NIST Interagency Report 7694

National Institute of
Standards and Technology
U.S. Department of Commerce

Specification for the Asset Reporting Format 1.1

Adam Halbardier
David Waltermire
Mark Johnson

NIST Interagency Report 7694

Specification for the Asset Reporting Format 1.1

Adam Halbardier
David Waltermire
Mark Johnson

C O M P U T E R S E C U R I T Y

Computer Security Division
Information Technology Laboratory
National Institute of Standards and Technology
Gaithersburg, MD 20899-8930

June 2011

U.S. Department of Commerce

Gary Locke, Secretary

National Institute of Standards and Technology

Dr. Patrick D. Gallagher, Director

Reports on Computer Systems Technology

The Information Technology Laboratory (ITL) at the National Institute of Standards and Technology (NIST) promotes the U.S. economy and public welfare by providing technical leadership for the nation's measurement and standards infrastructure. ITL develops tests, test methods, reference data, proof of concept implementations, and technical analysis to advance the development and productive use of information technology. ITL's responsibilities include the development of technical, physical, administrative, and management standards and guidelines for the cost-effective security and privacy of sensitive unclassified information in Federal computer systems. This Interagency Report discusses ITL's research, guidance, and outreach efforts in computer security and its collaborative activities with industry, government, and academic organizations.

National Institute of Standards and Technology Interagency Report 7694
29 pages (June 2011)

Acknowledgments

The authors, Adam Halbardier of Booz Allen Hamilton, David Waltermire of the National Institute of Standards and Technology (NIST) and Mark Johnson of Booz Allen Hamilton, wish to thank their colleagues who reviewed drafts of this document and contributed to its technical content. The authors would like to acknowledge Paul Cichonski and Harold Booth of NIST, John Wunder of the MITRE Corporation, Karen Scarfone of Scarfone Cybersecurity, Joseph Wolfkiel of the Defense Information Systems Agency (DISA), Jim Ronayne of Varen Technologies, Gary Newman of Belarc, and Rhonda Farrell of Booz Allen Hamilton for their keen and insightful assistance throughout the development of this document.

Abstract

This specification describes the Asset Reporting Format (ARF), a data model for expressing the transport format of information about assets and the relationships between assets and reports. The standardized data model facilitates the reporting, correlating, and fusing of asset information throughout and between organizations. ARF is vendor and technology neutral, flexible, and suited for a wide variety of reporting applications. The intent of ARF is to provide a uniform foundation for the expression of reporting results, fostering more widespread application of sound IT management practices. ARF can be used for any type of asset, not just IT assets.

Trademark Information

Table of Contents

List of Figures and Tables

1. Introduction

The Asset Reporting Format (ARF) is a data model to express the transport format of information about assets and the relationships between assets and reports. The standardized data model facilitates the reporting, correlating, and fusing of asset information throughout and between organizations. ARF is vendor and technology neutral, flexible, and suited for a wide variety of reporting applications.

ARF serves an important role in assisting organizations with sharing information about assets, both internally and externally. Data about assets often exists across many locations within an organization (e.g., databases, sensors). In addition, different groups within the organization often have different views of the information, and people within those groups often capture only the information pertinent to them. ARF helps enable organizations to quickly correlate the information from those disparate data sources, resulting in a more holistic view of an asset. Asset identification, defined in [Asset Identification], supports asset correlation, which allows organizations to match different representations of the same asset. By tying together asset identification, reports, and report requests, different ARF reports can be correlated by the asset identification component, and the reports can be fused. Conforming to ARF enhances an organization's ability to increase its situational awareness related to its asset infrastructure through better correlation and fusion across the enterprise.

In addition, ARF standardizes the way in which products produce and receive reports relating to assets, providing a common structure for asset reporting capabilities. Products that conform to ARF support a certain level of data interoperability, regardless of the domain the product is supporting, which better enables cross-product automation processing. ARF provides a common structure for relating reports about assets to assets, report requests, and other reports, leveraging the capabilities of other specifications to define the low-level data elements of the report. Much like the TCP/IP protocol that focuses on defining the structure of the message while leaving the detailed format of the payload to other specifications (e.g., HTTP, FTP, Java Messaging Service), ARF does not define the detailed data elements of the message. As such, ARF is intended to be free of any specific use case.

Since ARF exists at a level above the detailed data formats such as the eXtensible Configuration Checklist Description Format (XCCDF) and the Open Vulnerability and Assessment Language (OVAL), products from various domains, supporting various use cases, can conform to ARF while continuing to support their domains. Those products, which were once incompatible, would then share a level of interoperability. In this way, ARF cuts across domains within an enterprise, enabling data interoperability across disparate products.

1.1 Purpose and Scope

The purpose of this document is to define the Asset Reporting Format (ARF) specification, a data model to express the transport format of information about assets and the relationships between assets and reports. The scope of this document is to give an introduction to ARF version 1.1, define ARF's data model, and document the conformance requirements to comply with ARF. Other versions of ARF and the associated component specifications, including emerging specifications and future versions, are not addressed here. Future versions of ARF will be defined in distinct revisions of this document, each clearly labeled with a document revision number and the appropriate ARF version number.

1.2 Audience

The intended audience for this specification includes product vendors who are developing asset reporting applications, particularly those that will interoperate with other asset reporting solutions, as well as

organizations interested in building and tailoring customized, interoperable capabilities for asset reporting.

1.3 Document Structure

The remainder of this document is organized into the following major sections:

- Section 2 defines the terms used within this specification and provides a list of common abbreviations.
- Section 3 describes how this specification relates to other standards and specifications.
- Section 4 defines the conformance requirements for ARF.
- Section 5 provides an overview of the ARF data model constructs.
- Section 6 defines the relationships described in ARF.
- Section 7 defines the mechanism to link a report to an ARF object.
- Appendix A defines possible use-cases for ARF.
- Appendix B documents references to documents referenced normatively in ARF.
- Appendix C describes a sample workflow using ARF.
- Appendix D provides a sample XML file of an ARF report.

1.4 Document Conventions

Throughout this specification, whenever a specific term from the data model is referenced, as defined in section 5, the term is written in `Courier New` font. When referencing a specification listed in Appendix B, the name will be written between brackets, such as [Asset Identification].

