                     Definition of ROLIE CSIRT Extension
                      draft-banghart-mile-rolie-csirt-00

Abstract

   This document extends the Resource-Oriented Lightweight Information
   Exchange (ROLIE) core to add the information type categories and
   related requirements needed to support Computer Security Incident
   Response Team (CSIRT) use cases. The indicator and incident
   information types are defined as ROLIE extensions. Additional
   supporting requirements are also defined that describe the use of
   specific formats and link relations pertaining to the new information
   types.

Contributing to this document

   The source for this draft is being maintained in GitHub. Suggested
   changes should be submitted as pull requests at
   <https://github.com/CISecurity/ROLIE>. Instructions are on that page
   as well. Editorial changes can be managed in GitHub, but any
   substantial issues need to be discussed on the MILE mailing list.

Status of This Memo

   This Internet-Draft is submitted to IETF in full conformance with the
   provisions of BCP 78 and BCP 79.

   Internet-Drafts are working documents of the Internet Engineering
   Task Force (IETF). Note that other groups may also distribute
   working documents as Internet-Drafts. The list of current Internet-
   Drafts is at http://datatracker.ietf.org/drafts/current/.

   Internet-Drafts are draft documents valid for a maximum of six months
   and may be updated, replaced, or obsoleted by other documents at any
   time. It is inappropriate to use Internet-Drafts as reference
   material or to cite them other than as "work in progress."

   This Internet-Draft will expire on May 4, 2017.

Copyright Notice

Table of Contents

1. Introduction

   Threats to computer security are evolving ever more rapidly as time
   goes on. As software increases in complexity, the number of
   vulnerabilities in systems and networks can increase exponentially.
   Threat actors looking to exploit these vulnerabilities are making
   more frequent and more widely distributed attacks across a large
   variety of systems. The adoption of liberal information sharing
   amongst attackers allows a discovered vulnerability to be shared and
   used to attack a vulnerable system within a narrow window of time.
   As the skills and knowledge required to identify and combat these
   attacks become more and more specialized, even a well established and
   secure system may find itself unable to quickly respond to an
   incident. Effective identification of and response to a
   sophisticated attack requires open cooperation and collaboration
   between defending operators, software vendors, and end-users. To
   improve the timeliness of responses, automation must be used to
   acquire, contextualize, and put to use shared computer security
   information.

   CSIRTS share two primary forms of information: incidents and
   indicators. Using these forms of information, analysts are able to
   perform a wide range of activities both proactive and reactive to
   ensure the security of their systems.

   Incident information describes a cyber security incident. Such
   information may include attack characteristics, information about the
   attacker, and attack vector data. Sharing this information helps
   analysts within the sharing community to inoculate their systems
   against similar attacks, providing proactive protection.

   Indicator information describes the symptoms or necessary pre-
   conditions of an attack. Everything from system vulnerabilities to
   unexpected network traffic can help analysts secure systems and
   prepare for an attack. Making this information available for sharing
   aids in the proactive defense of systems both within an operating
   unit but also for any CSIRTs that are part of a sharing consortium.

   As a means to bring automation of content discovery and dissemination
   into the CSIRT domain, this specification provides an extension to
   the Resource-Oriented Lightweight Information Exchange (ROLIE) core
   [I-D.ietf-mile-rolie] designed to address CSIRT use cases. The
   primary purpose of this extension is to define two new information

types: incident, and indicator, along with formats and link relations
that support these information-types.

2. Terminology

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 [RFC2119].

Definitions for some of the common computer security-related
terminology used in this document can be found in Section 2 of
[RFC5070].

3. New information-types

This document defines the following two information types:

3.1. The "incident" information type

The "incident" information type represents any information describing
or pertaining to a computer security incident. This document uses
the definition of incident provided by [RFC4949]. Provided below is
a non-exhaustive list of information that may be considered to be an
incident information type.

o Timing information: start and end times for the incident and/or
   the response.

o Descriptive information: plain text or machine readable data that
   provides some degree of description of the incident itself.

o Response information: the methods and results of a response to the
   incident.

o Meta and contact information: data about the CSIRT that recorded
   the information, or the operator that enacted the response.

o Effect and result information: data that describes the effects of
   an incident, or what the final results of the incident are.

