I2RS working group
Internet-Draft
Intended status: Standards Track
Expires: December 19, 2014

S. Hares
Huawei
S. Brim
Consultant
N. Cam-Winget
Cisco
J. Halpern
Ericcson
D. Zhang
Q. Wu
Huawei
A. Abro
S. Asadullah
Cisco
J. Halpern
Ericcson
E. Yu
Cisco

June 17, 2014

I2RS Security Considerations
draft-hares-i2rs-security-01

Abstract

This presents an expansion of the security architecture found in 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 December 19, 2014.

Table of Contents

1. Introduction

The Interface to the Routing System (I2RS) provides read and write access to the information and state within the routing process and configuration process (as illustrated in the diagram in the architecture document within routing elements. The I2RS client

[I-D.ietf-i2rs-architecture] interacts with one or more I2RS agents to collect information from network routing systems. This security architecture expands on the security issues involved in the I2RS protocol's message exchange between the I2RS client and the I2RS agent which are described in [I-D.ietf-i2rs-architecture].

2. Definitions

This document utilizes the definitions found in the following drafts: [RFC4949], and [I-D.ietf-i2rs-architecture].

Specifically, this document utilizes the following definitions:

Access control

[RFC4949] describes access control as: a) protection of system resources against unauthorized access, b) process controlled by a security policy that permits access only by authorized entities (users, programs, process, or others) according to that policy, c) preventing unauthorized use of resource, d) using human controls to identify or admit properly authorized people to a SCIF, and e) limitations on between subjects and objections in a system. I2RS focuses on role-based access control (RBAC).

Authentication

[RFC4949] describes authentication as the process of verifying (i.e., establishing the truth of) an attribute value claimed by or for a system entity or system resource. Authentication has two steps: identify and verify.

Data Confidentiality

[RFC4949] describes data confidentiality has having two properties: a) data is not disclosed to system entities unless they have been authorized to know, and b) data is not disclosed to unauthorized individuals, entities or processes. The key point is that confidentiality implies that the originator has the ability to authorize where the information goes. Confidentiality is important for both read and write scope of the data.

Data confidentiality service

[RFC4949] also describes data confidentiality service as a security service that protects data against unauthorized disclosure. Please note that an operator can designate all people are authorized to view a piece of data which would mean a data confidentiality service would be essentially a null function.

Data Privacy

[RFC4949] describes data privacy as a synonym for data confidentiality. This I2RS document will utilize data privacy as a synonym for data confidentiality.

Mutual Authentication

[RFC4949] implies that mutual authentication exists between two interacting system entities. Mutual authentication in I2RS implies that both sides move from a state of mutual suspicion to mutually authenticated communication after having been identified and validated.

Mutual Suspicion

[RFC4949] defines mutual suspicion as a state that exists between two interacting system entities in which neither entity can trust the other to function correctly with regard to some security requirement.

Role

[RFC4949] describes role as a job function or employment position to which people or other system entities may be assigned in a system. In the I2RS interface, the I2RS agent roles relate to the roles that the I2RS client is utilizing. In the I2RS interface, the I2RS client negotiation is over the client's ability to access resources made available through the agent's particular role. Please refer to Figure 2 below. Existing work includes IETF work in ABFAB and HTTP related SAML work.

Role-based Access control

[RFC4949] describes role-based access control as an identity-based access control wherein the system entities that are identified and controlled are functional positions in an organization or process. Within [RFC4949] five relationships are discussed: 1) administrators to assign identities to roles, 2) administrators to assign permissions to roles, 3) administrators to assign roles to roles, 4) users to select identities in sessions, and 5) users to select roles in sessions. This document discusses I2RS use of Roles as Identities+Scope+Access where scope is the portion of the routing tree, and access is permissions to read or write (or both). Figure 1 below provides [RFC4949] the security view roles and assignments (page 254). Figure 2 provides the same conceptual view of role-based access control applied to I2RS's Combination of

roles and identities that allow read, write, or read-write access to I2RS agent functions.

Role hierarchy or Permissions inheritance

[RFC4949] describes the hierarchy of roles and identities in role-based access control shown in Figure 1 and described above. I2RS will used role-based access control as defined above, and shown in Figure 2.

Role certificate

[RFC4949] describes a role certificate as an organizational certificate that is issued to a system entity that is a member of the set of users that have identities that are assigned to the same role.

Security audit trail

[RFC4949] (page 254) describes a security audit trail as a chronological record of system activities that is sufficient to enable the reconstruction and examination of the sequence environments and activities surrounding or leading to an operation, procedure, or event in a security-relevant transaction from inception to final results. To apply this to the I2RS system, this implies that the processes on the I2RS client-I2RS Agent protocol and related actions on the I2RS-Agent can record a set of activity that will allow the reconstruction and examination of the sequence of environments and activities around actions caused by the I2RS protocol data streams.