The key words “MUST”, “MUST NOT”, “REQUIRED”, “SHALL”, “SHALL NOT”, “SHOULD”, “SHOULD NOT”, “RECOMMENDED”, “MAY”, and “OPTIONAL” in this document are to be interpreted as described in [RFC 2119].

Both inline and indented forms use qualified names to refer to specific XML elements. A qualified name associates a named element with a namespace. The namespace identifies the specific XML schema that defines (and consequently may be used to validate) the syntax of the element instance. A qualified name declares this schema to element association using the format ‘*prefix*:*element-name*’. The association of prefix to namespace is defined in the metadata of an XML document and generally will vary from document to document. In this specification, the conventional mappings listed in Table 1-1 are used.

Table 1-1: Conventional XML Mappings

Mappings Prefix	Namespace URI	Schema
ai	http://scap.nist.gov/schema/asset-identification/1.1	Asset Identification 1.1
arf	http://scap.nist.gov/schema/asset-reporting-format/1.1	Asset Reporting Format 1.1
core	http://scap.nist.gov/schema/reporting-core/1.1	SCAP Reporting Core 1.1
xlink	http://www.w3.org/1999/xlink	XLink
dsig	http://www.w3.org/2000/09/xmldsig#	XML Signature

2. Terms and Abbreviations

This section defines a set of common terms and abbreviations used within this specification.

2.1 Terms

Asset: Anything that has value to an organization, including, but not limited to, another organization, person, computing device, information technology (IT) system, IT network, IT circuit, software (both an installed instance and a physical instance), virtual computing platform (common in cloud and virtualized computing), and related hardware (e.g., locks, cabinets, keyboards).

Asset Identification: The attributes and methods necessary for uniquely identifying a given asset. A full explanation of asset identification is provided in [Asset Identification].

Asset Report: A collection of content (or link to content) about an asset.

Asset Report Request: A collection of structured information used as input to generate an asset report. An asset report request may be of any format and may have different contexts depending on the nature of the request. For instance, the request may be written in a control language that dictates how the request is to be propagated and executed. The request may also be written as a formal definition without reference to how the request is to be executed. The request may also be a prose description that must be interpreted and executed by a person. These examples are not exhaustive.

Asset Reporting Format (ARF) Report: The collection of all assets, report requests, reports, and relationships for a given instance of ARF.

2.2 Acronyms

ARF	Asset Reporting Format
DISA	Defense Information Systems Agency
DoD	Department of Defense
FISMA	Federal Information Security Management Act
FTP	File Transport Protocol
HIPAA	Health Insurance Portability and Accountability Act
HTTP	Hypertext Transport Protocol
IETF	Internet Engineering Task Force
IR	Interagency Report
IT	Information Technology
ITL	Information Technology Laboratory
NIST	National Institute of Standards and Technology
NSA	National Security Agency
OCIL	Open Checklist Interactive Language
OVAL	Open Vulnerability and Assessment Language
PCI	Payment Card Industry
RFC	Request for Comment
SCAP	Security Content Automation Protocol
SCG	Security Configuration Guide
SOX	Sarbanes-Oxley
STIG	Security Technical Implementation Guide
TCP/IP	Transmission Control Protocol/Internet Protocol
URI	Universal Resource Identifier

URL	Universal Resource Locator
W3C	World Wide Web Consortium
XCCDF	eXtensible Configuration Checklist Description Format
XLink	XML Linking Language
XML	Extensible Markup Language
XSD	XML Schema

3. Relationship to Existing Standards and Specifications

ARF's relationships to other selected specifications are described below.

1. Asset Identification – ARF leverages [Asset Identification] to identify assets for the ARF report. Asset identification is a critical component of ARF, as it enables the correlation and fusion of various ARF reports and report content. ARF provides a place to house assets defined using [Asset Identification], a place to house information about assets (asset reports), and a place to define relationships between those assets and asset reports. Therefore, [Asset Identification] is a crucial component in tying different asset reports together, because the common component is the asset identified using [Asset Identification].

2. Report Content and Report Request – While ARF defines the transport format and message structure for a report, the structure for the detailed data elements is left to other, lower-level specifications. The lower-level specifications are not enumerated here as the list is open-ended.

4. Conformance

Developers and organizations may want to build products in conformance with ARF so that users of those products, and consumers of the content generated by those products, have a guarantee about the format of the data that the product will produce and can consume. In addition, products that conform to this specification will be better able to interoperate and exchange reporting information with other products that conform to ARF.

Organizations may want to claim conformance with this specification to increase the ease of exchanging information with other organizations that also conform to this specification. In addition, an organization that conforms to this standard across its enterprise is better able to correlate and fuse data related to assets across its various domains, creating a more holistic view of its IT enterprise.

The following sections define the criteria for content and products to claim conformance with this specification.

4.1 Content Conformance

In order for an ARF report to be considered in compliance with this specification, the report MUST adhere to the following requirements:

1. The ARF report SHALL conform with all of the normative guidance provided in Section 5.
2. When identifying assets, the ARF report SHALL identify assets in a manner that is consistent with [Asset Identification].

4.2 Product Conformance

A product may claim conformance with this specification when it produces information about assets in a format that is consistent with this specification. To that end, a product claiming conformance to this specification MUST adhere to the following requirements:

1. If the product produces reports about assets, the product SHALL produce ARF reports that are compliant with this specification per Section 4.1.
2. If the product consumes reports about assets, the product SHALL consume ARF reports that are compliant with this specification per Section 4.1.
3. If the product consumes reports about assets, the product SHALL support the relationships defined in Section 6.

5. Data Model

This section documents the data model and data model requirements for ARF. A sample XML document in the format of this data model is located in Appendix D.

Figure 5-1 illustrates the organizational concept of ARF and the composition of its components.

