

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: September 9, 2015

J. Haas, Ed.
Juniper Networks
March 8, 2015

I2RS requirements for netmod/netconf
draft-haas-i2rs-netmod-netconf-requirements-01

Abstract

This document covers requests to the netmod and netconf Working Groups for functionality to support requirements to implement the I2RS architecture.

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 September 9, 2015.

Copyright Notice

Copyright (c) 2015 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.	I2RS Requirements	3
2.1.	Data Store Requirements	3
2.1.1.	A Separate Ephemeral Datastore	3
2.1.2.	Tagged Ephemeral Modules in the Running Datastore . .	4
2.1.3.	Permitting Existing Configuration State to be Made Optionally Ephemeral	4
2.2.	Mutual Authentication Requirements	5
2.2.1.	NETCONF over SSH	5
2.2.2.	NETCONF/RESTCONF over TLS	5
2.3.	Identity, Secondary-Identity Requirements; Priority Requirements; Implications	5
2.3.1.	Identity Requirements	5
2.3.2.	Priority Requirements	6
2.3.3.	Implications of Identities and Priorities on Internal State	6
2.4.	Access Control Model Requirements	6
2.4.1.	Data Store Implications	6
2.4.2.	I2RS Priority	6
2.5.	Connectivity Requirements	6
2.6.	Notification and Subscription Requirements	7
2.7.	Transaction Requirements	7
3.	IANA Considerations	7
4.	Security Considerations	7
5.	Acknowledgements	8
6.	Normative References	8
	Author's Address	9

1. Introduction

The Interface to the Routing System (I2RS) Working Group is chartered with providing architecture and mechanisms to inject into and retrieve information from the routing system. The I2RS Architecture document [I-D.ietf-i2rs-architecture] abstractly documents a number of requirements for implementing the I2RS requirements.

The I2RS Working Group has chosen to use the YANG data modeling language [RFC6020] as the basis to implement its mechanisms.

Additionally, the I2RS Working group has chosen to use the NETCONF [RFC6241] and its similar but lighter-weight relative RESTCONF [I-D.bierman-netconf-restconf] as the protocols for carrying I2RS.

While YANG, NETCONF and RESTCONF are a good starting basis for I2RS, there are some things needed from each of them in order for I2RS to be implemented.

Note that this draft does not attempt to address specific implementation of I2RS requirements that the existing YANG, RESTCONF and NETCONF mechanisms are expected to cover. A separate draft will be issued for the consumption of the I2RS Working Group for such cases.

2. I2RS Requirements

2.1. Data Store Requirements

One of the key mechanisms in I2RS is the ability to inject configuration state into a network element on an ephemeral basis. While at first glance this may seem equivalent to the writable-running datastore in NETCONF, running-config can be copied to a persistent data store, like startup config. The author wishes to prevent any action that would lead to preserving any configuration state entered via the I2RS agent across reboots. If state has to be restored, it should be solely by replay actions from I2RS client via I2RS agent.

A few options for implementing such ephemeral configuration suggest themselves, as do some possible problems with such an implementation:

1. A separate ephemeral datastore. The semantics of this datastore is that all configuration state is known ahead of time to not survive reboot and is not to be copied into persistent storage. Such a datastore could be referenced by NETCONF and RESTCONF using existing semantics, such as "target" and "source".
2. Configuration state in the existing running datastore where the module is "tagged ephemeral".
3. Permitting existing configuration to be optionally configured as ephemeral. As an example, the NETCONF server advertises in its <hello> message if it supports the specified YANG module persistently and/or ephemerally.

2.1.1. A Separate Ephemeral Datastore

The primary advantage of a fully separate datastore is that the semantics of its contents are always clearly ephemeral. It also provides strong segregation of I2RS configuration and operational state from the rest of the system within the network element.

The most obvious disadvantage of such a fully separate datastore is that interaction with the network element's operational or configuration state becomes significantly more difficult. As an example, a BGP I2RS use case would be the dynamic instantiation of a

BGP peer. While it is readily possible to re-use any defined groupings from an IETF-standardized BGP module in such an I2RS ephemeral datastore's modules, one cannot currently reference state from one datastore to another.