I2RS integrity

The data transfer as it is transmitted between client and agent cannot be modified by unauthorized parties without detection.

The following diagram is a variation of the [RFC4949] diagram on role-based security, and provides the context for the assumptions of security on the role-based work.

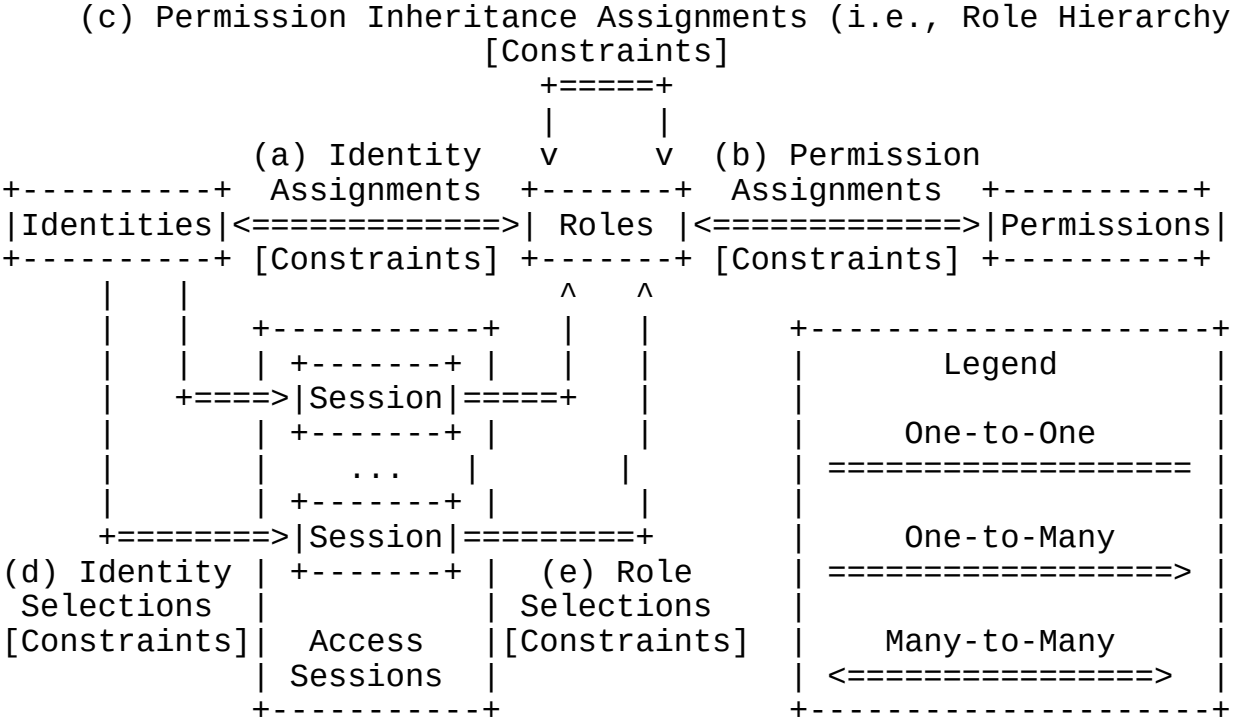

Figure 1 - Security definition of Role inheritance

3. Security Issues

The security for the I2RS protocol utilizes the role based access security for the I2RS clients access to the I2RS agent's data (read/ write). The I2RS [I-D.ietf-i2rs-architecture] treats the agent's notification stream or publication stream as a pre-authorized read. This security consideration document examines the major points:

I2RS roles and identities

This section looks at how I2RS roles and identities created by [I-D.ietf-i2rs-architecture], how I2RS model derived from the security model of role-based access control matches the [I-D.ietf-i2rs-architecture], and how Identities and roles get distributed?

Data Security

The data security section looks at incidents when the I2RS data stream will need confidentiality and message integrity, transport security, how role-based access control of I2RS data impacts the I2RS Information Model and Data Model design, and light weight clients who work without confidentiality.

Transport Requirements for Multiple data stream connections in I2RS

[I-D.ietf-i2rs-architecture] allows multiple data streams across one or more transports. This section examines the security issues surrounding those streams.

Subsequent sections will look how auditing, tracing and deployment scenarios impact the I2RS protocol.