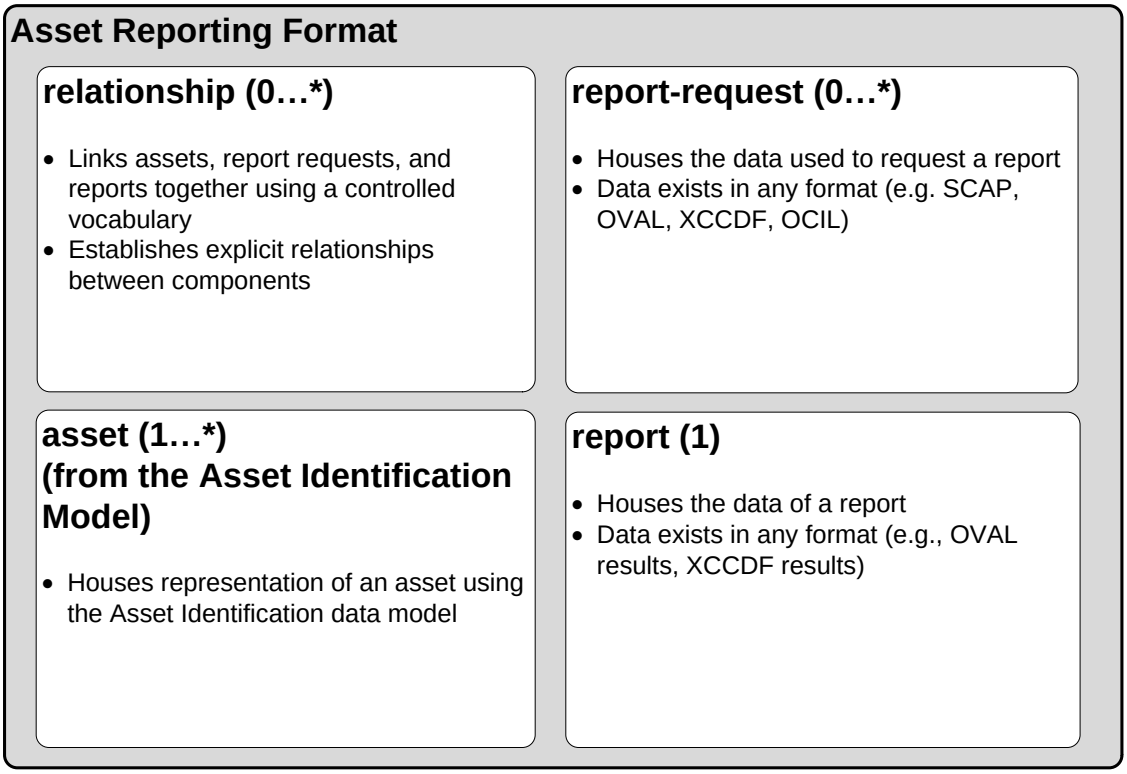

Figure 5-1: Notional Asset Reporting Format Model

For the purposes of this section, the term "content" can be understood to be the report data captured in the report section itself, or external data provided via a link. In addition, for the purposes of this section, the term "object" refers to the union of report requests, assets, and reports defined in an ARF report.

The intention of the ARF data model is to capture and encapsulate information about assets in a structured manner. The format is broken into four main sections.

1. The asset section includes asset identification information for one or more assets. The asset section simply houses assets independent of their relationships to reports. The relationship section can then link the report section to specific assets.
2. The report section contains one or more asset reports. An asset report is composed of content (or a link to content) about one or more assets.
3. The report-request section contains the asset report requests, which can give context to asset reports captured in the report section. The report-request section simply houses asset report requests independent of the report which was subsequently generated.
4. The relationship section links assets, reports, and report requests together with well-defined relationships. Each relationship is defined as {subject} {predicate} {object}, where {subject} is

the asset, report request, or report of interest, {predicate} is the relationship type being established, and {object} is one or more assets, report requests, or reports.

In order to comply with the ARF data model,

- The user MUST produce an XML `arf:asset-report-collection` element consistent with the data model described below.
- The XML element produced MUST validate against the XML schema for Asset Reporting Format 1.1 listed at http://scap.nist.gov/specifications/arf/index.html. In situations where the XML schema does not match the documented model in this specification, the XML schema takes precedence.

The following tables formalize the data model. The data contained in the tables are requirements and MUST be interpreted as follows:

- The "Element Name" field indicates the name for the XML element being described. Each element name has a namespace prefix indicating the namespace to which the element belongs. See Table 1-1 for a mapping of namespace prefixes to namespaces.
- The "Definition" field indicates the prose description of the element. The definition field MAY contain requirement words as indicated in [RFC 2119].
- The "Properties" field is broken into four subfields:
  - The "Name" column indicates the name of a property that MAY or MUST be included in the described element, in accordance with the cardinality indicated in the "Count" field
  - The "Type" column indicates the REQUIRED data type for the value of the property. There are two categories of types: literal and element. A literal type will indicate the type of literal as defined in [XML Schema]. An element type will reference the name of another element that ultimately defines that property.
  - The "Count" column indicates the cardinality of the property within the element. The property MUST be included in the element in accordance with the cardinality. If a range is given, and "n" is the upper-bound of the range, then the upper limit is unbounded.
  - The "Definition" column defines the property in the context of the element. The definition MAY contain requirement words as indicated in [RFC 2119].

Table 5-1: Element – arf:asset-report-collection

Element Name: arf:asset-report-collection				
Definition	The top-level container element that holds all of the information in an ARF report. This element need not be the root element of the XML document, as an ARF report MAY be included inside another document or as part of another report, but it SHALL be the root element of an ARF report.			
Properties	**Name**	**Type**	**Count**	**Definition**
	id	literal – NCName	0-1	The id for the collection of asset reports.
	report-requests	element – arf:report-requests	0-1	This is a container element that, when it exists in the arf:asset-report-collection, holds one or more report-request elements.
	assets	element – arf:assets	0-1	This is a container element that, when it exists in the arf:asset-report-collection, holds one or more asset elements.
	reports	element – arf:reports	1	This is a container element that contains one or more report elements.

	relationships	element – core:relationships	0-1	This is a container element that, when it exists in the arf:asset-report-collection, holds one or more relationship elements.
	extended-infos	element – ai:extended-infos	0-1	May be used as an extension point.

Table 5-2: Element – arf:report-requests

Element Name: arf:report-requests				
Definition	Contains a collection of report-request elements.			
Properties	**Name**	**Type**	**Count**	**Definition**
	report-request	element – arf:report-request	1-n	Contains an asset report request for at least one of the reports on this ARF report.

