

I2RS  
Internet-Draft  
Intended status: Informational  
Expires: April 01, 2014

J. Clarke  
G. Salgueiro  
C. Pignataro  
Cisco  
September 28, 2013

Interface to the Routing System (I2RS) Traceability: Framework and  
Information Model  
draft-clarke-i2rs-traceability-00

## Abstract

This document describes a framework for traceability in the Interface to the Routing System (I2RS) and information model for that framework. It specifies the motivation, requirements, use cases, and defines an information model for recording interactions between elements implementing the I2RS protocol. This framework provides a consistent tracing interface for components implementing the I2RS architecture to record what was done, by which component, and when. It aims to improve the management of I2RS implementations, and can be used for troubleshooting, auditing, forensics, and accounting purposes.

## Status of This Memo

This Internet-Draft is submitted 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 April 01, 2014.

## Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents

Internet-Draft

I2RS Traceability

September 2013

(<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

## Table of Contents

<a href="#">1.</a>	Introduction . . . . .	<a href="#">2</a>
<a href="#">2.</a>	Terminology and Conventions . . . . .	<a href="#">3</a>
<a href="#">3.</a>	Motivation and Use Cases . . . . .	<a href="#">3</a>
<a href="#">4.</a>	Information Model . . . . .	<a href="#">4</a>
<a href="#">4.1.</a>	I2RS Traceability Framework . . . . .	<a href="#">4</a>
<a href="#">4.2.</a>	I2RS Trace Log Mandatory Fields . . . . .	<a href="#">5</a>
<a href="#">4.3.</a>	I2RS Trace Log Extensibility and Optional Fields . . . . .	<a href="#">6</a>
<a href="#">4.4.</a>	I2RS Trace Log Syntax . . . . .	<a href="#">6</a>
<a href="#">5.</a>	Examples . . . . .	<a href="#">7</a>
<a href="#">6.</a>	Operational Guidance . . . . .	<a href="#">8</a>
<a href="#">6.1.</a>	Trace Log Creation . . . . .	<a href="#">8</a>
<a href="#">6.2.</a>	Trace Log Temporary Storage . . . . .	<a href="#">8</a>
<a href="#">6.3.</a>	Trace Log Rotation . . . . .	<a href="#">9</a>
<a href="#">6.4.</a>	Trace Log Retrieval . . . . .	<a href="#">9</a>
<a href="#">6.4.1.</a>	Retrieval Via Syslog . . . . .	<a href="#">9</a>
<a href="#">6.4.2.</a>	Retrieval Via I2RS Information Collection . . . . .	<a href="#">9</a>
<a href="#">6.4.3.</a>	Retrieval Via I2RS Pub-Sub . . . . .	<a href="#">10</a>
<a href="#">7.</a>	IANA Considerations . . . . .	<a href="#">10</a>
<a href="#">8.</a>	Security Considerations . . . . .	<a href="#">10</a>
<a href="#">9.</a>	References . . . . .	<a href="#">10</a>
<a href="#">9.1.</a>	Normative References . . . . .	<a href="#">10</a>
<a href="#">9.2.</a>	Informative References . . . . .	<a href="#">11</a>
	Authors' Addresses . . . . .	<a href="#">11</a>

[1.](#) Introduction

The architecture for the Interface to the Routing System ([[I-D.ietf-i2rs-architecture](#)]) specifies that I2RS Clients wishing to retrieve or change routing state on a routing element MUST authenticate to an I2RS Agent. The I2RS Client will have a unique identity it provides for authentication, and should provide another, opaque identifier for applications (or actors) communicating through it. The programming of routing state will produce in a return code

containing the results of the specified operation and associated reason(s) for the result. All of this is critical information to be used for understanding the history of I2RS interactions.

This document specifies requirements, use cases, and describes the information model for traceability attributes that must be collected and logged by I2RS Agents in order to effectively record I2RS interactions. In this context, effective troubleshooting means being able to identify what operation was performed by a specific I2RS Client, what was the result of the operation, and when that operation was performed.

Discussions about the retention of the data logged as part of I2RS traceability, while important, are outside of the scope of this document.

## 2. Terminology and Conventions

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](#)].

The architecture specification for I2RS [[I-D.ietf-i2rs-architecture](#)] defines additional terms used in this document that are specific to the I2RS domain, such as "I2RS Agent", "I2RS Client", etc. The reader is expected to be familiar with the terminology and concepts defined in [[I-D.ietf-i2rs-architecture](#)].