3.1. Security roles and Identities for the I2RS client and I2RS Agent

All I2RS clients and I2RS agents MUST have at least one unique identifier that uniquely identifies each party. The I2RS protocol MUST utilize these identifiers for mutual identification of the client and agent. An I2RS agent, upon receiving an I2RS message from a client, must confirm that the client has a valid identity. The client, upon receiving an I2RS message from an agent, must confirm the I2RS identity.

The distribution of security identity is taken up in the section below. To provide context for that discussion let us look at how I2RS roles are linked to that identity/identifier.

Role = identity + routing tree + Read/Write/R-W

Role security for an agent combines agent identity plus the potential read scope plus the potential write scope. The potential read scope is the routing attributes/variables within a data model (for example BGP peer information) or a set of data models (RIB Data Mode and the BGP peer information) that an agent may potential read. A notification or an event stream is considered a set of read scope data sent via different methodology. A write scope is something the client may write.

Role security exists even if multiple transport connections are being used between the I2RS client and I2RS agent (per [I-D.ietf-i2rs-architecture]). These transport message streams may start/stop without affecting the existence of the client/agent data exchange. TCP supports a single stream of data. SCTP [RFC4960] provides security for multiple streams plus end-to-end transport of data.

(Editor: Additional WG discussion will need to focus on how different deployments impact the transport layers, and the messages sizes (E.g. UDP’s limited size). Use case descriptions will guide this discussion.)

3.1.1. I2RS Role-Based Access Control

Figure 2 show a model of the I2RS role-based access control environment. This model is a variation of the [RFC4949] diagram on role-based security shown in Figure 1. Portions of this model are outside the scope of the I2RS protocol, but are part of the deployment environment of the I2RS protocol. For example, the I2RS identity repository is a logical construct of an entity that keeps all the identities. This logical entity may be implemented in deployments of I2RS in many ways. One simple way is the administrator securely transferring a file with identities and Roles to the client and agent. An automated way may be seen within the security identity distribution protcools in the IETF (AAA, ABFAB, etc). The important point is the Roles (Identity + Rib-portion + Scope (Read, Write, R/W) is passed within the I2RS environment in an manner consistent to the logical constrains in this model.