Table 5-3: Element – arf:report-request

Element Name: arf:report-request				
Definition	Contains a report request for at least one of the reports on this ARF report collection.			
Properties	**Name**	**Type**	**Count**	**Definition**
	content	element – any XML	Only 1 of content or remote-resource	The content element requires the inclusion of embedded XML content that MUST be in a namespace other than the ARF namespace and SHALL be used to represent a report request in a third party model. "any XML" is not defined elsewhere in this document because it can be any XML element. One, and only one, of either this element or the remote-resource element MUST be provided on each report request.
	remote-resource	element – arf:remote-resource		A link to the data of this report request. This element SHOULD be used when the report request data is external to this ARF report collection. The link MUST resolve to an XML element representing a report request. One, and only one, of either this element or the content element MUST be provided on each report request.
	id	literal - NCName	1	An ID that MUST be unique among all IDs within this asset-report-collection. This ID MUST be referenced at least once in a relationship in this ARF report collection.

Table 5-4: Element – arf:assets

Element Name: arf:assets				
Definition	Contains a collection of asset elements.			
Properties	**Name**	**Type**	**Count**	**Definition**
	asset	element – arf:asset	1-n	Contains an arf:asset element.

Table 5-5: Element – arf:asset

Element Name: arf:asset				
Definition	Contains an ai:asset that represents an asset to be identified.			
Properties	**Name**	**Type**	**Count**	**Definition**
	asset	element – ai:asset	Only 1 of asset or remote-resource	An ai:asset that represents an asset to be identified. One, and only one, of either this element or the remote-resource element MUST be provided on each arf:asset.
	remote-resource	element – arf:remote-resource		A link to an ai:asset element. This element SHOULD be used when the asset identifying information is external to this ARF report collection. The link MUST resolve to an XML element representing an ai:asset. One, and only one, of either this element or the asset element MUST be provided on each arf:asset.
	id	literal – NCName	1	An ID that MUST be unique among all IDs within this asset-report-collection. This ID MUST be referenced at least once in a relationship in this ARF report.

Table 5-6: Element – arf:reports

Element Name: arf:reports				
Definition	Contains a collection of report elements.			
Properties	**Name**	**Type**	**Count**	**Definition**
	report	element – arf:report	1-n	Contains the content of a report.

Table 5-7: Element – arf:report

Element Name: arf:report				
Definition	Contains the content of a report.			
Properties	**Name**	**Type**	**Count**	**Definition**
	content	element – any XML	Only 1 of content or remote-resource	The content element requires the inclusion of embedded XML content that must be in a namespace other than the ARF namespace and is used to represent a report in a third party model. "any XML" is not defined elsewhere in this document because it can be any XML element. One, and only one, of either this element or the remote-resource element MUST be provided on each report.
	remote-resource	element – arf:remote-resource		A link to the data of this report. This element SHOULD be used when the report data is external to this ARF report collection. The link MUST resolve to an XML element representing an asset report. One, and only one, of either this element or the content element MUST be provided on each report.

	id	literal - NCName	1	An ID that MUST be unique among all IDs within this asset-report-collection. This ID MUST be referenced at least once in a relationship.

Table 5-8: Element – core:relationships

Element Name: core:relationships				
Definition	Contains a collection of relationships between the report content and assets, report requests, and other reports.			
Properties	**Name**	**Type**	**Count**	**Definition**
	relationship	element – core:relationship	1-n	Contains a relationship between a subject and object(s) assets.

Table 5-9: Element – core:relationship

Element Name: core:relationship				
Definition	Contains a relationship between the subject and object(s) assets.			
Properties	**Name**	**Type**	**Count**	**Definition**
	ref	literal - NCName	1-n	This element MUST identify the object of this relationship by specifying the ID of an element in this asset-report-collection element. Depending on the type of relationship being asserted, there MAY be additional restrictions on which types of objects may be referenced, but that will be documented with the vocabulary term (see Section 6.1).
	type	literal - QName	1	The type of relationship that is being specified. The QName MUST refer to a term in a controlled vocabulary. The controlled vocabulary is identified by the namespace URI of the QName, and the term in that controlled vocabulary is specified by the local name of the QName. It is helpful, though not required, that when the namespace URI and local name are concatenated, the resulting URI is dereferenceable and points to a location that defines the term. See Section 6 for additional details.
	scope	literal - token	0-1	Determines how to interpret multiple ref elements in the relationship. This element MUST contain the string "inclusive" or "exclusive". When this element is not provided, its default value is "inclusive". When "inclusive" is specified, this relationship should be understood to exist between the report data of this report, and the collection of objects identified by the ref elements on this relationship. When "exclusive" is specified, this relationship should be understood to exist between the report data and each object identified by the ref elements individually. See Section 6.2 for more information.
	subject	literal – NCName	1	This property MUST identify the subject of the relationship by specifying the ID of the asset, report request, or report. Depending on the type of relationship being asserted, there may be additional restrictions on which type of object may be referenced, but that will be documented with the vocabulary term.

Table 5-10: Element – arf:extended-infos

Element Name: arf:extended-infos				
Definition	Contains a collection of extended-info elements.			
Properties	**Name**	**Type**	**Count**	**Definition**
	extended-info	element – arf:extended-info	1-n	An extension point that allows content creators to store information that does not fit elsewhere in the ARF data model.

Table 5-11: Element – arf:extended-info

Element Name: arf:extended-info				
Definition	Holds information that does not fit elsewhere in ARF. The contents of this element can be any XML in a namespace other than the ARF namespace.			
Properties	**Name**	**Type**	**Count**	**Definition**
	id	literal – xs:NCName	1	An ID that MUST be unique among all IDs within this asset-report-collection.

Table 5-12: Element – arf:remote-resource

Element Name: arf:remote-resource				
Definition	Contains a link to external content.			
Properties	**Name**	**Type**	**Count**	**Definition**
	xlink:type	literal – token	1	Value MUST be fixed to “simple”, indicating that this is a simple XLink element.
	xlink:href	literal – any URI	1	A URI indicating the remote content. Producers MUST ensure that consumers of the content will be able to resolve the URI.

6. Relationships

The ARF data model allows for explicit relationships to be defined between objects in the ARF report.