For example, XPath queries are done in the context document of the datastore in question and thus it is impossible for an I2RS model to fulfil a "must" or "when" requirement in the BGP module in the standard data stores. To implement such a mechanism would require appropriate semantics for XPath.

2.1.2. Tagged Ephemeral Modules in the Running Datastore

Presume a YANG keyword that flagged an entire module as being ephemeral. In such a case, entire modules could be crafted for I2RS (and other) purposes wherein the configuration state in the module had ephemeral properties. The primary property is that copy operations would not be able to cause the I2RS state to persist.

An obvious issue with this is the muddying of the semantics of existing NETCONF/RESTCONF operations. For example, get-config is expected to return the configuration state for the network element, but the knowledge that the configuration state may not persist is important. This may require alterations to get-config (and similar commands) along with the ambiguity of copy-config not picking up the ephemeral modules.

Providing additional parameters to the various configuration related operations in NETCONF/RESTCONF would likely be required.

2.1.3. Permitting Existing Configuration State to be Made Optionally Ephemeral

In YANG, configuration state is distinguished from operational state using "config true" vs. "config false". One way to implement I2RS state would be to introduce a third option, "config ephemeral", to configuration.

A form of this option was previously discussed in [I-D.rfernando-i2rs-yang-mods]. The suggestion of "config ephemeral" is made instead due to potential non-I2RS interest in this feature at the microphone during the IETF-90 session of netmod in Toronto, Canada.

2.2. Mutual Authentication Requirements

"Mutual authentication between the I2RS Client and I2RS Agent is required. An I2RS Client must be able to trust that the I2RS Agent is attached to the relevant Routing Element so that write/modify operations are correctly applied and so that information received from the I2RS Agent can be trusted by the I2RS Client."

Implementing the mutual authentication requirements for I2RS in each of the underlying protocols and their transports have some implications to be discussed.

2.2.1. NETCONF over SSH

The NETCONF service over SSH is believed to provide the necessary mutual authentication services required by I2RS. Per [RFC6242]: "The identity of the SSH server MUST be verified and authenticated by the SSH client according to local policy before password-based authentication data or any configuration or state data is sent to or received from the SSH server. The identity of the SSH client MUST also be verified and authenticated by the SSH server according to local policy to ensure that the incoming SSH client request is legitimate before any configuration or state data is sent to or received from the SSH client. Neither side should establish a NETCONF over SSH connection with an unknown, unexpected, or incorrect identity on the opposite side."

2.2.2. NETCONF/RESTCONF over TLS

Agent validation of the I2RS client is mandated over TLS in an I2RS context. The client shall also validate the Agent using its server certificate.

2.3. Identity, Secondary-Identity Requirements; Priority Requirements; Implications

2.3.1. Identity Requirements

I2RS requires clients to have an identity. This identity will be used by the Agent authentication mechanism over the appropriate protocol.

I2RS also permits clients to have a secondary identity which may be used for troubleshooting. This secondary identity is an opaque value. [I-D.ietf-i2rs-traceability] provides an example of how the secondary identity can be used for traceability.

If in support of the I2RS priority requirements if it is determined to be necessary to annotate state with the client that added it, the secondary identity SHOULD be included for troubleshooting.

2.3.2. Priority Requirements

To support Multi-Headed Control, I2RS requires that there be a decidable means of arbitrating the correct state of data when multiple clients attempt to manipulate the same piece of data. This is done via a priority mechanism with the highest priority winning. This priority may vary on a per-node or subtree basis based for a given identity.

2.3.3. Implications of Identities and Priorities on Internal State

Given the requirements for I2RS identities and priority arbitration, I2RS configured state must have "meta-data" that includes the identity that caused it to come into being. Agents must also be able to map priority on a particular piece of configuration state vs. the identity provisioning it for arbitration purposes. Such mapping might be represented as part of the "meta-data" or potentially a distinct mapping database of identity vs. priority vs. configuration state. Such a mapping may be implemented using an extension to the NETCONF Access Control Model [RFC6536].