Note again that this list is not exhaustive, any information that in
is the abstract realm of an incident should be classified under this
information-type.

3.2. The "indicator" information type

   The "indicator" information type represents computer security
   indicators or any information surrounding them. This document uses
   the definition of indicator provided by [RFC4949]. Some examples of
   indicator information is provided below, but note that indicator is
   defined in an abstract sense, to be understood as a flexible and
   widely-applicable definition.

   o Specific vulnerabilities that indicate a vector for attack.

   o Signs of malicious reconnaissance.

   o Definitions of patterns of other indicators.

   o Events that may indicate an attack and information regarding those
      events.

   o Meta information about the collecting agent.

   This list is intended to provide examples of the indicator
   information-type, not to define it.

4. Usage of CSIRT Information Types in the Atom Publishing Protocol

   These requirements apply when a ROLIE repository contains any
   Collections with categories with scheme attributes of either CSIRT
   information type, or if the CSIRT information types appear in the
   Categories document.

4.1. / (forward slash) Resource URL

   The forward slash resource URL MUST be supported as defined in
   Section 5.5 [I-D.ietf-mile-rolie]. Note that this is a stricter
   requirement than the core document.

5. Usage of CSIRT Information Types in the atom:feed element

   This document does not define any additional requirements for Feeds.

6. Usage of CSIRT Information Types in an atom:entry

   This document defines the following requirements for any Entries that
   are of the CSIRT information type categories.

6.1. Use of the atom:link element

   These sections define requirements for atom:link elements in Entries.
   Note that the requirements are determined by the information type
   that appears in either the Entry or in the parent Feed.

6.1.1. Link relations for the 'incident' information-type

   If the category of an Entry is the incident information type, then
   the following requirements MUST be followed for inclusion of
   atom:link elements.

Name	Description	Conformance
indicators	Provides a link to a collection of zero or more instances of cyber security indicators that are associated with the resource.	SHOULD
evidence	Provides a link to a collection of zero or more resources that provides some proof of attribution for an incident. The evidence may or may not have any identified chain of custody.	SHOULD
attacker	Provides a link to a collection of zero or more resources that provides a representation of the attacker.	SHOULD
vector	Provides a link to a collection of zero or more resources that provides a representation of the method used by the attacker.	SHOULD

              Table 1: Link Relations for Resource-Oriented Lightweight Indicator
                                    Exchange

6.1.2. Link relations for the 'indicator' information-type

   If the category of an Entry is the indicator information type, then
   the following requirements MUST be followed for inclusion of
   atom:link elements.

```
+----------+----------------------------------------+------------+
| Name     | Description                            | Conformance |
+----------+----------------------------------------+------------+
| incidents | Provides a link to a collection of zero | SHOULD     |
|          | or more instances of incident          |            |
|          | representations associated with the    |            |
|          | resource.                              |            |
+----------+----------------------------------------+------------+
```

      Table 2: Link Relations for Resource-Oriented Lightweight Indicator
                                  Exchange

6.1.3. Link relations for both information-types

   Regardless of the category of an Entry, the following requirements
   MUST be followed for inclusion of atom:link elements.

```
+----------------------+----------------------------+------------+
| Name                 | Description                 | Conformance |
+----------------------+----------------------------+------------+
| assessments          | Provides a link to a        | MAY        |
|                      | collection of zero or more  |            |
|                      | resources that represent    |            |
|                      | the results of executing a  |            |
|                      | benchmark.                  |            |
| reports              | Provides a link to a        | MAY        |
|                      | collection of zero or more  |            |
|                      | resources that represent    |            |
|                      | RID reports.                |            |
| traceRequests        | Provides a link to a        | MAY        |
|                      | collection of zero or more  |            |
|                      | resources that represent    |            |
|                      | RID traceRequests.          |            |
| investigationRequests | Provides a link to a        | MAY        |
|                      | collection of zero or more  |            |
|                      | resources that represent    |            |
|                      | RID investigationRequests.  |            |
+----------------------+----------------------------+------------+
```

      Table 3: Link Relations for Resource-Oriented Lightweight Indicator
                                  Exchange

6.2. Use of the rolie:format element

   This document defines two additional valid values for the 'ns'
   attribute of the rolie:format element. Both of these formats are
   valid for either information type.

