

I2RS
Internet-Draft
Intended status: Informational
Expires: November 2, 2016

J. Clarke
G. Salgueiro
C. Pignataro
Cisco
May 1, 2016

Interface to the Routing System (I2RS) Traceability: Framework and
Information Model
draft-ietf-i2rs-traceability-09

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 November 2, 2016.

Copyright Notice

Copyright (c) 2016 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

(<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

1.	Introduction	2
2.	Terminology and Conventions	3
3.	Motivation	3
4.	Use Cases	3
5.	Information Model	4
5.1.	I2RS Traceability Framework	4
5.2.	I2RS Trace Log Mandatory Fields	6
5.3.	End of Message Marker	8
6.	Examples	9
7.	Operational Guidance	9
7.1.	Trace Log Creation	9
7.2.	Trace Log Temporary Storage	10
7.3.	Trace Log Rotation	10
7.4.	Trace Log Retrieval	10
7.4.1.	Retrieval Via Syslog	11
7.4.2.	Retrieval Via I2RS Information Collection	11
7.4.3.	Retrieval Via I2RS Pub-Sub	11
8.	IANA Considerations	12
9.	Security Considerations	12
10.	Acknowledgments	12
11.	References	12
11.1.	Normative References	13
11.2.	Informative References	13
	Authors' Addresses	13

[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 identity for applications communicating through it. The

programming of routing state will produce 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 describes use cases for I2RS traceability. Based on these use cases, the document proposes an information model and reporting requirements to provide for effective recording of 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](#)].

3. Motivation

As networks scale and policy becomes an increasingly important part of the control plane that creates and maintains the forwarding state, operational complexity increases as well. I2RS offers more granular and coherent control over policy and control plane state, but it also removes or reduces the locality of the policy that has been applied to the control plane at any individual forwarding device. The ability to automate and abstract 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

[\[I-D.ietf-i2rs-problem-statement\]](#))).

4. Use Cases

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

Clarke, et al.

Expires November 2, 2016

[Page 3]

Internet-Draft

I2RS Traceability

May 2016

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 policies that require saving logging information for operational or regulatory compliance 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 operations;
- o Standardized structured data format for writing common tools;
- o Common reference for automated testing and incident reporting;
- o Real-time monitoring and troubleshooting;

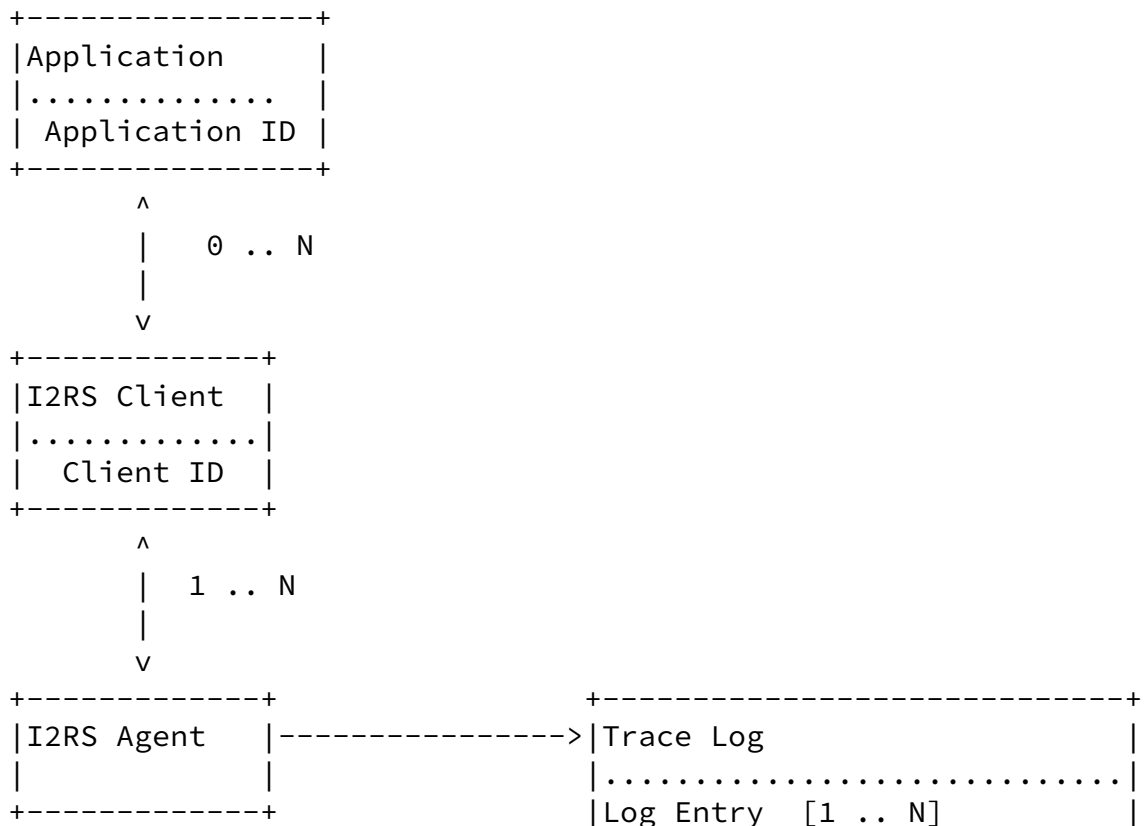
- o Enhanced network audit, management and forensic analysis capabilities.

