation in any way without being detected.

608 Since many applications of SOAP do not use sessions, this sort of authentication of author (as
609 opposed to authentication of sender) may need to be combined with information from the
610 transport layer to confirm that the sender and the author are the same party in order to prevent a
611 weaker form of “MITM as eavesdropper”.

612 Another implementation would depend on a SOAP transport that provides, or is implemented on
613 a lower layer that provides, bilateral authentication. The example of this is again SOAP over
614 HTTP over TLS/SSL with both server- and client-side certificates required.

615 Additionally, the validity interval of the assertions returned functions as an adjustment on the
616 degree of risk from MITM attacks. The shorter the valid window of the assertion, the less
617 damage can be done if it is intercepted.

618 **5.3.2. Specifics of SOAP over HTTP**

619 Since the SOAP binding requires that conformant applications support HTTP over TLS/SSL
620 with a number of different bilateral authentication methods such as Basic over server-side SSL,
621 certificate-backed authentication over server-side SSL, these methods are always available to
622 mitigate threats in cases where other lower-level systems are not available and the above listed
623 attacks are considered significant threats.

624 This does not mean that use of HTTP over TLS with some form of bilateral authentication is
625 mandatory.. If an acceptable level of protection from the various risks can be arrived at through
626 other means (for example, by an IPsec tunnel), full TLS with certificates is not required.
627 However, in the majority of cases for SOAP over HTTP, using HTTP over TLS with bilateral
628 authentication will be the appropriate choice.

629 Note, however, that the use of transport-level security (such as the SSL or TLS protocols under
630 HTTP) only provides confidentiality and/or integrity and/or authentication for “one hop”. For
631 models where there may be intermediaries, or the assertions in question need to live over more
632 than one hop, the use of HTTP with TLS/SSL does not provide adequate security.

633 **5.4. Profiles for SAML**

634 The SAML bindings specification [**SAMLBind**] in addition defines profiles for SAML, which
635 are sets of rules describing how to embed and extract SAML assertions into a framework or

636 protocol. Currently there are three profiles for SAML that are sanctioned by the OASIS SAML
637 Committee:

- 638 • Two web browser-based profiles that support single sign-on (SSO):
 - 639 ○ The browser/artifact profile for SAML
 - 640 ○ The browser/POST profile for SAML
- 641 • The SOAP profile for SAML

642 **5.4.1. *Web Browser-Based Profiles***

643 The following sections describe security considerations that are common to the browser/artifact
644 and browser/POST profiles for SAML.

645 Note that user authentication at the source site is explicitly out of scope, as are all issues that
646 arise from it. The key notion is that the source system entity must be able to ascertain that it is
647 the same authenticated client system entity that it is interacting with in the next interaction step.
648 One way to accomplish this is for these initial steps to be performed using TLS as a session layer
649 underneath the protocol being used for this initial interaction (likely HTTP).

650 **5.4.1.1. *Eavesdropping***

651 The possibility of eavesdropping exists in all web browser cases. In cases where confidentiality
652 is required (bearing in mind that any assertion that is not sent securely, along with the requests
653 associated with it, is available to the malicious eavesdropper), HTTP traffic needs to take place
654 over a transport that ensures confidentiality. HTTP over TLS/SSL [RFC2246] and the IP
655 Security Protocol [IPsec] meet this requirement.

656 The following sections provide more detail on the eavesdropping threat.

657 **5.4.1.1.1. *Theft of the User Authentication Information***

658 In the case where the subject authenticates to the source site by revealing authentication
659 information, for example, in the form of a password, theft of the authentication information will
660 enable an adversary to impersonate the subject.

661 In order to avoid this problem, the connection between the subject's browser and the source site
662 must implement a confidentiality safeguard. In addition, steps must be taken by either the subject
663 or the destination site to ensure that the source site is genuinely the expected and trusted source
664 site before revealing the authentication information. Using HTTP over TLS can be used to
665 address this concern.

666 **5.4.1.1.2. *Theft of the Bearer Token***

667 In the case where the authentication assertion contains the assertion bearer authentication
668 protocol identifier, theft of the artifact will enable an adversary to impersonate the subject.