6.1 Relationship Terms

Each relationship is defined as {subject} {predicate} {object}, where {subject} is the principal asset, report request, or report of interest, {predicate} is the relationship type being established, and {object} is one or more of the following: assets, report requests, or reports. As defined in the data model, the predicate (i.e., the value of the type field on the relationship element) MUST be a qualified name that refers to a term in a controlled vocabulary. Specified below are terms defined in a controlled vocabulary for ARF. It is not required that content producers use the terms defined below, but all ARF compliant implementations MUST be capable of processing all of the terms defined in this section. Content producers SHOULD use these terms when possible. Content producers MAY use terms defined in other vocabularies, but ARF compliant tools are not required to understand any relationships beyond the ones listed in this section.

All terms listed in Table 6-1exist in the controlled vocabulary identified by http://scap.nist.gov/specifications/arf/vocabulary/relationships/1.0#. The definition of each term can also be found at the URL created when concatenating the URL and the term together. The table MUST be interpreted as follows:

- The "Term" column indicates the local-name in the QName of the term being identified.
- The "Domain" column indicates the exhaustive set of subject types that MAY be referenced by a relationship of that type. A relationship of that type MUST reference a subject of the type indicated in "Domain" for that relationship.
- The "Range" column indicates the exhaustive set of object types that MAY be referenced by a relationship of that type. A relationship of that type MUST reference an object of the type indicated in "Range" for that relationship.
- The "Description" column contains a prose description of the relationship type. This column MAY contain requirement words as indicated in [RFC 2119].

Table 6-1: Controlled Vocabulary Defined for ARF

Term	Domain	Range	Description
isAbout	arf:report	ai:asset	The data in the report is about the asset.
retrievedFrom	arf:report	ai:asset	The data in the report was retrieved from the asset. This relationship will generally be used when the asset identifies some data store that houses information. This relationship indicates that the data came from the asset, but the asset did not create, or in any other way produce, the data, other than to supply the stored data.
createdBy	arf:report	ai:asset	The data in the report was created by the asset. This relationship SHOULD refer to the tool that created the content originally.
hasSource	arf:report	ai:asset	The report contains knowledge from the asset. This relationship refers to the asset that supplied the knowledge to create the report content. This relationship implies that the asset is authoritative about the information.

Term	Domain	Range	Description
recordedBy	arf:report	ai:asset	The information in the report was recorded by the asset. This relationship will usually be used when the report content is data about a digital event that is captured by an asset (e.g., a piece of software fires a digital event).
initiatedBy	arf:report	ai:asset	The information in the report was initiated by the asset. This relationship will usually be used when the report content represents a digital event and that digital event is initiated by the asset.
createdFor	arf:report	ai:report-request	The report was created because of the report request. This relationship will usually be used to associate request and response type data.
hasMetadata	arf:report	ai:report	The subject report has additional metadata that is represented in the object report.

Content producers who choose not to use the terms listed in Table 6-1, or to use other terms in addition to those listed in Table 6-1, MAY do so while still remaining compliant to this specification. In that case, though, the content producers SHALL use terms defined in a controlled vocabulary. The controlled vocabulary SHALL be identified using a URI. Concatenating the controlled vocabulary URI with a term in the vocabulary MAY create a dereferenceable URI that points to a definition for that term. This is often accomplished by using an HTTP URL for the controlled vocabulary URI, and ending that URL in "#" or "/". For instance, http://scap.nist.gov/specifications/arf/vocabulary/relationships/1.0#isAbout is a dereferenceable link to the definition of "isAbout".

6.2 Relationship Scope

In order to provide an additional level of robustness, the ARF data model allows relationships to be defined in two ways: inclusive and exclusive. When a relationship is defined in ARF, it may be defined between a single subject and multiple objects. The way the relationship is interpreted with respect to those objects is described by the scope of the relationship. An inclusive relationship means that the relationship is defined and understood to be between the subject and the collection of objects. An exclusive relationship means that multiple relationships are established with one relationship between the subject and each object.

This XML fragment shows an inclusive relationship:

```
<rc:relationship type="rel:isAbout" subject="subject1" scope="inclusive">
    <rc:ref>object1</rc:ref>
    <rc:ref>object2</rc:ref>
    <rc:ref>object3</rc:ref>
</rc:relationship>
```

Figure 6-1 illustrates the relationship defined in the XML fragment above.

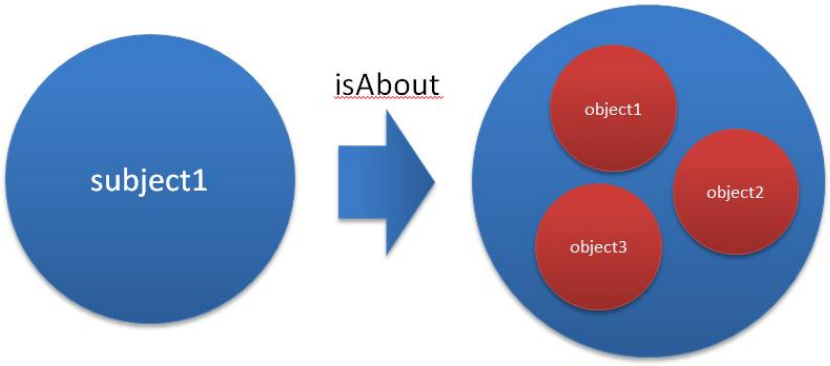

Figure 6-1: Demonstrate Inclusive Relationship

The example in Figure 6-1 illustrates that subject1 isAbout the collection of objects 1, 2, and 3.

This XML fragment shows an exclusive relationship:

```
<rc:relationship type="rel:isAbout" subject="subject1" scope="exclusive">
    <rc:ref>target1</rc:ref>
    <rc:ref>target2</rc:ref>
    <rc:ref>target3</rc:ref>
</rc:relationship>
```

Figure 6-2 illustrates the relationship defined in the XML fragment above.