2.4. Access Control Model Requirements

2.4.1. Data Store Implications

As noted above, one of the possible options for implementing the I2RS ephemeral behavior is a separate data store. However, this clashes with Section 3.2 of [RFC6536] which limits itself to the well-known data stores.

2.4.2. I2RS Priority

A likely implementation of priority arbitration would be to extend the NACM model to also contain criteria for I2RS priority.

2.5. Connectivity Requirements

I2RS does not require clients to maintain active communication channels with their agents. Agents thus require the ability to open communication channels back to clients to satisfy previously requested information.

[I-D.ietf-netconf-call-home] describes a mechanism by which NETCONF may "phone home" using SSH and TLS.

While NETCONF notifications currently permit a different client to establish a session to an agent specifically for notification purposes, the I2RS use case typically expects that provisioning of notifications is centrally managed and that systems receiving the notifications should not need to be individually to be provisioned.

2.6. Notification and Subscription Requirements

[I-D.ietf-i2rs-pub-sub-requirements] has been published and contains the requirements for I2RS for publication and subscription services.

2.7. Transaction Requirements

Each transaction should be treated as atomic and providing full functionality. If the configuration change is not functionally complete, then the transaction should fail and be rolled back (rollback 0). Example, I2RS agents wants to configure BGP:

```
routing-options {
  autonomous-system autonomous-system;
}
protocols {
  bgp {
    group group-name {
      peer-as autonomous-system;
      type type;
      neighbor address;
    }
  }
}
```

If a statement like neighbor address is missing or is mis-formatted, like 300.127.5.23, configuration is not functional, transaction should fail and rollback 0 should be performed by the I2RS agent on the ephemeral config store. If the neighbor address is in the transaction, but the address is not reachable or similar, transaction is accepted, but notification will be sent that BGP peering can't be established.

3. IANA Considerations

This document introduces no new considerations to IANA.

4. Security Considerations

TBD

5. Acknowledgements

This document is an attempt to distill length conversations on the I2RS mailing list for an architecture that was for a long period of time a moving target. Some individuals in particular warrant specific mention for their extensive help in providing the basis for this document:

- o Alia Atlas
- o Andy Bierman
- o Martin Bjorklund
- o Dean Bogdanavich
- o Rex Fernando
- o Joel Halpern
- o Susan Hares
- o Thomas Nadeau
- o Juergen Schoenwaelder
- o Kent Watsen

6. Normative References

[I-D.bierman-netconf-restconf]

Bierman, A., Bjorklund, M., Watsen, K., and R. Fernando, "RESTCONF Protocol", draft-bierman-netconf-restconf-04 (work in progress), February 2014.

[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-05 (work in progress), July 2014.

[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-00 (work in progress), March 2015.

- [I-D.ietf-i2rs-traceability]
Clarke, J., Salgueiro, G., and C. Pignataro, "Interface to the Routing System (I2RS) Traceability: Framework and Information Model", draft-ietf-i2rs-traceability-02 (work in progress), March 2015.
- [I-D.ietf-netconf-call-home]
Watson, K., "NETCONF Call Home", draft-ietf-netconf-call-home-00 (work in progress), September 2014.
- [I-D.rfernando-i2rs-yang-mods]
Fernando, R., Pals, P., Madhayyan, M., and A. Clemm, "YANG modifications for I2RS", draft-rfernando-i2rs-yang-mods-00 (work in progress), February 2013.
- [RFC6020] Bjorklund, M., "YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)", RFC 6020, October 2010.
- [RFC6241] Enns, R., Bjorklund, M., Schoenwaelder, J., and A. Bierman, "Network Configuration Protocol (NETCONF)", RFC 6241, June 2011.
- [RFC6242] Wasserman, M., "Using the NETCONF Protocol over Secure Shell (SSH)", RFC 6242, June 2011.
- [RFC6536] Bierman, A. and M. Bjorklund, "Network Configuration Protocol (NETCONF) Access Control Model", RFC 6536, March 2012.

Author's Address

Jeffrey Haas (editor)
Juniper Networks

Email: jhaas@juniper.net