The IP addresses used in the example in this document correspond to the documentation address blocks 192.0.2.0/24 (TEST-NET-1), 198.51.100.0/24 (TEST-NET-2) and 203.0.113.0/24 (TEST-NET-3) as described in [[RFC5737](#)].

## 3. Motivation and Use Cases

As networks grow, so too does the scale and complexity of routing systems. I2RS offers a standard, programmatic interface allowing improved automation and more granular control over an increasingly dynamic routing and signaling state. This ability to automate even complex policy-based controls highlights the need for an equally

scalable traceability function to provide event-level granularity of the routing system compliant with the requirements of I2RS (Section 5 of [[I-D.ietf-i2rs-problem-statement](#)]).

An obvious motivation for I2RS traceability is the need to troubleshoot and identify root-causes of problems in these increasingly complex routing systems. For example, since I2RS is a high-throughput multi-channel, full duplex and highly responsive interface, I2RS Clients may be performing a large number of operations on I2RS Agents concurrently or at nearly the same time and quite possibly in very rapid succession. As these many changes are made, the network reacts accordingly. These changes might lead to a

race condition, performance issues, data loss, or disruption of services. In order to isolate the root cause of these issues it is critical that a network operator or administrator has visibility into what changes were made via I2RS at a specific time.

Some network environments have strong auditing requirements for configuration and runtime changes. Other environments have internal policies to save logging information for operational considerations. These requirements therefore demand that I2RS provides an account of changes made to network element routing systems.

As I2RS becomes increasingly pervasive in routing environments, a traceability model offers significant advantages and facilitates the following use cases:

- o Automated event correlation, trend analysis, and anomaly detection.
- o Trace log storage for offline (manual or tools) analysis.
- o Improved accounting of routing system transactions.
- o Standardized structured data format for writing common tools.
- o Common reference for automated testing and incident reporting.
- o Real-time monitoring and troubleshooting.
- o Enhanced network audit, management and forensic analysis

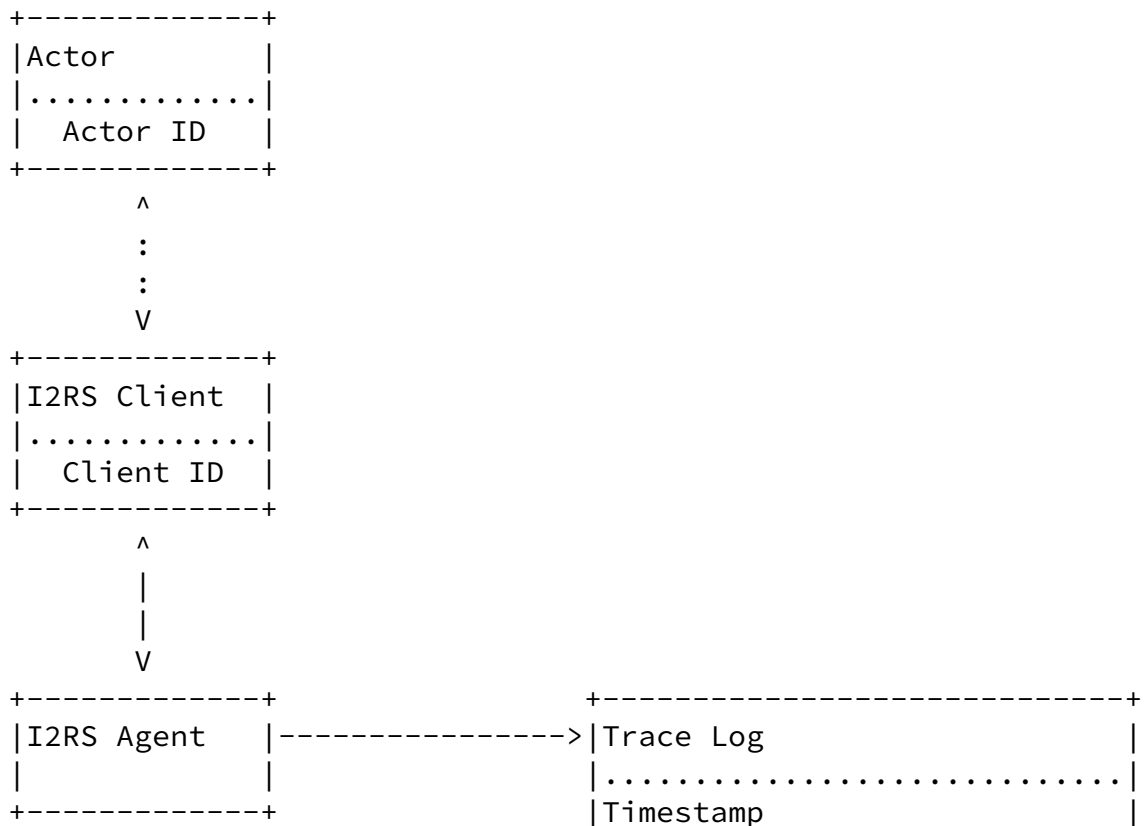
capabilities.