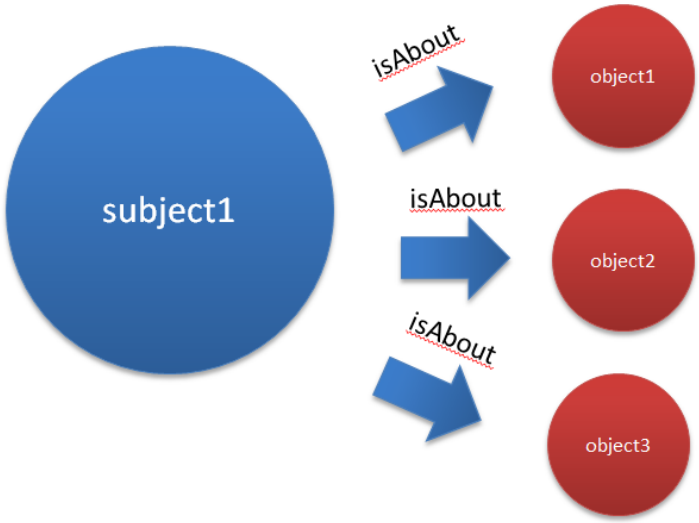

Figure 6-2: Demonstrate Exclusive Relationship

The example in Figure 6-2 illustrates that subject1 isAbout object1, subject1 isAbout object2, and subject1 isAbout object3 individually. Unlike the inclusive relationship, the exclusive relationship establishes 3 separate relationships.

7. Referencing ARF Objects

The data model is built on a hierarchical design where the asset-report-collection holds the reports, assets, and report requests, as well as the relationships linking those objects together. The details of the reports, assets, and report requests are defined by other, lower-level specifications. This hierarchical design is intended to create a logical division between ARF and the lower-level specifications. One advantage is that report content can be created without considering if the report will be included in a larger ARF report. This gives content creators the freedom to develop a report independent of the final report packaging.

While the design is intended to decouple reports from the ARF wrapper, an element is included in the ARF schema that allows a level of coupling if a report creator desires such a feature. The ARF schema includes an `<arf:object-ref>` element, as documented below. A content creator MAY include the element in a report. The element references an object (i.e., another report, asset, or report request) by the ID assigned to that object in the asset report collection. The semantics of the reference MUST be defined by the data model defining the report. Table 7-1 defines the `<arf:object-ref>` element.

Table 7-1: Element – arf:object-ref

Element Name: arf:object-ref				
Definition	Defines a reference to an ARF object (e.g. a report, asset, or report request)			
Properties	**Name**	**Type**	**Count**	**Definition**
	ref-id	literal – token	1	Value MUST be the ID of an asset, report, or report request ARF object located in the same <arf:asset-report-collection> ancestor as this element.

The following XML fragment demonstrates the use of arf:object-ref.

```
<arf:asset-report-collection>
    ...
    <arf:assets>
        <arf:asset id="asset1">
            <ai:computing-device>
                ...
            </ai:computing-device>
        </arf:asset>
    </arf:assets>
    <arf:reports>
        <arf:report id="report1">
            <arf:content>
                <sample-report xmlns="http://tempuri.org">
                    <report-is-about>
                        <arf:object-ref ref-id="asset1"/>
                    </report-is-about>
                    ...
                </sample-report>
            </arf:content>
        </arf:report>
    </arf:reports>
</arf:asset-report-collection>
```

In the XML fragment above, report1 contains a sample-report. The sample-report defines an element `<report-is-about>` that uses the object-ref element to refer to the asset1 element in ARF. The semantics of the report-is-about element are defined by the sample report schema. The arf:object-ref element functions as a pointer to the asset1 element so that the asset1 data does not need to be duplicated in the sample-report.

Appendix A—Use Cases

The following use cases are related to ARF. As such, the requirements needed to fulfill each use case have been taken into consideration when deciding which components should be included in this specification. The following list of use cases is not exhaustive.

Asset Discovery and Inventory Management

It is often important for an organization to track inventory of assets it owns and operates. Inventory management may be a requirement imposed upon an organization by another entity or it may be necessary for internal compliance and auditing. In any case, it is essential that an organization have the ability to discover and track assets, along with critical information related to those assets. ARF supports this use case by providing a container to identify the assets (i.e., using the `assets` element) discovered and hold information about those assets.

Vulnerability Management

Vulnerability management and continuous monitoring are important topics in organizations, especially as they push for better visibility into their security postures. ARF provides a format that allows organizations to report on vulnerability related information about assets at any level of abstraction, whether it is about a specific asset, a group of assets, or an entire organization. Figure A-1 shows a sample vulnerability management workflow.

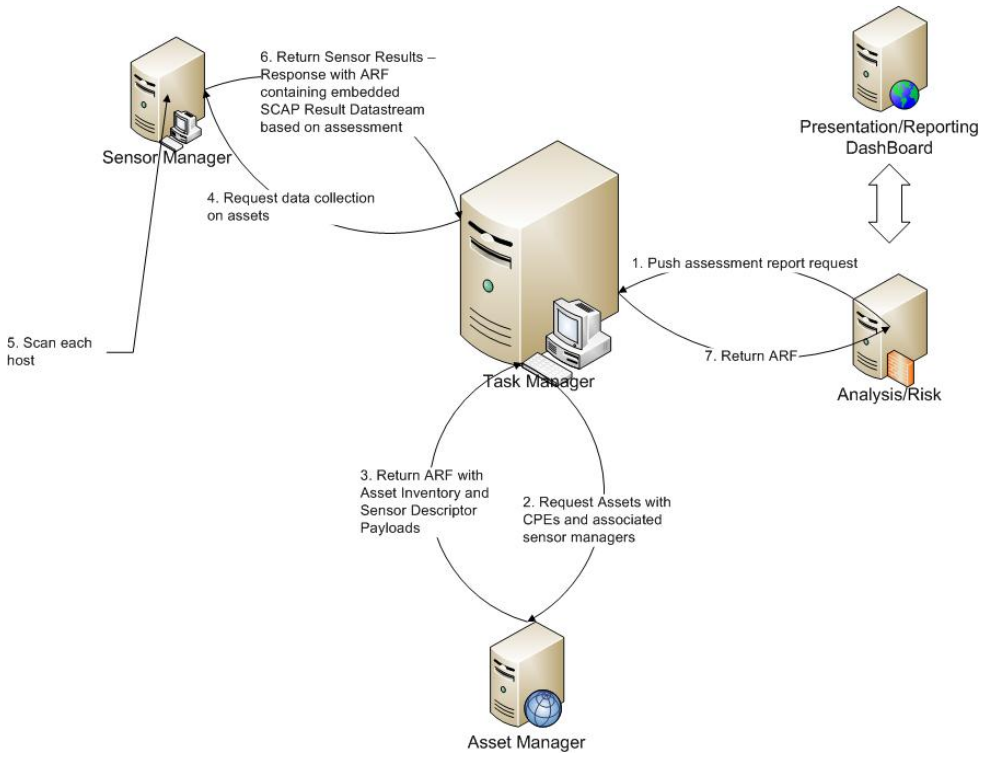

Figure A-1: Sample Vulnerability Management Workflow

The sample workflow shown in Figure A-1 illustrates that a dashboard view of the vulnerable state of an enterprise could be fed by ARF reports generated through a series of request and response transactions with different actors in the infrastructure.

