

This is an attempt to have a consistent story around SML references. (sml:ref="true"; not including sml:keyrefs.) This covers SML Bugzilla bugs 4658, 4673, 4682, 4683, 4780/4795, 4834, 4865, 4884, 4976.

Some of concepts and suggestions are already implemented in the first public working draft of the specifications. They are included here to make this proposal complete. Resolutions to Bugzilla bugs are marked with the corresponding bug number.

1. What's an SML reference?

An element information item in an SML instance document is an SML reference if and only if it has an attribute information item whose {name} is "ref" and whose {namespace} is <sml namespace> and whose {normalized value}, after whitespace normalization using "collapse" following schema rules, is either "true" or "1".

Note: this doesn't depend on having a schema, and will not be affected if a schema is used to produce a PSVI.

Rationale: schemas may not always exist. People need to understand SML documents with SML semantics (just like people can process XML documents without a schema; perform XSLT without a schema.)

2. How are reference elements handled?

If element R is a non-null SML reference, then for reference-handling purposes, R is processed as a unit by SML consumers (who know how to resolve such reference; including their implementations of the deref() function). Note that the above only applies to reference-handling. Often, R may also have non-reference-scheme content and/or attributes (collectively: information items) that may be used in other contexts for other purposes, in which case it doesn't need to be viewed as a unit.

Producers decide which reference scheme(s) to use; in order to be useful, at least one reference scheme used by the producer of an SML-IF document should be understood by each consumer of the document. Each consumer decides which reference scheme(s) to recognize, and this choice **MUST** be consistent between the consumer's implementation of every reference scheme the consumer recognizes and deref() components, which may be separate implementations.

Every consumer must follow a set of rules when processing a reference, but there is a minor difference between the processing required of validators and that required more generally of all consumers. This happens when a single reference R contains multiple schemes recognized by the consumer or the recognized scheme(s) resolve(s) to more than one target. Since a particular consumer may only use some of the SML-defined semantics, it may choose to short-circuit the reference resolution process after finding one scheme it recognizes that resolves successfully; a validator, since it is expected to implement **all** the SML-defined semantics, must evaluate all reference schemes that it recognizes. It is perfectly acceptable for a general consumer to follow exactly the same rules required for a validator, e.g. to facilitate code re-use between validator and non-validating consumers.

The general form for handling non-null references is shown below. Note: this assumes that all attempts to resolve a reference will always return 0, 1 or many targets. That is, the attempt never fails (but resolution may fail) and all model documents are always reachable.

1. If the consumer recognizes no scheme used in the reference, then R is unresolved
2. Assume the consumer recognizes R as using N schemes supported by the consumer, then (answers bug [4976]):
 - 2.1. The validator consumer **MUST** attempt to resolve R using all N schemes, and
 - 2.1.1. If none of the recognized schemes resolves, then R is unresolved.

- 2.1.2. If at least one of the recognized schemes resolves to more than one target element, then the model is invalid.
 - 2.1.3. If one scheme resolves to a target that's different from the target resolved by another scheme, then the model is invalid.
 - 2.1.4. If one scheme resolves and another doesn't, then the model is invalid.
 - 2.1.5. If none of the above is true (that is, all recognized schemes resolve to the same one and only one target element, call it T), then R is resolved to T.
- 2.2. Non-validator consumers are not required to attempt to resolve all N schemes. Their behavior in this case is implementation-defined, but it has to satisfy the following constraints:
- 2.2.1. It either treats R as unresolved, or resolves R to a single target T.
 - 2.2.2. If none of the recognized schemes resolves, then R is unresolved.
 - 2.2.3. If R is treated as unresolved, then at least the condition of one the above 2.1.1~2.1.4 is true for those schemes that are attempted by the consumer.
 - 2.2.4. If R resolves to T, then T MUST be among the target(s) resolved by at least one recognized scheme.

In summary, the consumer has to obey the "at most one target" rule, and it can't lie about the result. The above rules allow a general consumer to behave in (and not limited to) any of the following ways or any combination of them:

- Only use the first scheme it recognizes, and ignore other schemes.
- Try the recognized schemes one by one until one resolves.
- If a scheme or multiple schemes resolve to more than one target, make R unresolved.
- If a scheme or multiple schemes resolve to more than one target, pick one of them as the target for R.

Note: implementation-defined means a consumer MUST clearly *document* its behavior.

Note: 2.2.3 above indicates non-validator consumers are *not* required to attempt to resolve any recognized schemes.