#### 4. Information Model

##### 4.1. I2RS Traceability Framework

This section describes a framework for I2RS traceability based on the I2RS Architecture. Some notable elements on the architecture are highlighted herein.

The interaction between the optional northbound actor, I2RS Client, I2RS Agent, the Routing System and the data captured in the I2RS trace log is shown in Figure 1.



	^		Client ID	
		^	Actor ID	
Operation +		Result Code	Client Address	
Op Data			Operation	
V			Operation Data	
	V		Result Code	
	+-----+		+-----+	
			Routing	
			System	
	+-----+		+-----+	

Figure 1: I2RS Interaction Trace Log Capture

#### 4.2. I2RS Trace Log Mandatory Fields

In order to ensure that each I2RS interaction can be properly traced back to the Client that made the request at a specific point in time, the following information **MUST** be collected and stored by the Agent.

The list below describes the fields captured in the I2RS trace log.

**Timestamp:** The specific time, adhering to [\[RFC3339\]](#) format, at which the I2RS transaction occurred. Given that many I2RS transactions can occur in rapid succession, the use of fractional seconds **MUST** be used to provide adequate granularity.

**Client Identifier:** The I2RS Client identifier used to authenticate the Client to the I2RS Agent.

**Actor Identifier:** This is an opaque identifier that may be known to the Client from a northbound controlling application. This is used to trace the northbound actor driving the actions of the Client. The Client may not provide this identifier to the Agent if there is no external actor driving the Client. However, this field **MUST** be logged. If the Client does not provide an actor ID, then the Agent **MUST** log an empty field.

**Client Address:** This is the network address of the client that connected to the Agent. For example, this may be an IPv4 or IPv6 address. [Note: will I2RS support interactions that have no network address? If so this field will need to be updated.]

Operation: This is the I2RS operation performed. For example, this may be an add route operation if a route is being inserted into a routing table.

Operation Data: This field comprises the data passed to the Agent to complete the desired operation. For example, if the operation is a route add operation, the Operation Data would include the route prefix, prefix length, and next hop information to be inserted as well as the specific routing table to which the route will be added. The operation data can also include interface information. Some operations MAY not provide operation data, and in those cases this field MUST be logged as a NULL string.

Result Code: This field holds the result of the operation. In the case of RIB operations, this MUST be the return code as specified in Section 4 of [[I-D.nitinb-i2rs-rib-info-model](#)]. The operation may not complete with a result code in the case of a timeout. If the operation fails to complete, it MUST still log the attempted operation with an appropriate result code (e.g., a result code indicating a timeout).

#### [4.3.](#) I2RS Trace Log Extensibility and Optional Fields

[NOTE: This section is TBD based on further development of I2RS WG milestones.]

#### [4.4.](#) I2RS Trace Log Syntax

The following describes the trace log information model using Augmented Backus-Naur Form (ABNF) syntax [[RFC5234](#)]:

```
i2rs-trace-log = timestamp client-id actor-id client-addr operation
                operation-data result-code
client-id      = uuid
actor-id      = byte-string
byte-string   = *( %x01-09 / %x0B-0C / %x0E-FF )
               ; any byte except NUL, CR or LF
client-addr   = IP-literal / IPv4address
operation     = ALPHA *( ALPHA / "_" / "-" / "(" / ")" )
```

```
operation-data = *VCHAR
result-code    = 1*( ALPHA / DIGIT / "-" / "_" / "(" / ")" )
```

The ABNF syntax rules <IP-literal>, <IPv4address>, and <sub-delims> are specified in [[RFC3986](#)]. The core rules <DIGIT>, <ALPHA> and <VCHAR> are used as described in [Appendix B of \[RFC5234\]](#).

Here, <uuid> is specified in [[RFC4122](#)]. The idea of using a UUID for the Client identifier ensures the ID is unique not just in the scope of the current I2RS Agent, but across Agents as well. This ensures that two clients that are unaware of each other will not allocate the same Client ID. That does not preclude two Clients acting as one for purposes of high availability from sharing the same UUID as generated by one one of the Clients. [Note: the format of the Client ID has not yet been decided by the I2RS WG. This is subject to change.]