6.2.1. IODEF Format

   If, and only if, the content of the Entry contains an IODEF document,
   then the 'ns' attribute of the rolie:format element MUST be 'TODO-
   IODEF-URN'.

   Any entry using the IODEF format MUST conform to the requirements in
   Section 6.3

   The Incident Object Description Exchange Format (IODEF) is a format
   for representing computer security information commonly exchanged
   between Computer Security Incident Response Teams (CSIRTs) or other
   operational security teams.

   IODEF conveys indicators, incident reports, response activities, and
   related meta-data in an XML serialization. This information is
   formally structured in order to support and encourage automated
   machine-to-machine security communication, as well as enhanced
   processing at the endpoint.

   The full IODEF specification provides further high-level discussion
   and technical details.

6.2.2. STIX Format

   If, and only if, the content of the Entry contains a STIX document,
   then the 'ns' attribute of the rolie:format element MUST be 'TODO-
   STIX-URN'.

   STIX is a structured language for describing a wide range of security
   resources. STIX approaches the problem with a focus on flexibility,
   automation, readability, and extensibility.

   The use of STIX as the content of an Entry does not impose any
   additional requirements on ROLIE implementations.

6.3. Additional requirements for use of IODEF

   This section provides the normative requirements for usage of the
   IODEF format.

6.3.1. The IODEF Document

   An IODEF document that is carried in an Atom Entry SHOULD NOT contain
   a <relatedActivity> element. Instead, the related activity SHOULD be
   available via a link rel=related.

An IODEF document that is carried in an Atom Entry SHOULD NOT contain a <history> element. Instead, the related history SHOULD be available via a link rel="history" (todo: or a fully qualified link rek name). The associated href MAY leverage OpenSearch to specify the required query.

6.3.2. Category Element

A collection or entry containing IODEF incident content MUST contain at least two <category> element. One cateogry elemnt must have the scheme attribute be equal to 'urn:ietf:params:rolie:category:iodef:purpose' and the other 'urn:ietf:params:rolie:category:iodef:restriction'. This metadata provides valuable metadata for searching and organization.

When the scheme attribute of this element is 'urn:ietf:params:rolie:category:iodef:purpose', the term attribute MUST be constrained as per section 3.2 of IODEF, e.g. traceback, mitigation, reporting, or other.

When the scheme attribute of this element is 'urn:ietf:params:rolie:category:iodef:restriction', the term attribute MUST be constrained as per section 3.2 of IODEF, e.g. public, need-to-know, private, default.

6.3.3. Entry Elements

AUTHORS NOTE: This section is TODO.

An entry containing an IODEF payload MUST contain aa <ROLIE-CSIRT:ID> element. This element The ID element for an Atom entry SHOULD be established via the concatenation of the value of the name attribute from the IODEF <IncidentID> element and the corresponding value of the <IncidentID> element. This requirement ensures a simple and direct one-to-one relationship between an IODEF incident ID and a corresponding Feed entry ID and avoids the need for any system to maintain a persistent store of these identity mappings.

(todo: Note that this implies a constraint on the IODEF document that is more restrictive than the current IODEF schema. IODEF section 3.3 requires only that the name be a STRING type. Here we are stating that name must be an IRI. Possible request to update IODEF to constrain, or to support a new element or attribute).

6.3.4. User Authorization

   When the content model for the Atom <content> element of an Atom
   Entry contains an <IODEF-Document>, then authorization MUST be
   adjudicated based upon the Atom <category> element(s), whose values
   have been mapped as per Section 6.3.2.

6.3.5. Expectation and Impact Classes

   It is frequently the case that an organization will need to triage
   their investigation and response activities based upon, e.g., the
   state of the current threat environment, or simply as a result of
   having limited resources.

   In order to enable operators to effectively prioritize their response
   activity, it is RECOMMENDED that feed implementers provide Atom
   categories that correspond to the IODEF Expectation and Impact
   classes. The availability of these feed categories will enable
   clients to more easily retrieve and prioritize cyber security
   information that has already been identified as having a specific
   potential impact, or having a specific expectation.