5. Information Model

5.1. I2RS Traceability Framework

This section describes a framework for I2RS traceability based on the I2RS Architecture. Some notable elements of the architecture are in this section.

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



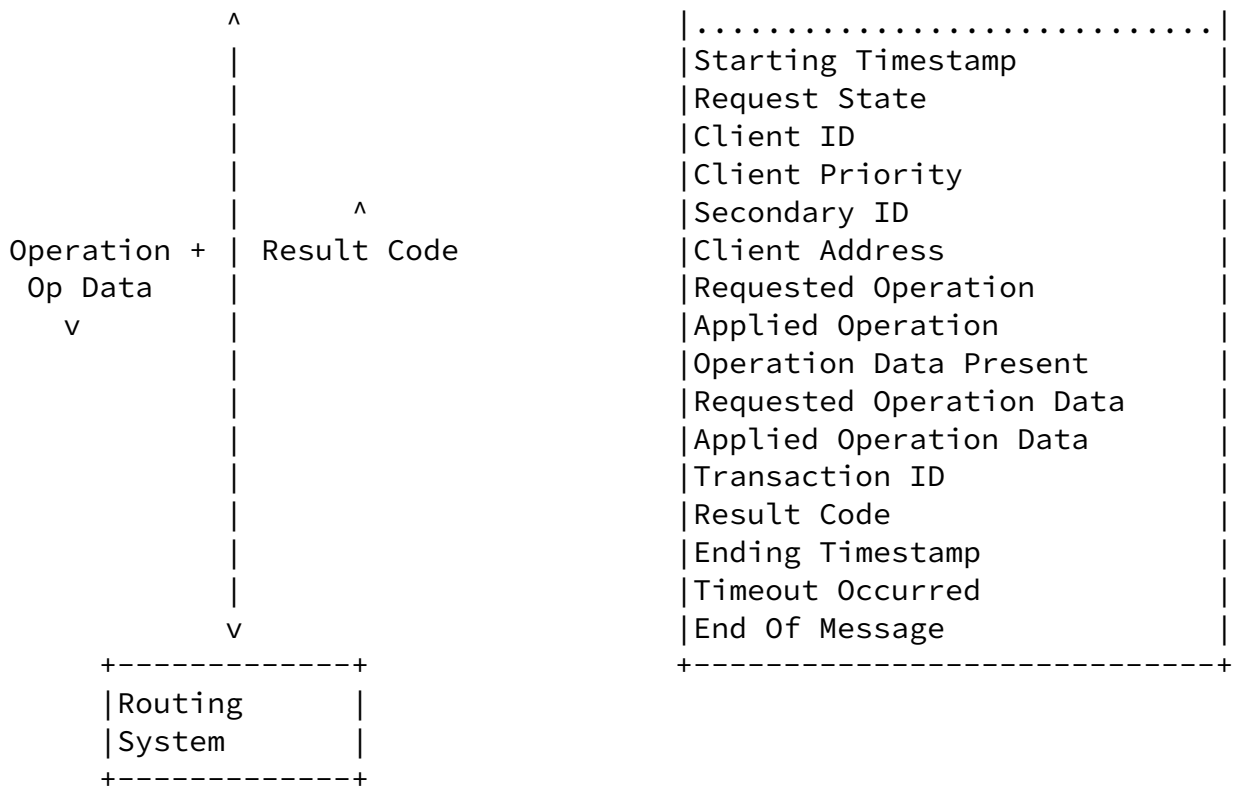


Figure 1: I2RS Interaction Trace Log Capture

5.2. 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.

Entry ID: This is a unique identifier for each entry in the I2RS trace log. Since multiple operations can occur from the same Client at the same time, it is important to have an identifier that can be unambiguously associated to a specific entry.

Starting Timestamp: The specific time at which the I2RS operation entered the specified Request State within the Agent. The time is

passed in the [[RFC3339](#)] format. Given that many I2RS operations can occur in rapid succession, the use of fractional seconds MUST be used to provide adequate granularity. Fractional seconds SHOULD be expressed using human-readable 32-bit second and 32-bit microsecond granularity in second.microsecond format. In the case when the trace log entry specifies a Request State of COMPLETED this time will reflect when the operation was first received by the I2RS Agent.