The <timestamp> field is defined in [[RFC3339](#)]. As stated in [Section 4.2](#) the fractional second format MUST be used to provide proper granularity.

The values for <operation>, <operation-data> and <result-code> are suggestions as the protocol has not been defined yet. By making these human-readable (as opposed to opcodes) the log becomes more easily consumable by operators and administrators trying to troubleshoot issues relating to I2RS. Even in cases where the operations or codes might appear as opcodes on the wire, their textual translations MUST be included in the log. The opcodes themselves MAY appear in parentheses after the textual representation.

## [5](#). Examples

Here is a proposed sample of what the fields might look like in an I2RS trace log. This is only an early proposal. These values are subject to change.



Client ID: 5CEF1870-0326-11E2-A21F-0800200C9A66  
Actor ID: com.example.RoutingApp  
Client Address: 192.0.2.2  
Operation: ROUTE\_ADD  
Operation Data: PREFIX 203.0.113.0 PREFIX-LEN 24 NEXT-HOP  
198.51.100.1  
Result Code: SUCCESS(0)

## [6.](#) Operational Guidance

Specific operational procedures regarding temporary log storage, rollover, retrieval, and access of I2RS trace logs is out of scope for this document. Organizations employing I2RS trace logging are responsible for establishing proper operational procedures that are appropriately suited to their specific requirements and operating environment. In this section we only provide fundamental and generalized operational guidelines that are implementation-independent.

### [6.1.](#) Trace Log Creation

The I2RS Agent interacts with the Routing and Signaling functions of the Routing Element. Since the I2RS Agent is responsible for actually making the routing changes on the associated network device, it creates and maintains a log of transactions that can be retrieved to troubleshoot I2RS-related impact to the network.

### [6.2.](#) Trace Log Temporary Storage

The trace information may be temporarily stored either in an in-memory buffer or as a file local to the Agent. Care should be given to the number of I2RS transactions expected on a given agent so that the appropriate storage medium is used and to maximize the effectiveness of the log while not impacting the performance and health of the Agent. [Section 6.3](#) talks about rotating the trace log in order to preserve the transaction history without exhausting Agent or network device resources. It is perfectly acceptable, therefore, to use both an in-memory buffer for recent transactions while rotating or archiving older transactions to a local file.

It is outside the scope of this document to specify the implementation details (i.e., size, throughput, data protection, privacy, etc.) for the physical storage of the I2RS log file. Data retention policies of the I2RS traceability log is also outside the scope of this document.

### [6.3.](#) Trace Log Rotation

In order to prevent the exhaustion of resources on the I2RS Agent or its associated network device, the I2RS trace log implementation MAY choose to implement a configurable rotation schedule. It SHOULD be possible to do file rotation based on either time or size of the current trace log. If file rollover is supported, multiple archived log files SHOULD be supported in order to maximize the troubleshooting and accounting benefits of the trace log.

### [6.4.](#) Trace Log Retrieval

Implementors are free to provide their own, proprietary interfaces and develop custom tools to retrieve and display the I2RS trace log. These may include the display of the I2RS trace log as Command Line Interface (CLI) output. However, a key intention of defining this information model is to establish an implementor-agnostic and consistent interface to collect I2RS trace data. Correspondingly, retrieval of the data should also be made implementor-agnostic.

The following three sections describe potential ways the trace log can be accessed. At least one of these three MUST be used, with the I2RS mechanisms being preferred as they are implementor-independent approaches to retrieving the data.

#### [6.4.1.](#) Retrieval Via Syslog

The syslog protocol [[RFC5424](#)] is a standard way of sending event notification messages from a host to a collector. However, the protocol does not define any standard format for storing the messages, and thus implementors of I2RS tracing would be left to define their own format. So, while the data contained within the syslog message would adhere to this information model, and may be consumable by a human operator, it would not be easily parseable by a machine. Therefore, syslog MAY be employed as a means of retrieving or disseminating the I2RS trace log contents.

#### [6.4.2.](#) Retrieval Via I2RS Information Collection

[Section 6.7](#) of the I2RS architecture [[I-D.ietf-i2rs-architecture](#)] defines a mechanism for information collection. The information collected includes obtaining a snapshot of a large amount of data from the network element. It is the intent of I2RS to make this data available in an implementor-agnostic fashion. Therefore, the I2RS trace log SHOULD be made available via the I2RS information collection mechanism either as a single snapshot or via a

subscription stream.