Digital Events Analysis

Tracking and reporting on digital events across an organization can be an important task for understanding the real-time security posture of an enterprise. While this use case is not well developed within the security automation space at this time, it is still an important use case to consider when reporting on assets. In the context of ARF, digital event reporting is supported through relationships such as "recordedBy" and "initiatedBy", which, in conjunction with lower-level specifications focused on digital events, can help an organization report more consistently in this realm.

Compliance Assessment

Organizations are often required to maintain compliance with one or more internal or external policies related to its IT assets. These include such policies as Payment Card Industry (PCI) data security standards, Health Insurance Portability and Accountability Act (HIPAA) privacy and security rules, Sarbanes-Oxley (SOX) compliance, Federal Information Security Management Act (FISMA) reporting, and organizational policies such as Security Technical Implementation Guides (STIG) and National Security Agency (NSA) Security Configuration Guides (SCG), among others. Each of these policies has different scopes and intentions, but all of them require some level of compliance reporting. It is imperative that organizations have the resources and standards necessary to ensure compliance with these, or other, policies. Compliance with these policies often includes complex reporting and aggregation of data across the enterprise on a variety of attributes related to assets, and their relationships. ARF provides a standardized format that allows organizations to maintain a level of interoperability of data across its enterprise when reporting compliance metrics at various levels of abstraction. In addition, when ARF is adopted across organizations, the data about assets that those organizations report has a level of interoperability as well. Multiple ARF reports could be provided to consuming technologies for correlation, aggregation and higher-level reporting.

Business Information

This use case supports the exchange and/or sharing of asset information with other systems that could consume and communicate ARF data either actively or store the results for later use. Information exchanged could include organizational affiliation, function or role, and organizational security context.

Appendix B—Normative References

This appendix lists the normative references for the ARF specification.

[Asset Identification] NIST Interagency Report (IR) 7693 - Asset Identification 1.1, May 2011. See: http://scap.nist.gov/specifications/ai/index.html

[RFC 2119] Internet Engineering Task Force (IETF) Request for Comment (RFC) 2119: Key words for use in RFCs to Indicate Requirement Levels, March 1997. See: http://www.ietf.org/rfc/rfc2119.txt

[XLink] World Wide Web Consortium (W3C) XML Linking Language (XLink) Version 1.0, 27 June, 2001. See http://www.w3.org/TR/xlink/

[XML] W3C Recommendation Extensible Markup Language (XML) 1.0 (Fifth Edition), 26 November 2008. See: http://www.w3.org/TR/REC-xml/

[XML Schema] W3C Recommendation XML Schema, 28 October 2004. See: http://www.w3.org/XML/Schema.html

Appendix C—Sample Workflow

Figure C-1 shows one possible workflow in which an ARF report could be generated and used. It does not mandate a particular workflow nor exclude the use of other possible workflows.

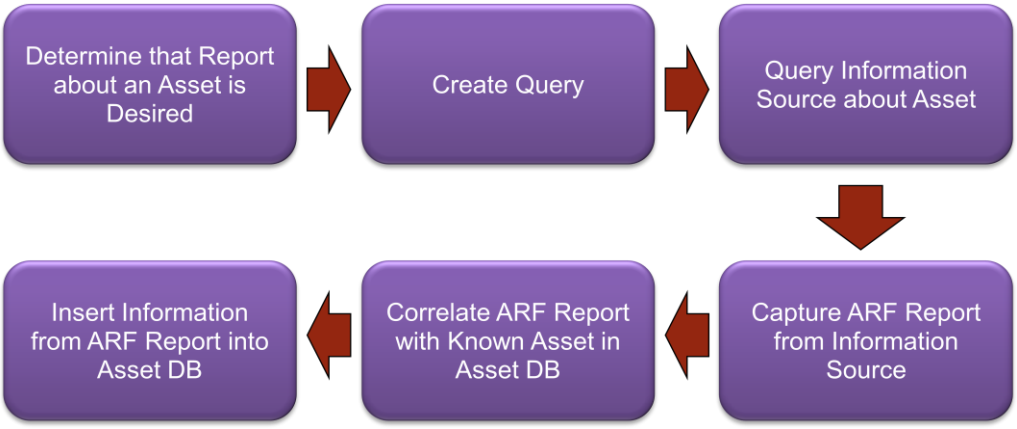

Figure C-1: Sample Workflow Using ARF

The workflow begins when a person decides that information about an asset is desired. The person then creates a query to retrieve that information. The query could be a request of data from a database, instructions to execute in a sensor, or any other applicable request. The information source then accepts that request and returns an ARF report about the asset. The information source could be a database, sensor (e.g., network packet capture device, event manager, intrusion detection system, thermometer), etc. The asset identification information contained in the report is then used to correlate the content of the report with other information about the asset already captured in an asset database. Once the information is correlated, the new information from the report can be inserted into the asset database and associated with the proper asset.