Request State: The state of the given operation within the I2RS Agent state machine between the specified Starting and Ending Timestamps. This can be one of the following values:

PENDING: The request has been received and queued for processing.

IN PROCESS: The request is currently being handled by the I2RS Agent.

COMPLETED: The request has reached a terminal point.

In the case of the COMPLETED state, the Starting and Ending Timestamps will cover the entire duration of the operation including time spent in the PENDING and IN PROCESS states.

Every state transition MAY be logged unless doing so will put an undue performance burden on the I2RS Agent. However, an entry with Request State set to COMPLETED MUST be logged for all operations.

Client Identity: The I2RS Client identity used to authenticate the Client to the I2RS Agent.

Client Priority: The I2RS Client priority assigned by the access control model that authenticates the Client. For example, this can be set by the NETCONF Access Control Model (NACM) as described in [[RFC6536](#)].

Secondary Identity: This is an opaque identity that may be known to the Client from a northbound controlling application. This is used to trace the northbound application driving the actions of

the Client. The Client may not provide this identity to the Agent if there is no external application driving the Client. However, this field MUST be logged even if the Client does not provide a Secondary Identity. In that case, the field will be logged with an empty value.

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.

Requested Operation: This is the I2RS operation that was requested to be performed. For example, this may be an add route operation if a route is being inserted into a routing table. This may not be the operation that was actually applied to the Agent.

Applied Operation: This is the I2RS operation that was actually performed. This can differ from the Requested Operation in cases where the Agent cannot satisfy the Requested Operation. This field may not be logged unless the Request State is COMPLETED.

Operation Data Present: This is a Boolean field that indicates whether or not addition per-Operation Data is present.

Requested 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. If Operation Data is provided, then the Operation Data Present field MUST be set to TRUE. Some operations may not provide operation data. In those cases, the Operation Data Present field MUST be set to FALSE, and this field MUST be empty. This may not represent the data that was used for the operation that was actually applied on the Agent.

Applied Operation Data: This field comprises the data that was actually applied as part of the Applied Operation. If the Agent cannot satisfy the Requested Operation with the Requested Operation Data, then this field can differ from the Requested

Operation Data. This field may not be logged unless the Request

State is COMPLETED.

Transaction ID: The Transaction Identity represents that this particular operation is part of a long-running I2RS transaction that can consist of multiple, related I2RS operations. Using this value, one can relate multiple log entries together as they are part of a single, overall I2RS operation.

Result Code: This field holds the result of the operation once the Request State is COMPLETED. In the case of RIB operations, this MUST be the return code as specified in Section 4 of [[I-D.ietf-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.

Timeout Occurred: This is a Boolean field that indicates whether or not a timeout occurred in the operation. When this is true, the value of the Ending Timestamp MUST be set to the time the Agent recorded for the timeout occurrence. This field may not be logged unless the Request State is COMPLETED.

Ending Timestamp: The specific time at which the I2RS operation exited the specified Request State within the I2RS Agent. The time is passed in the [[RFC3339](#)] format. Given that many I2RS operations can occur in rapid succession, the use of fractional seconds MUST be used to provide adequate granularity. Fractional seconds SHOULD be expressed using human-readable 32-bit second and 32-bit microsecond granularity in second.microsecond format.

End Of Message: Each log entry SHOULD have an appropriate End Of Message (EOM) indicator. See section [Section 5.3](#) below for more details.

[5.3](#). End of Message Marker

Because of variability within I2RS trace log fields, implementors MUST use a format-appropriate end of message (EOM) indicator in order to signify the end of a particular record. That is, regardless of format, the I2RS trace log MUST provide a distinct way of distinguishing between the end of one record and the beginning of another. For example, in a linear formatted log (similar to syslog) the EOM marker may be a newline character. In an XML formatted log, the schema would provide for element tags that denote beginning and end of records. In a JSON formatted log, the syntax would provide record separation (likely by comma-separated array elements).

[6.](#) Examples

This section shows a sample of what the fields and values could look like.

```
Entry ID: 1
Starting Timestamp: 2013-09-03T12:00:01.21+00:00
Request State: COMPLETED
Client ID: 5CEF1870-0326-11E2-A21F-0800200C9A66
Client Priority: 100
Secondary ID: com.example.RoutingApp
Client Address: 2001:db8:c0c0::2
Requested Operation: ROUTE_ADD
Applied Operation: ROUTE_ADD
Operation Data Present: TRUE
Requested Operation Data: PREFIX 2001:db8:feed:: PREFIX-LEN 64
                          NEXT-HOP 2001:db8:cafe::1
Applied Operation Data: PREFIX 2001:db8:feed:: PREFIX-LEN 64
                          NEXT-HOP 2001:db8:cafe::1
Transaction ID: 2763461
Result Code: SUCCESS(0)
Timeout Occurred: FALSE
Ending Timestamp: 2013-09-03T12:00:01.23+00:00
```

[7.](#) 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.

