

# I2RS Traceability

draft-ietf-i2rs-traceability-00

**I2RS Virtual Interim**

January 8, 2015

Joe Clarke, Gonzalo Salgueiro, Carlos Pignataro

# Overview

## “ **Describes a framework and IM for I2RS traceability**

- **Motivation and use cases for traceability**
- **Useful for troubleshooting, accounting, auditing, etc.**
- **Consistent tracing fields to be supported by implementations**
- **Suggests protocols used for data export**

## “ **Current Status**

- **Adopted as a working group draft on December 16, 2014**

# Current Feedback – 1 / 12

- Log entry only describes one timestamp. Which is it (client request, agent request, client reply, agent reply)? [Nobo Akiya]

**PLAN:** Add an additional timestamp to differentiate agent request/completion (I2RS logging happens at the agent)

## Current Feedback – 2 / 12

- In section 7.4.1, the draft describes syslog export. RFC5424 structured data elements should be used to encode fields. [Alex Clemm]

**PLAN:** Update section to refer to RFC5424 and that structured data elements SHOULD be used when using syslog data retrieval

# Current Feedback – 3 / 12

- The Service Model for traceability should be included as it covers items such as log size, log policy, etc. [Kwang-koog Lee]

**PLAN:** Get clarification on what is meant here. Is there a specific service model to reference or is the ask just to include a field for a service model name in the log entry?

# Current Feedback – 4 / 12

- We say that NULL MUST be used if an operation has no operation data; but the operation data may be NULL itself. [Ignas Bagdonas]

**PLAN:** Define a Boolean field to indicate whether or not operational data exists.

# Current Feedback – 5 / 12

- **Asynchronous, long running, blocking operations. Client request may not always be processed synchronously or within a bounded amount of time. To keep Operation and Result Code values in the same record may require buffering the trace log entries, and that may result in additional resource load on the agent and network element. [Ignas Bagdonas]**

**PLAN:** Yes, logging may require additional resources. We can mention that in the draft (small sentence). In terms of timing, see issue #1. We will have multiple timestamps to account for possible delays or blocking.

## Current Feedback – 6 / 12

- **Blocking on traceability information export. Traceability information export is a valuable diagnostics tool, but that is not the main function of the I2RS agent, and network element as such. Possible blocking of traceability component should not block the operation of the agent. [Ignas Bagdonas]**

**PLAN:** We will add text to this effect.

# Current Feedback – 7 / 12

- **Temporary on-element traceability data storage requirements. Related to the blocking point above and depending on many implementation aspects it may be more practical to store the traceability data and export/buffer it periodically than to do it synchronously with the requests. In such case the amount of resources required for such temporary storage must not interfere with normal operation of the agent itself.**  
[Ignas Bagdonas]

**PLAN:** Storage implementation is outside the scope of our draft (and stated so already). We are not planning changes.

# Current Feedback – 8 / 12

- **Intermixed operations.** I2RS agent may respond to incoming requests non-sequentially, different operations may take different amount of time required for completion. Batching of traceability data export would need to account for a possibility of signaled operation still being processed at the time of export. [Ignas Bagdonas]

**PLAN:** Implementation of export is outside the scope of this draft. We are not planning to make changes for this item.

# Current Feedback – 9 / 12

- **Timestamp granularity. RFC3339 defines subsecond granularity in timestamps but leaves the granularity of it aside. While this is highly implementation dependent, the nature of multiple and rapid operations would tend to ask for a recommended minimum granularity of trace records to be specified. While not enforcing, it could be recommended to support UNIX style 32.32 bit second.microsecond or 64 bit nanosecond timestamp granularity represented in RFC3339 format. [Ignas Bagdonas]**

**PLAN:** We will add text that an implementation **MUST** use milliseconds and **SHOULD** use nanoseconds.

# Current Feedback – 10 / 12

- **Section 7.1: The term 'transaction' in this paragraph seems to describe the internal machinery of the agent that will likely be dependent on many implementation factors and possibly not having much meaning outside the context of such implementation if exported via the traceability mechanism. The I2RS operation level transactions typically would be controlled by the Actor and/or Client, and would not be visible to the Agent. Could you clarify the meaning of the transaction term as used in this context? [Ignas Bagdonas]**

**PLAN:** The term “transaction” is used in Section 7.9, which describes error handling. However, we will update the text to use “operation” instead of “transaction”.

# Current Feedback – 11 / 12

- **More-trace-logs-follow** marker. An operation may return in multiple (sub-)results, possibly spread over a longer period of time compared to request processing and initial trace entry generation. A mechanism for recording into trace log that more output will follow at some later time would be **useful**. [Ignas Bagdonas]

**PLAN:** The authors were thinking to generate a log entry when the operation completes or times out without over-complicating the generation. Is there strong WG desire to chunk logs like this?

# Current Feedback – 12 / 12

- **Request/response traceability split, sub-response identifiers. A single request operation may result in more than one actual operation performed and more than one response being returned. Supporting such trace records would need to have a request and response correlation identifiers and ability to identify multiple responses. [Ignas Bagdonas]**

**PLAN:** Each message has an entry ID. We can cross-link messages by including a related message field. However, we should get overall WG consensus on this item.

# Next Steps

- Get clarification/consensus on items from mailing list
- Perform planned updates
- Publish a -01 by early February
- Ask for WGLC at/before the next IETF

This is a relatively simple document