669 Each of the following methods decreases the likelihood of this happening:

- 670 • The destination site implements a confidentiality safeguard on its connection with the
671 subject's browser.
- 672 • The subject or destination site ensures (out of band) that the source site implements a
673 confidentiality safeguard on its connection with the subject's browser.
- 674 • The destination site verifies that the subject's browser was directly redirected by a source
675 site that directly authenticated the subject.
- 676 • The source site refuses to respond to more than one request for an assertion
677 corresponding to the same assertion ID.
- 678 • If the assertion contains a condition element of type AudienceRestrictionConditionType
679 that identifies a specific domain, then the destination site verifies that it is a member of
680 that domain.
- 681 • The connection between the destination site and the source site, over which the assertion
682 ID is passed, is implemented with a confidentiality safeguard.
- 683 • The destination site, in its communication with the source site, over which the assertion
684 ID is passed, must verify that the source site is genuinely the expected and trusted source
685 site.

686 **5.4.1.2. Replay**

687 The possibility of a replay attack exists for this set of profiles. A replay attack can be used either
688 to attempt to deny service or to retrieve information fraudulently. The specific countermeasures
689 depend on which specific profile is being used, and thus are discussed in Sections 5.4.2.1 and
690 5.4.3.1.

691 **5.4.1.3. Message Insertion**

692 Message insertion attacks are not a general threat in this set of profiles.

693 **5.4.1.4. Message Deletion**

694 Deleting a message during any step of the interactions between the browser, SAML assertion
695 issuer, and SAML assertion consumer will cause the interaction to fail. It results in a denial of
696 some service but does not increase the exposure of any information.

697 The SAML bindings and profiles specification provides no countermeasures for message
698 deletion.

699 **5.4.1.5. Message Modification**

700 The possibility of alteration of the messages in the stream exists for this set of profiles. Some
701 potential undesirable results are as follows:

- 702 • Alteration of the initial request can result in rejection at the SAML issuer, or creation of
703 an artifact targeted at a different resource than the one requested

- 704 • Alteration of the artifact can result in denial of service at the SAML consumer.
705 • Alteration of the assertions themselves while in transit could result in all kinds of bad
706 results (if they are unsigned) or denial of service (if they are signed and the consumer
707 rejects them).

708 To avoid message modification, the traffic needs to be transported by means of a system that
709 guarantees message integrity from endpoint to endpoint.

710 For the web browser-based profiles, the recommended method of providing message integrity in
711 transit is the use of HTTP over TLS/SSL with a cipher suite that provides data integrity
712 checking.

713 **5.4.1.6. Man-in-the-Middle**

714 Man-in-the-middle attacks are particularly pernicious for this set of profiles. The MITM can
715 relay requests, capture the returned assertion (or artifact), and relay back a false one. Then the
716 original user cannot access the resource in question, but the MITM can do so using the captured
717 resource.

718 Preventing this threat requires a number of countermeasures. First, using a system that provides
719 strong bilateral authentication will make it much more difficult for a MITM to insert himself into
720 the conversation.

721 However the possibility still exists of a MITM who is purely acting as a bidirectional port
722 forwarder, and eavesdropping on the information with the intent to capture the returned assertion
723 or handler (and possibly alter the final return to the requester). Putting a confidentiality system in
724 place will prevent eavesdropping. Putting a data integrity system in place will prevent alteration
725 of the message during port forwarding.

726 For this set of profiles, all the requirements of strong bilateral session authentication,
727 confidentiality, and data integrity can be met by the use of HTTP over TLS/SSL if the TLS/SSL
728 layer uses an appropriate cipher suite (strong enough encryption to provide confidentiality, and
729 supporting data integrity) and requires X509v3 certificates for authentication.

730 **5.4.2. Browser/Artifact Profile**

731 Many specific threats and counter-measures for the Browser/Artifact profile are documented
732 normatively in the SAML bindings specification [**SAMLBind**] Section 4.1.1.7. Additional non-
733 normative comments are included below.

734 **5.4.2.1. Replay**

735 The threat of replay as a reuse of an artifact is addressed by the requirement that each artifact is a
736 one-time-use item. Systems should track cases where multiple requests are made referencing the
737 same artifact, as this situation may represent intrusion attempts.

738 The threat of replay on the original request that results in the assertion generation is not
739 addressed by SAML, but should be mitigated by the original authentication process.

740 **5.4.3. Browser/POST Profile**

741 Many specific threats and counter-measures for the Browser/POST profile are documented
742 normatively in the SAML bindings specification [**SAMLBind**] Section 4.1.2.5. Additional non-
743 normative comments are included below.

744 **5.4.3.1. Replay**

745 Replay attacks amount to resubmission of the form in order to access a protected resource
746 fraudulently. The profile mandates that the assertions transferred have the one-use property at the
747 destination site, preventing replay attacks from succeeding.