   Support for these categories may also enable efficiencies for
   organizations that already have established (or plan to establish)
   operational processes and workflows that are based on these IODEF
   classes.

6.3.6. Search

   Implementers SHOULD support search based upon the IODEF AlternativeID
   class as a search parameter.

   Implementers SHOULD support search based upon the four timestamp
   elements of the IODEF Incident class: <startTime>, <EndTime>,
   <DetectTime>, and <ReportTime>.

   Implementers MAY support additional search capabilities based upon
   any of the remaining elements of the IODEF Incident class, including
   the <Description> element.

   Collections that support use of the RID schema as a content model in
   the Atom member entry <content> element (e.g. in a report resource
   representation reachable via the "report" link relationship) MUST
   support search operations that include the RID MessageType as a
   search parameter, in addition to the aforementioned IODEF schema
   elements, as contained within the <ReportSchema> element.

7. IANA Considerations

7.1. incident information-type

   IANA has added an entry to the "ROLIE Security Resource Information
   Type Sub-Registry" registry located at
   <https://www.iana.org/assignments/rolie/category/information-type> .

   The entry is as follows:

      name: incident

      index: TBD

      reference: This document, Section 3.1

7.2. indicator information-type

   IANA has added an entry to the "ROLIE Security Resource Information
   Type Sub-Registry" registry located at
   <https://www.iana.org/assignments/rolie/category/information-type> .

   The entry is as follows:

      name: indicator

      index: TBD

      reference: This document, Section 3.2

8. Security Considerations

   This document implies the use of ROLIE in high-security use cases, as
   such, added care should be taken to fortify and secure ROLIE
   repositories and clients using this extension. The guidance in the
   ROLIE core specification is strongly recommended, and implementers
   should consider adding additional security measures as they see fit.

9. Normative References

   [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate
              Requirement Levels", BCP 14, RFC 2119,
              DOI 10.17487/RFC2119, March 1997,
              <http://www.rfc-editor.org/info/rfc2119>.

   [RFC5070] Danyliw, R., Meijer, J., and Y. Demchenko, "The Incident
              Object Description Exchange Format", RFC 5070,
              DOI 10.17487/RFC5070, December 2007,
              <http://www.rfc-editor.org/info/rfc5070>.

   [I-D.ietf-mile-rolie]
              Field, J., Banghart, S., and D. Waltermire, "Resource-
              Oriented Lightweight Information Exchange", draft-ietf-
              mile-rolie-03 (work in progress), July 2016.

   [RFC4949] Shirey, R., "Internet Security Glossary, Version 2",
              FYI 36, RFC 4949, DOI 10.17487/RFC4949, August 2007,
              <http://www.rfc-editor.org/info/rfc4949>.

Appendix A. Non-Normative Examples

   TODO

A.1. Use of Link Relations

   A key benefit of using the RESTful architectural style is the ability
   to enable the client to navigate to related resources through the use
   of hypermedia links. In the Atom Syndication Format, the type of the
   related resource identified in a <link> element is indicated via the
   "rel" attribute, where the value of this attribute identifies the
   kind of related resource available at the corresponding "href"
   attribute. Thus, in lieu of a well-known URI template the URI itself
   is effectively opaque to the client, and therefore the client must
   understand the semantic meaning of the "rel" attribute in order to
   successfully navigate. Broad interoperability may be based upon a
   sharing consortium defining a well-known set of Atom Link Relation
   types. These Link Relation types may either be registered with IANA,
   or held in a private registry.

   Individual CSIRTs may always define their own link relation types in
   order to support specific use cases, however support for a core set
   of well-known link relation types is encouraged as this will maximize
   interoperability.

   In addition, it may be beneficial to define use case profiles that
   correspond to specific groupings of supported link relationship
   types. In this way, a CSIRT may unambiguously specify the classes of
   use cases for which a client can expect to find support.

   The following sections provide non-normative examples of link
   relation usage. Three distinct cyber security information sharing
   use case scenarios are described. In each use case, the unique
   benefits of adopting a resource-oriented approach to information

sharing are illustrated. It is important to note that these use
cases are intended to be a small representative set and is by no
means meant to be an exhaustive list. The intent is to illustrate
how the use of link relationship types will enable this resource-
oriented approach to cyber security information sharing to
successfully support the complete range of existing use cases, and
also to motivate an initial list of well-defined link relationship
types.