Note: only validators are required to report "multiple target" errors. This means the current "3.1.2.1 At Most One Target" constraint needs to be updated to only cover generic consumers. Validator behavior needs to be moved to a different section.

Any attempt to resolve a reference is limited to the *current* model. That is, target MUST be within the current model.

So any given reference falls into one of 3 categories: null, resolved, and unresolved. (More about "null" in section 4.) Suggest we delete the definition of "dangling reference" and use "unresolved reference" instead, because the former doesn't carry more information than the latter. For unresolved references, it's not always possible to tell whether it's because the reference is wrongly specified or whether the target is outside of the model. (But "dangling" would imply the latter case.)

Implementations of the deref() function should be governed by the same rules as non-validator consumers. A deref() implementation should not report errors (which isn't allowed by XPath 1.0 anyway). For a given reference element, it MAY either return nothing (unresolved) or one target. It MUST NOT return more than 1 target for each reference element. (Answers bug [4683].)

Rationale: the reference is a unit, because we should not constrain how schemes are defined. e.g. maybe my scheme uses multiple <my:uri> elements, and I specify "the reference resolves to the first <my:uri> that actually resolves". Then even though there could be many <my:uri> sub-elements, they are *not* multiple instances of the same scheme and don't constitute an error if more than one resolves.

There is a question about "when no scheme is recognized, should it be viewed as a null or unresolved reference?" The above gives the answer "unresolved", because this case is distinct from "this reference doesn't have a value (null)".

Because all schemes used by the same reference element are meant to target the same element, it should then be an error for cases described in 2.3, which should be reported by validators. But for reasons like performance, general consumers may not want to try to resolve all schemes, hence the implementation-defined behavior.

3. How are schemes defined?

Again because references are units, we don't need to get into details in terms of how new schemes should be defined. In particular:

- Schemes may overlap. That is, the same sub-element/attribute can be used by multiple schemes.
- Schemes can use attributes. (More on this follows.)

The only questions a scheme definition needs to and MUST answer are: (Answers bug [4865].)

- How is a scheme recognized. That is, there is a set of rules that, when satisfied, identify elements with `sml:ref="true"` as instances of this scheme. (e.g. for the "uri/iri" scheme, the rule can include: there is one and only one sub-element named `<sml:uri>`, whose value is `xs:anyURI`.) That is, either a reference element is an instance of (or an example of, or uses) a scheme, or it is not. "How many instances" is never a question.
- How is a scheme de-referenced. That is, for an element R using this scheme, a deterministic approach is defined to resolve R to a set of target element nodes. This MUST cover all the reference elements recognized as an instance of this scheme.

Because rules are specified in terms of the reference elements, we don't need to deal with "what if there are 2 instances of the same scheme in a single reference R". (Answers bug [4658].)

We may want to add a note to suggest that scheme definitions SHOULD NOT use overlapping elements/attributes without good reasons.

Rationale: again, we don't want to constrain how people want to define their schemes, especially for documents that pre-date SML. All we can require is that given an SML reference element R (with `sml:ref="true"`), a scheme must be defined in a way that a) consumers can tell whether R uses such a scheme, and b) if so how to resolve R using the scheme to a set of nodes. (The set may contain more than 1 node, which will result in an error because SML references must resolve to a single node, but no one else cares about why there were more than 1.)

4. How are null references identified and handled?

"Null" references are distinct from unresolved references. A null reference is an explicit declaration of intent by the producer that the reference itself does not exist, and a processing directive (NOT a hint) to consumers not to search the reference for reference scheme information items.

Propose to introduce another attribute, say, `sml:nilref`, with a boolean value (following the same rules as `sml:ref`). This attribute is only allowed when `sml:ref="true"`; if `sml:nilref`'s value is "true", then the reference is viewed as "null". This is the only case where a reference element denotes a "null" reference. (Answers bug [4884].) In the SML schema, define a global `sml:nilref` attribute with `xs:boolean` type.

When an reference element is recognized as "null", then consumers MUST not attempt to resolve it. Any scheme-related content (including attributes) recognized by the consumer is ignored for reference-related purposes. The question of whether a null reference is resolved or not is undefined; it is an ill-formed question.

"targetRequired" should also apply to "null" references. That is, it's an error if `targetRequired=true` is specified and the corresponding reference element R has `sml:nilref=true`. (Answers bug [4780/4795].)

Rationale: Null references are different from empty references (e.g. content of the reference element is intentionally empty; this is an explicit declaration of intent from the producer). Empty references are just normal references (whose value happens to be empty).