748 **5.4.4. SOAP Profile**

749 This profile defines methods for securely attaching SAML assertions to a SOAP document.
750 SOAP documents are used in multiple contexts, specifically including cases where the message
751 is transported without an active session, the message can be persisted, and the message is routed
752 through a number of intermediaries. Such a general context of use suggests that users of this
753 profile must be concerned with a variety of threats. In particular, no consideration has been given
754 to the issue of sender or receiver authentication. Therefore, if required, the sender may need to
755 authenticate the receiver using some authentication technique dependent on the context of use.
756 Further, the receiver may need to authenticate the sender using some techniques dependent on
757 the context of use. In the latter case, there is a possibility that the receiver may authenticate the
758 sender utilizing the attached SAML assertions as a credential together with other information.

759 The SAML bindings and profiles specification [**SAMLBind**], Section 4.2.3, provides more
760 information about security considerations for this profile.

761 **5.4.4.1. Holder of Key**

762 This profile has one or more authorities issuing assertions that contain <SubjectConfirmation>
763 elements that basically say “This assertion is valid if it is presented with proof that the presenter
764 is the holder of the specified key”.

765 A sender inserts these assertions in a message and the entire message (payload and assertions)
766 are digitally signed using the specified key—thus providing proof to the receiver that the sender
767 of the message held the key specified in the assertions.

768 **5.4.4.1.1. Eavesdropping**

769 Eavesdropping continues to be a threat in the same manner as for the SAML SOAP binding, as
770 discussed in Section 5.3.1.1. The routable nature of SOAP adds the potential for a large number
771 of steps and actors in the course of a message’s lifetime, which means that the potential
772 incidences of eavesdropping are increased as the number of possible times a message is in transit
773 increases.

774 The persistent nature of SOAP messages adds an additional possibility of eavesdropping, in that
775 stored items can be read from their store.

776 To provide maximum protection from eavesdropping, assertions should be encrypted in such a
777 way that only the intended audiences can view the material. This removes threats of
778 eavesdropping in transit, but does not remove risks associated with storage by the receiver or
779 poor handling of the clear text by the receiver.

780 **5.4.4.1.2. *Replay***

781 Binding of assertions to a document opens the door to replay attacks by a malicious user. Issuing
782 a `HolderOfKey` assertion amounts to “blessing the user’s key” for the purpose of binding
783 assertions to documents. Once a `HolderOfKey` assertion has been issued to a user, that user can
784 bind it to any document or documents he chooses.

785 While each assertion is signed, and bound by a second signature into a document, which prevents
786 a malicious third-party (who has no access to the private key required for the binding signature)
787 from binding the assertions to arbitrary documents, there is nothing preventing a malicious **user**
788 (who by definition has access to the private key) from detaching a signed assertion from the
789 document it arrived in and rebinding it to another document.

790 There are two lines of defense against this type of attack. The first is to consider carefully to
791 whom you issue `HolderOfKey` assertions (can they be trusted with the right to attach the
792 assertion to any document?) and what kind of assertions you issue as `HolderOfKey` assertions
793 (do you want to give up control over the binding of this particular statement to a given
794 document?). The second is a short lifetime on the assertion, to narrow the window of opportunity
795 for this attack.

796 The capture and resubmission of the entire message (SAML assertions and business payload) is a
797 threat. One counter-measure is to add information about time, or a sequence number to the
798 digital signature included in the SOAP header. The receiver can use this information to detect
799 duplicate messages.

800 **5.4.4.1.3. *Message Insertion***

801 There is no message insertion attack at the level of the `HolderOfKey` format of the SOAP profile.

802 **5.4.4.1.4. *Message Deletion***

803 There is no message deletion attack at the level of the `HolderOfKey` format of the SOAP profile.

804 **5.4.4.1.5. *Message Modification***

805 The double signing in this profile prevents most message modification attacks. The receiver is
806 always able to verify the signature on the assertion itself (and should be able to verify that the
807 key used in that signing act is associated with the putative signer by means of X509v3 certificate,
808 Certificate Revocation List checks, and so on), which provides a guarantee that the assertion is
809 unaltered.

810 The receiver can also verify the binding signature to ensure that the message to which the
811 assertion is attached is unaltered.

812 The profile is secure against modification within the context of an existing trust relationship. The
813 remaining threats (compromised keys, revoked certificates being used, and so on) are outside the
814 scope of SAML.

815 Note that the threat of message modification by the holder of the key exists, as discussed in the
816 discussion of replay attacks in Section 5.4.4.1.2.

