REPUTE BoF

Murray S. Kucherawy
<msk@cloudmark.com>

Note Well

Any submission to the IETF intended by the Contributor for publication as all or part of an IETF Internet-Draft or RFC and any statement made within the context of an IETF activity is considered an "IETF Contribution". Such statements include oral statements in IETF sessions, as well as written and electronic communications made at any time or place, which are addressed to:

- The IETF plenary session
- The IESG, or any member thereof on behalf of the IESG
- Any IETF mailing list, including the IETF list itself, any working group or design team list, or any other list functioning under IETF auspices
- Any IETF working group or portion thereof
- The IAB or any member thereof on behalf of the IAB
- The RFC Editor or the Internet-Drafts function

All IETF Contributions are subject to the rules of RFC 5378 and RFC 3979 (updated by RFC 4879).

Statements made outside of an IETF session, mailing list or other function, that are clearly not intended to be input to an IETF activity, group or function, are not IETF Contributions in the context of this notice.

Please consult RFC 5378 and RFC 3979 for details.

A participant in any IETF activity is deemed to accept all IETF rules of process, as documented in Best Current Practices RFCs and IESG Statements.

A participant in any IETF activity acknowledges that written, audio and video records of meetings may be made and may be available to the public.

REPUTE BoF Agenda

1. Introduction and administrative notes - 5 min
2. Problem statement and scope - 15 min
3. Discussion of work so far - 15 min
4. Discussion of proposed charter - 15 min
5. Questions - 10 min
   - Are the problem statement and scope clear?
   - Is this appropriate for IETF work?
   - Who will work on it?

Problem Space

- Several security technologies today need to be able to determine the *value* of an identifier
  - *Value* is intentionally undefined here because it will mean different things in different contexts
- Most common examples of these include IP addresses and domain names used when evaluating email
  - Is mail from example.com desirable?
  - Does 1.2.3.4 appear to be hacked?

Problem Space

- We have some very rudimentary reputation services already
  - DNSBLs (RFC5782) can indicate presence of an IP address on a block list
  - Vouch-By-Reference (RFC5518) can indicate that one domain vouches for another in certain specific contexts
- But these replies are binary only

Problem Space

- What's missing is something more general and extensible
  - Extend easily into new application spaces and sub-spaces
  - Allow something other than a Boolean expression of membership in a set
  - A means of indicating the service's confidence in its own answer

Problem Space

- This is a logical follow-on to technologies like DKIM, SPF and Sender-ID that authenticate a domain name on a piece of email
  - "OK, so I know this mail really came from (or was authorized by) example.com… so what?"

Problem Statement

- “Security applications need a general framework for querying one or more authorities about the reputation of an identifier. This framework must be extensible into new applications, and must support a plurality of possible assertions within each application, and an expression of confidence in the answer.”

What's In Scope

- The framework, including media types and other definitions
- A lightweight protocol for simple queries and replies, and a heavyweight protocol for more structured or detailed queries and replies
  - We'll try re-use existing wire protocols to move the data around; only talking about payload here
- Initial application definitions, such as for email

What's Not In Scope

- How a reputation service provider computes reputation or collects data to do so
  - Method of collection (feedback)
  - Specific data to be collected
- How a reputation consumer makes use of the reply from a reputation service
  - Fully a local policy/configuration decision

The Story So Far

- Reputation was specifically declared out-of-scope for DKIM, SPF and Sender-ID even though it's an obvious application making use of those technologies
- Now that those are all out and stable, it's the right time to start working on this new layer of those applications
- No current working groups are operating in this space
  - Doesn't seem to fit in APPSAWG, which wasn't meant for big projects

The Story So Far

- Created a non-WG mailing list about a year ago (domainrep@ietf.org) where discussion has taken place
  - Original scope was reputation about domain names only, hence the name; some people on the list urged us to "think bigger"
  - Draft charter circulated there
- Initial suite of drafts now available
- Approached APPS area ADs
  - Reputation is an applications layer service
  - Should have a SEC AD as advisor to cover privacy and security of the queries and responses

Drafts Available

- draft-kucherawy-reputation-model
  - Defines the overall concept, basic definitions, and framework
- draft-kucherawy-reputation-media-type
  - Defines a new media type to contain the full reply in XML form
  - Some people would rather see JSON; we can discuss this

Drafts Available

- draft-kucherawy-reputation-query-dns
  - Defines a query/response format using TXT records in the DNS
- draft-kucherawy-reputation-query-http
  - Defines a query/response format using HTTP and XML, and specifies an XML schema
- draft-kucherawy-reputation-query-udp
  - Template for a document to contain a UDP query/reply format (currently mostly empty)
- draft-kucherawy-reputation-vocabulary-identities
  - Initial “vocabulary” for providing reputation service about identities found in email

Draft Charter

- Basically says all of the above, and sets out milestones matching the drafts that have been started:
  - Framework and definitions (informational)
  - Media type
  - Lightweight query
  - Heavyweight query
  - Email vocabulary
  - Protocol for reporting reputation data (feedback)
    - Old expired draft from mimedefang project

Questions To Explore

- Is the problem space identified correctly?
  - What are we missing?
- Is the problem statement the right one?
  - Too narrow, too broad?
- Is the scope right?
  - Should more stuff be declared in or out of scope?

Questions To Explore

- Who will be able to review documents?
- Who will be willing to edit documents?
- Who will be willing to participate in development and testing of prototypes?
- Who would be interested in co-chairing a working group?
- What are some reasonable dates for the proposed milestones?