A.1.1. Use Case: Incident Sharing

This section provides a non-normative example of an incident sharing
use case.

In this use case, a member CSIRT shares incident information with
another member CSIRT in the same consortium. The client CSIRT
retrieves a feed of incidents, and is able to identify one particular
entry of interest. The client then does an HTTP GET on that entry,
and the representation of that resource contains link relationships
for both the associated "indicators" and the incident "history", and
so on. The client CSIRT recognizes that some of the indicator and
history may be relevant within her local environment, and can respond
proactively.

Example HTTP GET response for an incident entry:

```
<?xml version="1.0" encoding="UTF-8"?>
<entry>
  <id>http://www.example.org/csirt/private/incidents/123456</id>
  <title>Sample Incident</title>
  <link href="http://www.example.org/csirt/private/incidents/123456"
    rel="self"/>
  <link href="http://www.example.org/csirt/private/incidents/123456"
    rel="alternate"/>
  <published>2012-08-04T18:13:51.0Z</published>
  <updated>2012-08-05T18:13:51.0Z</updated>

  <link href="http://www.example.org/csirt/private/incidents/123456"
    rel="edit"/>

  <!-- The links to indicators related to this incident,
       and the history of this incident, and so on.... -->
  <link href="http://www.example.org/csirt/private/incidents/123456
    /relationships/indicators" rel="indicators"/>
  <link href="http://www.example.org/csirt/private/incidents/123456
    /relationships/history" rel="history"/>
  <link href="http://www.example.org/csirt/private/incidents/123456
    /relationships/campaign" rel="campaign"/>

  <!-- navigate up to the full collection.
       Might also be rel="collection" as per IANA registry -->
  <link href="http://www.example.org/csirt/private/incidents" rel="up"/>

  <content type="application/xml" src="example.org/123456/source">
  <!-- Content provided here as example, the content tag is only a
       link to this file. -->
    <iodef:IODEF-Document lang="en"
      xmlns:iodef="urn:ietf:params:xml:ns:iodef-1.0">
      <iodef:Incident purpose="traceback" restriction="need-to-know">
        <iodef:IncidentID name="http://www.example.org/csirt/private/
          incidents">123456</iodef:IncidentID>
        <!-- ...additional incident data.... -->
        </iodef:Incident>
    </iodef:IODEF-Document>
  </content>
</entry>
```

   As can be seen in the example response, the Atom <link> elements
   enable the client to navigate to the related indicator resources,
   and/or the history entries associated with this incident.

A.1.2. Use Case: Collaborative Investigation

   This section provides a non-normative example of a collaborative
   investigation use case.

   In this use case, two member CSIRTs that belong to a closed sharing
   consortium are collaborating on an incident investigation. The
   initiating CSIRT performs an HTTP GET to retrieve the service
   document of the peer CSIRT, and determines the collection name to be
   used for creating a new investigation request. The initiating CSIRT
   then POSTs a new incident entry to the appropriate collection URL.
   The target CSIRT acknowledges the request by responding with an HTTP
   status code 201 Created.

   Example HTTP GET response for the service document:

```
HTTP/1.1 200 OK
Date: Fri, 24 Aug 2012 17:09:11 GMT
Content-Length: 934
Content-Type: application/atomsvc+xml;charset="utf-8"

<?xml version="1.0" encoding="UTF-8"?>
<service xmlns="http://www.w3.org/2007/app"
         xmlns:atom="http://www.w3.org/2005/Atom">
    <workspace xml:lang="en-US"
      xmlns:xml="http://www.w3.org/XML/1998/namespace">
      <atom:title type="text">RID Use Case Requests</atom:title>
      <collection
        href="http://www.example.org/csirt/RID/InvestigationRequests">
          <atom:title type="text">Investigation Requests</atom:title>
          <accept>application/atom+xml; type=entry</accept>
      </collection>
      <collection href="http://www.example.org/csirt/RID/TraceRequests">
          <atom:title type="text">Trace Requests</atom:title>
          <accept>application/atom+xml; type=entry</accept>
      </collection>
      <!-- ...and so on.... -->
    </workspace>
</service>
```

   As can be seen in the example response, the Atom <collection>
   elements enable the client to determine the appropriate collection
   URL to request an investigation or a trace.