817 **5.4.4.1.6. *Man-in-the-Middle***

818 An MITM attack is impossible for the `HolderOfKey` format of the SOAP profile, since the
819 assertion specifies the key that must be used for the binding signature, and the assertion itself is
820 protected against tampering by a signature.

821 The MITM can eavesdrop (if communication is not protected by some confidentiality scheme)
822 but cannot alter the document without detection.

823 Note that a MITM could alter parts of the document unprotected by the signature (i.e. the other
824 header elements within the `<Signature>` element). For example, a MITM could remove an
825 included `<KeyInfo>` block from a `<Signature>` without affecting the validity of the signature.
826 Theoretically this could force an XKMS lookup or other network call that could be perverted to
827 malicious ends. However this does not pose a threat for the `HolderOfKey` profile since (1) the
828 assertion has issuer info (so you know who originated the assertion came) (2) the signed
829 assertion includes the key for the binding signature.

830 **5.4.4.2. *Sender Vouches***

831 This profile has one or more authorities issuing assertions that contain `<SubjectConfirmation>`
832 elements that basically say “Trust these if you trust the issuer and the entity who signed them”.

833 A collects these assertions and inserts them in a message. The sender then signs over the entire
834 message, with the signature being used to indicate that these assertions (which are themselves
835 signed by their issuers) are vouched for by the sender.

836 **5.4.4.2.1. *Eavesdropping***

837 Eavesdropping continues to be a threat in the same manner as for the SAML SOAP binding, as
838 discussed in Section 5.3.1.1. The routable nature of SOAP adds the potential for a large number
839 of steps and actors in the course of a message’s lifetime, which means that the potential
840 incidences of eavesdropping are increased as the number of possible times a message is in transit
841 increases.

842 The persistent nature of SOAP messages adds an additional possibility of eavesdropping, in that
843 persisted items can be read from their store.

844 To provide maximum protection from eavesdropping, assertions should be encrypted in such a
845 way that only the intended audiences can view the material. This removes threats of
846 eavesdropping in transit, but does not remove risks associated with storage by the receiver or
847 poor handling of the clear text by the receiver.

848 **5.4.4.2.2. *Replay***

849 The fact that the sender does all binding prevents a variety of replay attacks that reuse the
850 assertion with different documents. In this case the assertions are directly signed into the
851 document, so separating them from the document for reuse would not benefit a malicious user.
852 (i.e. The assertions are only as valid as the binding signature of the sender, so reusing them with
853 a different key does not pose a risk).

854 Authorities should note that once a “SenderVouches” assertion has been issued, there is no
855 control over who may use it. Any entity coming into contact with the assertion can separate these
856 assertions and use them by signing them with their own keys. Consumers of SenderVouches
857 assertions must, therefore, carefully decide which senders to allow to vouch for what assertions.

858 The capture and resubmission of the entire message (SAML assertions and business payload) is a
859 threat. One counter-measure is to add information about time, or a sequence number to the
860 digital signature included in the SOAP header. The receiver can use this information to detect
861 duplicate messages.

862 **5.4.4.2.3. Message Insertion**

863 There is no message insertion attack at the level of the `SenderVouches` format of the SOAP
864 profile.

865 **5.4.4.2.4. Message Deletion**

866 There is no message insertion attack at the level of the `SenderVouches` format of the SOAP
867 profile.

868 **5.4.4.2.5. Message Modification**

869 The binding signature should prevent any message modification attacks. Selection of what parts
870 of the document to sign should be made carefully with the possibility of this attack in mind.

871 Receivers should consider only the portions of the document actually bound by signature to the
872 assertions as valid with respect to the assertions.

873 **5.4.4.2.6. Man-in-the-Middle**

874 The requirement for a signature here should prevent MITM attacks. Note that the verifiability of
875 the signature is key to this step: Not only must a receiver be able to verify that a document was
876 signed with a key, but he also needs to be able to verify the binding of key to identity. This may
877 be accomplished by including an X509v3 certificate with the digital signature, which the receiver
878 verifies by some means (XKMS, OCSP, CRLs) and further maps onto a known identity for the
879 signer.

880 If this step is skipped, then MITM becomes a possibility: The MITM captures the original
881 document, alters it, and passes along this new document signed with a key that purports to be
882 from the original sender (but which is actually held by the MITM).

883 The MITM can eavesdrop (if communication is not protected by some confidentiality scheme)
884 but cannot alter the document without detection.