[7.1.](#) 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 operations that can be retrieved to troubleshoot I2RS-related impact to the network.

[7.2.](#) 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 operations 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. Client requests may not always be processed synchronously or within a bounded time period. Consequently, to ensure that trace log fields, such as "Operation" and "Result Code", are part of the same trace log record it may require buffering of the trace log entries. This buffering may result in additional resource load on the Agent and the network element.

[Section 7.3](#) discusses rotating the trace log in order to preserve the operation history without exhausting Agent or network device resources. It is perfectly acceptable, therefore, to use both an in-memory buffer for recent operations while rotating or archiving older operations 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.

[7.3.](#) Trace Log Rotation

In order to prevent the exhaustion of resources on the I2RS Agent or its associated network device, it is RECOMMENDED that the I2RS Agent implements trace log rotation. The details on how this is achieved are left to the implementation and outside the scope of this document. However, 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.

[7.4.](#) 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 a vendor-agnostic and consistent interface to collect I2RS trace data. Correspondingly, retrieval of the data should also be made vendor-agnostic.

Despite the fact that export of I2RS trace log information could be an invaluable diagnostic tool for off-box analysis, exporting this information MUST NOT interfere with the ability of the Agent to process new incoming operations.

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 vendor-independent approaches to retrieving the data.

7.4.1. 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. Syslog MAY be employed as a means of retrieving or disseminating the I2RS trace log contents.

If syslog is used for trace log retrieval, then existing logging infrastructure and capabilities of syslog [[RFC5424](#)] should be leveraged without the need to define or extend existing formats. For example, the various fields described in [Section 5.2](#) SHOULD be modeled and encoded as Structured Data Elements (referred to as "SD-ELEMENT"), as described in [Section 6.3.1 of \[RFC5424\]](#).

7.4.2. 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.

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

[Section 7.6](#) of the I2RS architecture [[I-D.ietf-i2rs-architecture](#)] goes on to describe notification mechanisms for a feed of changes happening within the I2RS layer. Specifically, the requirements for a publish-subscribe system for I2RS are defined in [[I-D.ietf-i2rs-pub-sub-requirements](#)]. 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. 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.

[8. IANA Considerations](#)

This document makes no request of IANA.

[9. 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.

10. Acknowledgments

The authors would like to thank Alia Atlas for her initial feedback and overall support for this work. Additionally, the authors acknowledge Alvaro Retana, Russ White, Matt Birkner, Jeff Haas, Joel Halpern, Dean Bogdanovich, Ignas Bagdonas, Nobo Akiya, Kwang-koog Lee, Sue Hares, Mach Chen, and Alex Clemm for their reviews, contributed text, and suggested improvements to this document.

11. References

Clarke, et al. Expires November 2, 2016 [Page 12]

Internet-Draft I2RS Traceability May 2016

11.1. Normative References

[I-D.ietf-i2rs-architecture]

Atlas, A., Halpern, J., Hares, S., Ward, D., and T. Nadeau, "An Architecture for the Interface to the Routing System", [draft-ietf-i2rs-architecture-13](#) (work in progress), February 2016.

[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-10](#) (work in progress), February 2016.

[I-D.ietf-i2rs-pub-sub-requirements]

Voit, E., Clemm, A., and A. Prieto, "Requirements for Subscription to YANG Datastores", [draft-ietf-i2rs-pub-sub-requirements-06](#) (work in progress), April 2016.

[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>>.

11.2. Informative References

- [I-D.ietf-i2rs-rib-info-model] Bahadur, N., Kini, S., and J. Medved, "Routing Information Base Info Model", [draft-ietf-i2rs-rib-info-model-08](#) (work in progress), October 2015.
- [RFC3339] Klyne, G. and C. Newman, "Date and Time on the Internet: Timestamps", [RFC 3339](#), DOI 10.17487/RFC3339, July 2002, <<http://www.rfc-editor.org/info/rfc3339>>.
- [RFC5424] Gerhards, R., "The Syslog Protocol", [RFC 5424](#), DOI 10.17487/RFC5424, March 2009, <<http://www.rfc-editor.org/info/rfc5424>>.
- [RFC6536] Bierman, A. and M. Bjorklund, "Network Configuration Protocol (NETCONF) Access Control Model", [RFC 6536](#), DOI 10.17487/RFC6536, March 2012, <<http://www.rfc-editor.org/info/rfc6536>>.

Authors' Addresses

Clarke, et al.

Expires November 2, 2016

[Page 13]

Internet-Draft

I2RS Traceability

May 2016

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

Gonzalo Salgueiro

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

Email: gsalguei@cisco.com

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

Email: cpignata@cisco.com