   The client CSIRT then POSTs a new entry to the appropriate feed
   collection. Note that the <content> element of the new entry may
   contain a RID message of type "InvestigationRequest" if desired,
   however this would NOT be required. The entry content itself need

only be an IODEF document, with the choice of the target collection
resource URL indicating the callers intent. A CSIRT would be free to
use any URI template to accept investigationRequests.

```
POST /csirt/RID/InvestigationRequests HTTP/1.1
Host: www.example.org
Content-Type: application/atom+xml;type=entry
Content-Length: 852

<?xml version="1.0" encoding="UTF-8"?>
<entry xmlns="http://www.w3.org/2005/Atom">
  <title>New Investigation Request</title>
  <id>http://www.example2.org/csirt/private/incidents/123456</id>
  <!-- id and updated not guranteed to be preserved -->
  <!-- may want to profile that behavior in this document -->
  <updated>2012-08-12T11:08:22Z</updated>
  <author><name>Name of peer CSIRT</name></author>
  <content type="application/xml">
    <iodef:IODEF-Document lang="en"
      xmlns:iodef="urn:ietf:params:xml:ns:iodef-1.0">
      <iodef:Incident purpose="traceback" restriction="need-to-know">
      <iodef:IncidentID name="http://www.example2.org/csirt/
        private/incidents">123</iodef:IncidentID>
        <!-- ...additional incident data.... -->
      </iodef:Incident>
    </iodef:IODEF-Document>
  </content>
</entry>
```

The receiving CSIRT acknowledges the request with HTTP return code
201 Created.

```
HTTP/1.1 201 Created
Date: Fri, 24 Aug 2012 19:17:11 GMT
Content-Length: 906
Content-Type: application/atom+xml;type=entry
Location: http://www.example.org/csirt/RID/InvestigationRequests/823
ETag: "8a9h9he4qphqh"

<?xml version="1.0" encoding="UTF-8"?>
<entry xmlns="http://www.w3.org/2005/Atom">
  <title>New Investigation Request</title>
  <id>http://www.example.org/csirt/RID/InvestigationRequests/823</id>
  <!-- id and updated not guranteed to be preserved -->
  <!-- may want to profile that behavior in this document -->
  <updated>2012-08-12T11:08:30Z</updated>
  <published>2012-08-12T11:08:30Z</published>
  <author><name>Name of peer CSIRT</name></author>
  <content type="application/xml">
    <iodef:IODEF-Document lang="en"
      xmlns:iodef="urn:ietf:params:xml:ns:iodef-1.0">
      <iodef:Incident purpose="traceback" restriction="need-to-know">
      <iodef:IncidentID name="http://www.example.org/csirt/private
        /incidents">123</iodef:IncidentID>
        <!-- ...additional incident data.... -->
      </iodef:Incident>
    </iodef:IODEF-Document>
  </content>
</entry>
```

Consistent with HTTP/1.1 RFC, the location header indicates the URL
of the newly created InvestigationRequest. If for some reason the
request were not authorized, the client would receive an HTTP status
code 403 Unauthorized. In this case the HTTP response body may
contain additional details, if an as appropriate.

A.1.3. Use Case: Cyber Data Repository

   This section provides a non-normative example of a cyber security
   data repository use case.

   In this use case a client accesses a persistent repository of cyber
   security data via a RESTful usage model. Retrieving a feed
   collection is analogous to an SQL SELECT statement producing a result
   set. Retrieving an individual Atom Entry is analogous to a SQL
   SELECT statement based upon a primary key producing a unique record.
   The cyber security data contained in the repository may include
   different data types, including indicators, incidents, benchmarks, or
   any other related resources. In this use case, the repository is
   queried via HTTP GET, and the results that are returned to the client

may optionally contain URL references to other cyber security
resources that are known to be related. These related resources may
also be persisted locally, or they may exist at another (remote)
cyber data respository.