Using the university example, if each course has a “primary teacher” and an “assistant teacher”, wouldn’t it be nice for the consumer/application to know “this course doesn’t have an assistant teacher” rather than “I (the consumer) don’t understand the way you (the producer) specified the assistant teacher, and I don’t know who (s)he is”?

As for recognizing “null” references, `xsi:nil` doesn't work unless we constrain scheme definitions to not use attributes (i.e., to use element content only). Also when schema validation is used, `xsi:nil` disallows any sub-element or textual content; but there may be cases where a null reference still has non-reference-related content.

This also answers the question "do we support schemes involving attributes"? Because whether a reference is null or not is controlled by `sml:nilref`, it can cover attributes well (whereas `xsi:nil` can't). (Answers bug [4682].)

“Null” should not be a special case to “targetRequired”. If we view “nilref” as similar to “xsi:nil” in schema, then “targetRequired=true” is like the opposite of “nillable”.

5. What's an SML reference type in schema?

There are 5 classes of (complex) types with regard to how they interact with the `sml:ref` attribute.

1. Those who are derived from the current `sml:refType`.
2. Those who have an explicit reference to the global attribute declaration with `use=required` and `fixed=true`
3. Those who have an explicit reference to the global attribute declaration, with either true or false value
4. Those who allow `sml:ref` via an attribute wildcard or an explicit reference to the global attribute declaration, with either true or false value
5. Any complex type / any element declaration, including those who don't allow `sml:ref` with a “true” value

SML processors or applications may choose to do different things for these different classes (e.g. static analysis, GUI-based tooling, etc.).

There is no need to define “SML reference type”. The only thing SML needs to define is, for which class(s) of types are `sml:acyclic` and `sml:targetXXX` allowed to be specified.

Discussed during 2007-08 F2F. The WG decided to choose class #5.

(Similar to schema disallowing defining new `xsi:` attributes, see [1]) SML can disallow defining `sml:` attributes (by user schemas). (This removes the possibility for a schema to define a local `sml:ref`, which would confuse consumers.)

[1] <http://www.w3.org/TR/xmlschema-1/#no-xsi>

Suggest to remove the special `sml:refType` (as it has no special meaning any more).

Rationale:

Class 1 makes it impossible to use existing schemas and make existing complex type reference types (because schema doesn't support multiple-inheritance). Classes 2~4 all make it hard for the scenario

where there is a base (possibly abstract) type with SML constraints specified, and sml:ref attribute is only allowed by types derived from this base type.

6. How to handle non-references

When an element is not a reference element (i.e. doesn't have sml:ref=true), how should it be handled in reference-related cases? There are 3 scenarios:

1. targetXXX constraints

Discussed during 2007-08 F2F; WG decided to make them trivially "satisfied" for non-references. That is, they are enforced when and only when the elements have sml:ref=true.

2. deref() (e.g. in identity constraints)

Should those non-refs be passed to deref()? Or an error should be issue and deref() is not called?

During the 2007-08 F2F, WG decided to answer "no error; deref() returns unresolved".

3. Schema default value for sml:ref

Then it's possible that in the input infoset there is no sml:ref, but in PSVI there is sml:ref=true. Should it be viewed as a reference element?

During the 2007-08 F2F, some members were worried about potential inter-op problems between performing and not performing schema assessment; others are not as worried, because such problem already exists between processors who support different reference schemes. Some other members also see the value in viewing defaulted sml:ref=true as references. This makes it possible to transform existing XML data into SML models, without having to modify all the instances and add sml:ref attributes to them.

This issue was **not** resolved during the F2F.

Rationale: to expand a little on item #1. Each element in the instance (whose corresponding element declaration has any one of the above constraints) can be in one of 4 categories:

- non-reference
- null reference
- unresolved reference
- resolved reference

And for the 4 constraints, we have the following 4 x 4 matrix applies:

Reference\Constraint	Acyclic	targetRequired	targetElement	targetType
Non-reference	Satisfied			
Null	Satisfied	Violated	Violated	Violated
Unresolved	Satisfied	Violated	Violated	Violated
Resolved	Check	Satisfied	Check	Check

There are 2 options for the "non-reference" row for the targetXXX constraints:

- a) "violated", because targetXXX means "it must be a reference and its target must ..."
- b) "satisfied", because targetXXX means "if it is a refernce then its target must ...".

WG picked answer (b), because for those who want #1 answer, they could have had use=required and fixed=true.