6. References

The following are cited in the text of this document:

- 888 **[Anonymity]** Anonymity, Unobservability, and Pseudonymity -- A Proposal for
889 Terminology
890 Andreas Pfitzmann, Marit Köhntopp
891 http://www.cert.org/IHW2001/terminology_proposal.pdf
- 892 **[FreeHaven]** The Free Haven Project: Distributed Anonymous Storage Service
893 Roger Dingledine & Michael J. Freedman & David Molnar
894 <http://www.freehaven.net/paper/node6.html>
895 <http://www.freehaven.net/paper/node7.html>
- 896 **[HTTTPR]** A Primer for HTTTPR: An overview of the reliable HTTP protocol
897 Stephen Todd, Francis Parr, Michael H. Conner
898 <http://www-106.ibm.com/developerworks/webservices/library/ws-phht/>
- 899 **[IPsec]** IETF IP Security Protocol Working Group,
900 <http://www.ietf.org/html.charters/ipsec-charter.html>.
- 901 **[Pooling]** Pooling Intellectual Capital: Thoughts on Anonymity, Pseudonymity, and
902 Limited Liability in Cyberspace
903 David G. Post
904 <http://www.cli.org/DPost/paper8.htm>
- 905 **[Rescorla-Sec]** E. Rescorla et al., *Guidelines for Writing RFC Text on Security*
906 *Considerations*, <http://www.ietf.org/internet-drafts/draft-rescorla-sec-cons-03.txt>.
- 908 **[RFC2246]** The TLS Protocol Version 1.0, <http://www.ietf.org/rfc/rfc2246.html>.
- 910 **[SAMLBind]** P. Mishra et al., *Bindings and Profiles for the OASIS Security Assertion*
911 *Markup Language (SAML)*, [http://www.oasis-](http://www.oasis-open.org/committees/security/docs/draft-sstc-bindings-model-07.pdf)
912 [open.org/committees/security/docs/draft-sstc-bindings-model-07.pdf](http://www.oasis-open.org/committees/security/docs/draft-sstc-bindings-model-07.pdf),
913 OASIS, December 2001.
- 914 **[SAMLCore]** Hallam-Baker, P. et al., *Assertions and Protocol for the OASIS Security*
915 *Assertion Markup Language (SAML)*, [http://www.oasis-](http://www.oasis-open.org/committees/security/docs/draft-sstc-core-21.pdf)
916 [open.org/committees/security/docs/draft-sstc-core-21.pdf](http://www.oasis-open.org/committees/security/docs/draft-sstc-core-21.pdf), OASIS,
917 December 2001.
- 918 **[SAMLGloss]** J. Hodges et al., *Glossary for the OASIS Security Assertion Markup*
919 *Language (SAML)*, [http://www.oasis-](http://www.oasis-open.org/committees/security/docs/draft-sstc-glossary-02.pdf)
920 [open.org/committees/security/docs/draft-sstc-glossary-02.pdf](http://www.oasis-open.org/committees/security/docs/draft-sstc-glossary-02.pdf), OASIS,
921 December 2001.
- 922 **[SRMPPres]** Message Queuing: Messaging Over The Internet
923 Shai Kariv
924 <http://www.microsoft.com/israel/events/teched/presentations/EN308.zip>

925 **[XMLEnc]** Donald Eastlake et al., *XML Encryption Syntax and Processing*,
926 <http://www.w3.org/TR/xmlenc-core/>, World Wide Web Consortium,
927 October 2001.

928 **[XMLSig]** D. Eastlake et al., *XML-Signature Syntax and Processing*,
929 <http://www.w3.org/TR/xmldsig-core/>, World Wide Web Consortium.

930

931 The following additional documents are recommended reading:

932 **[ebXML-MSS]** Message Service Specification: ebXML Transport, Routing & Packaging
933 Version 1.0 <http://www.ebxml.org/specs/ebMS.pdf>. Chapter 12 is the
934 material of interest.

935

936 **[ebXML-Risk]** ebXML Technical Architecture Risk Assessment v1.0,
937 <http://www.ebxml.org/specs/secRISK.pdf>.

938

939 **[Prudent]** Prudent Engineering Practice for Cryptographic Protocols,
940 <http://citeseer.nj.nec.com/abadi96prudent.html>.

941

942 **[Robustness]** Robustness principles for public key protocols,
943 <http://citeseer.nj.nec.com/2927.html>.

944 **Appendix A. Notices**

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

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

957 Copyright © The Organization for the Advancement of Structured Information Standards
958 [OASIS] 2001. All Rights Reserved.

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

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

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