Example HTTP GET request to a persistent repository for any resources
representing Distributed Denial of Service (DDOS) attacks:

```
GET /csirt/repository/ddos
Host: www.example.org
Accept: application/atom+xml
```

The corresponding HTTP response would be an XML document containing
the DDOS feed.

Example HTTP GET response for a DDOS feed:

```
HTTP/1.1 200 OK
Date: Fri, 24 Aug 2012 17:20:11 GMT
Content-Length: nnnn
Content-Type: application/atom+xml;type=feed;charset="utf-8"

<?xml version="1.0" encoding="UTF-8"?>
<feed xmlns="http://www.w3.org/2005/Atom"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xsi:schemaLocation="http://www.w3.org/2005/Atom
                        file:/C:/schemas/atom.xsd
                        urn:ietf:params:xml:ns:iodef-1.0
                        file:/C:/schemas/iodef-1.0.xsd"
    xml:lang="en-US">

    <generator version="1.0" xml:lang="en-US">
      emc-csirt-iodef-feed-service</generator>
    <id>http://www.example.org/csirt/repository/ddos</id>
    <title type="text" xml:lang="en-US">
      Atom formatted representation of a feed of known ddos resources.
      </title>
    <updated xml:lang="en-US">2012-05-04T18:13:51.0Z</updated>
    <author>
        <email>csirt@example.org</email>
        <name>EMC CSIRT</name>
    </author>

    <!-- By convention there is usually a self link for the feed -->
    <link href="http://www.example.org/csirt/repository/ddos"
      rel="self"/>
```

```
   <entry>
       <id>http://www.example.org/csirt/repository/ddos/123456</id>
       <title>Sample DDOS Incident</title>
       <link href="http://www.example.org/csirt/repository/ddos/123456"
         rel="self"/>          <!-- by convention -->
       <link href="http://www.example.org/csirt/repository/ddos/123456"
         rel="alternate"/>     <!-- required by Atom spec -->
       <link href="http://www.example.org/csirt/repository/ddos/987654"
         rel="related"/>       <!-- link to a related DDOS resource
         in this repository -->
       <link href="http://www.cyber-agency.gov/repository/
         indicators/1a2b3c" rel="related"/>
         <!-- link to a related DDOS resource in another repository -->
       <published>2012-08-04T18:13:51.0Z</published>
       <updated>2012-08-05T18:13:51.0Z</updated>
       <!-- The category is based upon IODEF
                   purpose and restriction attributes -->
       <category term="traceback" scheme="purpose" label="trace back"/>
       <category term="need-to-know" scheme="restriction"
         label="need to know" />
       <category term="ddos" scheme="ttp"
         label="tactics, techniques, and procedures"/>
       <summary>A short description of this DDOS attack, extracted
       from the IODEF Incident class, <description> element. </summary>
   </entry>

   <entry>
       <!-- ...another entry... -->
   </entry>

</feed>
```

   This feed document has two atom entries, one of which has been
   elided. The completed entry illustrates an Atom <entry> element that
   provides a summary of essential details about one particular DDOS
   incident. Based upon this summary information and the provided
   category information, a client may choose to do an HTTP GET operation
   to retrieve the full details of the DDOS incident. This example
   shows how a persistent repository may provide links to additional
   resources, both local and remote.

   Note that the provider of a persistent repostory is not obligated to
   follow any particular URL template scheme. The repository available
   at the hypothetical provider "www.example.com" uses a different URL
   pattern than the hypothetical repository available at "www.cyber-
   agency.gov". When a client de-references a link to resource that is
   located in a remote repository the client may be challenged for
   authentication credentials acceptable to that provider. If the two

repository providers choose to support a federated identity scheme or some other form of single-sign-on technology, then the user experience can be improved for interactive clients (e.g., a human user at a browser). However, this is not required and is an implementation choice that is out of scope for this specification.

Authors' Addresses

John P. Field
Pivotal Software, Inc.
625 Avenue of the Americas
New York, New York
USA

Phone: (646)792-5770
Email: jfield@pivotal.io

Stephen A. Banghart
National Institute of Standards and Technology
100 Bureau Drive
Gaithersburg, Maryland
USA

Phone: (301)975-4288
Email: sab3@nist.gov