### [6.4.3.](#) Retrieval Via I2RS Pub-Sub

[Section 6.7](#) of the I2RS architecture [[I-D.ietf-i2rs-architecture](#)] goes on to define a publish-subscribe mechanism for a feed of changes happening within the I2RS layer. I2RS Agents SHOULD support publishing I2RS trace log information to that feed as described in that document. Subscribers would then receive a live stream of I2RS interactions in trace log format and could flexibly choose to do a number of things with the log messages. For example, the subscribers could log the messages to a datastore, aggregate and summarize interactions from a single client, etc. Using pub-sub for the purpose of logging I2RS interactions augments the areas described by [[I-D.camwinget-i2rs-pubsub-sec](#)]. The full range of potential activities is virtually limitless and the details of how they are performed are outside the scope of this document, however.

## [7.](#) IANA Considerations

This document makes no request of IANA.

## [8.](#) Security Considerations

The I2RS trace log, like any log file, reveals the state of the entity producing it as well as the identifying information elements and detailed interactions of the system containing it. The information model described in this document does not itself introduce any security issues, but it does define the set of attributes that make up an I2RS log file. These attributes may contain sensitive information and thus should adhere to the security, privacy and permission policies of the organization making use of the I2RS log file.

It is outside the scope of this document to specify how to protect the stored log file, but it is expected that adequate precautions and security best practices such as disk encryption, appropriately restrictive file/directory permissions, suitable hardening and physical security of logging entities, mutual authentication, transport encryption, channel confidentiality, and channel integrity if transferring log files. Additionally, the potentially sensitive

information contained in a log file SHOULD be adequately anonymized or obfuscated by operators to ensure its privacy.

## [9.](#) References

### [9.1.](#) Normative References

[I-D.ietf-i2rs-architecture]

Clarke, et al.

Expires April 01, 2014

[Page 10]

---

Internet-Draft

I2RS Traceability

September 2013

Atlas, A., Halpern, J., Hares, S., Ward, D., and T. Nadeau, "An Architecture for the Interface to the Routing System", [draft-ietf-i2rs-architecture-00](#) (work in progress), August 2013.

[I-D.ietf-i2rs-problem-statement]

Atlas, A., Nadeau, T., and D. Ward, "Interface to the Routing System Problem Statement", [draft-ietf-i2rs-problem-statement-00](#) (work in progress), August 2013.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

[RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, [RFC 3986](#), January 2005.

[RFC4122] Leach, P., Mealling, M., and R. Salz, "A Universally Unique Identifier (UUID) URN Namespace", [RFC 4122](#), July 2005.

[RFC5234] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, [RFC 5234](#), January 2008.

### [9.2.](#) Informative References

[I-D.camwinget-i2rs-pubsub-sec]

Beck, K., Cam-Winget, N., and D. McGrew, "Using the Publish-Subscribe Model in the Interface to the Routing System", [draft-camwinget-i2rs-pubsub-sec-00](#) (work in progress), July 2013.

[I-D.nitinb-i2rs-rib-info-model]

Bahadur, N., Folkes, R., Kini, S., and J. Medved, "Routing Information Base Info Model", [draft-nitinb-i2rs-rib-info-model-02](#) (work in progress), August 2013.

[RFC3339] Klyne, G., Ed. and C. Newman, "Date and Time on the Internet: Timestamps", [RFC 3339](#), July 2002.

[RFC5424] Gerhards, R., "The Syslog Protocol", [RFC 5424](#), March 2009.

[RFC5737] Arkko, J., Cotton, M., and L. Vegoda, "IPv4 Address Blocks Reserved for Documentation", [RFC 5737](#), January 2010.

#### Authors' Addresses

Clarke, et al.

Expires April 01, 2014

[Page 11]

---

Internet-Draft

I2RS Traceability

September 2013

Joe Clarke  
Cisco Systems, Inc.  
7200-12 Kit Creek Road  
Research Triangle Park, NC 27709  
US

Phone: +1-919-392-2867  
Email: [jclarke@cisco.com](mailto:jclarke@cisco.com)

Gonzalo Salgueiro  
Cisco Systems, Inc.  
7200-12 Kit Creek Road  
Research Triangle Park, NC 27709  
US

Email: [gsalguei@cisco.com](mailto:gsalguei@cisco.com)

Carlos Pignataro  
Cisco Systems, Inc.  
7200-12 Kit Creek Road  
Research Triangle Park, NC 27709  
US

Email: cpignata@cisco.com

Clarke, et al.

Expires April 01, 2014

[Page 12]