Appendix D—Sample XML

This appendix provides a sample XML document in the ARF format. This example is for reference only.

```
<?xml version="1.0" encoding="UTF-8"?>
<asset-report-collection xmlns:ai="http://scap.nist.gov/schema/asset-identification/1.1"
    xmlns="http://scap.nist.gov/schema/asset-reporting-format/1.1"
    xmlns:core="http://scap.nist.gov/schema/reporting-core/1.1"
xmlns:xlink="http://www.w3.org/1999/xlink"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xsi:schemaLocation="http://scap.nist.gov/schema/asset-reporting-format/1.1
http://scap.nist.gov/schema/asset-reporting-format/1.1/asset-reporting-format_1.1.0.xsd">
    <core:relationships xmlns:arfvocab="http://scap.nist.gov/vocabulary/arf/relationships/1.0#">
        <core:relationship type="arfvocab:isAbout" subject="report_1">
            <core:ref>asset_1</core:ref>
        </core:relationship>
        <core:relationship type="arfvocab:createdFor" subject="report_1">
            <core:ref>report_request_1</core:ref>
        </core:relationship>
    </core:relationships>
    <report-requests>
        <report-request id="report request 1">
            <content>
                <Benchmark id="minimal-xccdf" xml:lang="en-US"
xmlns="http://checklists.nist.gov/xccdf/1.1"
                    xmlns:cpe="http://cpe.mitre.org/dictionary/2.0"
xmlns:dc="http://purl.org/dc/elements/1.1/"
                    xmlns:xhtml="http://www.w3.org/1999/xhtml">
                    <status date="2009-12-01">draft</status>
                    <title>Test Title</title>
                    <description>
                        <xhtml:strong>Test Description</xhtml:strong>
                    </description>
                    <notice id="test-notice">Test Notice</notice>
                    <reference href="http://testreference1">
                        <dc:publisher>Test Publisher1</dc:publisher>
                        <dc:identifier>Test Identifier1</dc:identifier>
                    </reference>
                    <platform idref="cpe:/o:microsoft:windows vista"/>
                    <version>Test Version</version>
                    <metadata>
                        <dc:creator>Test Creator</dc:creator>
                        <dc:publisher>Test Publisher</dc:publisher>
                        <dc:contributor>Test Contributor</dc:contributor>
                        <dc:source>http://scap.nist.gov/</dc:source>
                    </metadata>
                    <Profile id="test profile1">
                        <title>Test Title for Profile 1</title>
                        <description>Test Description for Profile 1</description>
                        <select idref="test_rule1" selected="true"/>
                    </Profile>
                    <Rule id="test_rule1" selected="true" weight="10.0">
                        <title>Test Title for Rule 1</title>
                        <description>Test Description for Rule 1</description>
                        <ident system="http://cce.mitre.org">CCE-2466-1</ident>
                        <check system="http://oval.mitre.org/XMLSchema/oval-definitions-5">
                            <check-content-ref href="minimal-oval.xml"
name="oval:gov.nist.test.compliance:def:1"/>
                        </check>
                    </Rule>
                </Benchmark>
            </content>
        </report-request>
    </report-requests>
    <assets>
        <asset id="asset_1">
            <ai:computing-device>
                <ai:connections>
                    <ai:connection>
                        <ai:ip-address>
```

```
                            <ai:ip-v4>192.168.2.10</ai:ip-v4>
                        </ai:ip-address>
                    </ai:connection>
                </ai:connections>
                <ai:fqdn>comp1234.tempuri.org</ai:fqdn>
            </ai:computing-device>
        </asset>
    </assets>
    <reports>
        <report id="report_1">
            <content>
                <TestResult xmlns="http://checklists.nist.gov/xccdf/1.1" id="minimal-xccdf-
1280857747215"
                    version="Test Version" test-system="cpe:/a:nist:scap scanner:1.0"
                    start-time="2010-08-03T13:44:07.657-04:00" end-time="2010-08-03T13:49:07.657-
04:00">
                    <benchmark href="minimal-xccdf"/>
                    <title xml:lang="en-US">SCAP automated assessment for checklist minimal-xccdf
performed at Tuesday,
                        August 3, 2010</title>
                    <organization>National Institute of Standards and Technology</organization>
                    <identity authenticated="1" privileged="1">administrator</identity>
                    <profile idref="test_profile1"/>
                    <target>0:0:0:0:0:0:0:1</target>
                    <target>127.0.0.1</target>
                    <target>host.domain.tld</target>
                    <target-address>0:0:0:0:0:0:0:1</target-address>
                    <target-address>127.0.0.1</target-address>
                    <target-address>192.168.222.1</target-address>
                    <target-facts>
                        <fact name="urn:xccdf:fact:asset:identifier:host name"
type="string">0:0:0:0:0:0:0:1</fact>
                        <fact name="urn:xccdf:fact:asset:identifier:fqdn"
type="string">0:0:0:0:0:0:0:1</fact>
                        <fact name="urn:xccdf:fact:asset:identifier:ipv6"
type="string">0:0:0:0:0:0:0:1</fact>
                        <fact name="urn:xccdf:fact:asset:identifier:host name"
type="string">127.0.0.1</fact>
                        <fact name="urn:xccdf:fact:asset:identifier:fqdn"
type="string">127.0.0.1</fact>
                        <fact name="urn:xccdf:fact:asset:identifier:ipv4"
type="string">127.0.0.1</fact>
                        <fact name="urn:xccdf:fact:asset:identifier:mac" type="string"/>
                        <fact name="urn:xccdf:fact:asset:identifier:host name"
type="string">host.domain.tld</fact>
                        <fact name="urn:xccdf:fact:asset:identifier:fqdn"
type="string">host.domain.tld</fact>
                        <fact name="urn:xccdf:fact:asset:identifier:ipv4"
type="string">192.168.222.1</fact>
                        <fact name="urn:xccdf:fact:asset:identifier:mac"
type="string">00:50:56:c0:00:01</fact>
                    </target-facts>
                    <rule-result idref="test_rule1" time="2010-08-03T13:49:07.650-04:00"
weight="10.0">
                        <result>pass</result>
                        <ident system="http://cce.mitre.org">CCE-2466-1</ident>
                        <instance>host.domain.tld</instance>
                        <check system="http://oval.mitre.org/XMLSchema/oval-definitions-5">
                            <check-content-ref href="minimal-oval-res.xml"
name="oval:gov.nist.test.compliance:def:1"/>
                        </check>
                    </rule-result>
                    <score maximum="1" system="urn:xccdf:scoring:flat-unweighted">1</score>
                    <score maximum="10" system="urn:xccdf:scoring:flat">10</score>
                </TestResult>
            </content>
        </report>
    </reports>
</asset-report-collection>
```