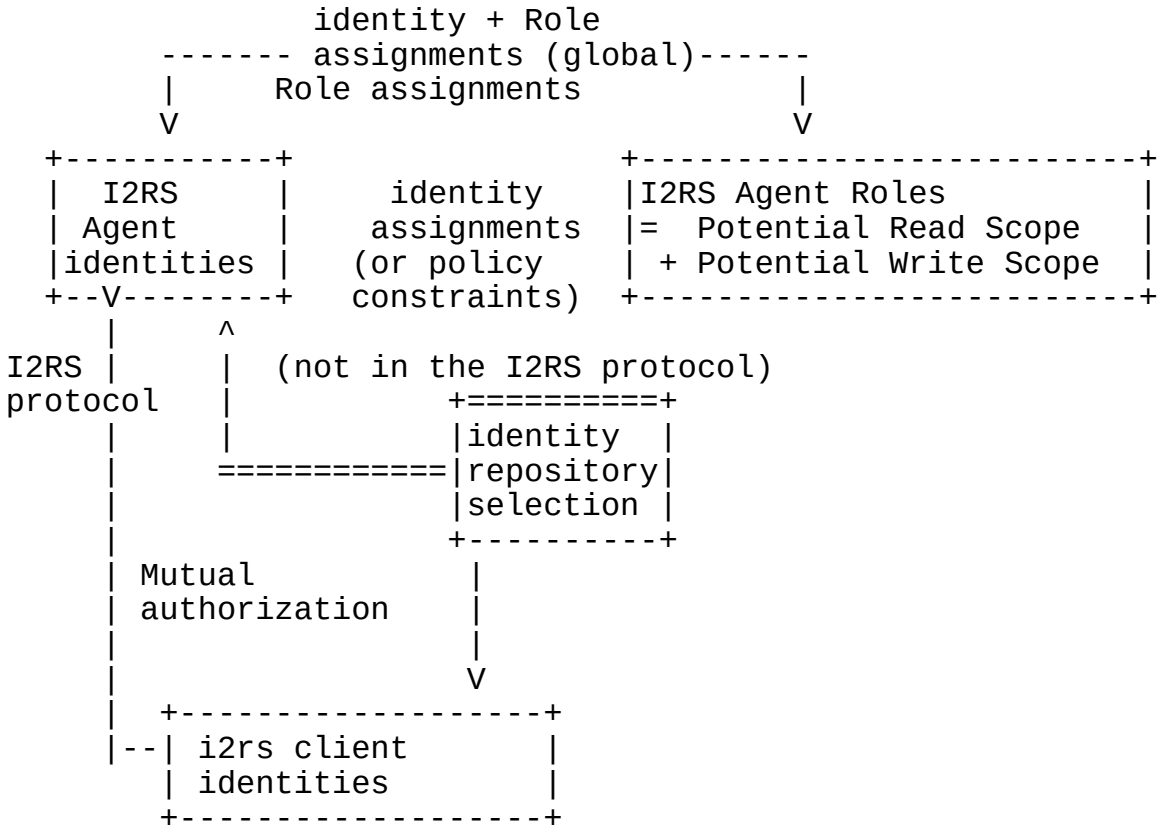

Figure 2 - I2RS Role Based Access Model

3.1.2. Identities

This document suggests that identity distribution and the loading of these identities into I2RS agent and I2RS Client occur outside the I2RS protocol. The I2RS protocol SHOULD assume some mechanism (IETF or private) will distribute identities and that the I2RS client/agent will load the identities prior to the I2RS protocol establishing a connection between I2RS client and I2RS agent.

Each Identity will be linked (via internal policy) to one or more roles. The context of the I2RS client-agent communication is based on an role which may/may not require message confidentiality, message integrity protection, or replay attack protection.

I2RS clients may be used by multiple applications to configure routing via I2RS agents, receive status reports, turn on the I2RS audit stream, or turn on I2RS traceability. An I2RS client software could arrange to store multiple secure identities and use the identity to insure that the "Status-only" application process only uses the client identity for status notification no matter what role that identity takes on. Multiple identities provide some secondary level support for the application-client, but may grow the number of identities. The multiple identities per client could also be used for multiple levels of security for the data passed between an I2RS client and agent as either: a) confidential, b) authorized with message integrity protection, c) authorized without message integrity protection, and or d) no protection. See the section below for additional discussions on these options.

Editor’s note: The WG needs to discuss the scaling properties of the out of band establishment of identities (that is outside the I2RS protocol).

3.2. I2RS Data Security

I2RS data security involves determining of the I2RS client to I2RS agent data transfer needs to be confidential, or have message integrity, or support an end-to-end integrity (in the case of stacked clients). This section discuss the consideration of I2RS data security.

It is assumed that all I2RS data security mechanisms used for protecting the I2RS packets needs to be associated with proper key management solutions. A key management solution needs to guarantee that only the entities having sufficient privileges can get the keys to encrypt/decrypt the sensitive data. In addition, the key management mechanisms need to be able to update the keys before they

have lost sufficient security strengths, without breaking the connection between the agents and clients.

3.2.1. Data Confidentiality Requirements

In a critical infrastructure, certain data within routing elements is sensitive and R/W operations on such data must be controlled in order to protect its confidentiality. For example, most carriers do not want a router's configuration and data flow statistics known by hackers or their competitors. While carriers may share peering information, most carriers do not share configuration and traffic statistics. To achieve this, access control to sensitive data needs to be provided, and the confidentiality protection on such data during transportation needs to be enforced.

It is normal to protect the confidentiality of the sensitive data during transportation by encrypting them. Encryption obscures the data transported on the wire and protects them against eavesdropping attacks. Because the encryption itself cannot guarantee the integrity or fresh of data being transported, in practice, confidentiality protection is normally provided with integrity protection.

3.2.2. Message Integrity Requirements

An integrity protection mechanism for I2RS should be able to ensure 1) the data being protected are not modified without detection during its transportation and 2) the data is actually from where it is expected t come from 3) the data is not repeated from some earlier interaction of the protocol. That is, when both confidentiality and integrity of data is properly protected, it is possible to ensure that encrypted data are not modified or replayed without detection.

As a part of integrity protection, the replay protection approaches provided for I2RS must consider both online and offline attackers, and have sufficient capability to deal with intra connection and inter-connection attacks. For instance, when using symmetric keys, sequence numbers which increase monotonically could be useful to help in distinguishing the replayed messages, under the assistance of signatures or MACs (dependent on what types of keys are applied). In addition, in the cases where only offline attacker is considered, random nonce could be effective.

3.2.3. End-to-End Data Integrity: Data or Transport

The I2RS protocol is concerned with I2RS client-agent exchange. End-to-end confidentiality requires at least transport layer security. In a simple case of a I2RS Client to a single I2RS agent transfer,

the I2RS client puts the data in to the secure transport message and the I2RS agent takes it out of the transport message.

In the case of a stacked client where the I2RS-client1 talks to a I2RS-agent1-I2RS-client2, the data that transfers between the I2RS-agent-1 and I2RS-client-2 is outside the scope of the I2RS protocol. However, it is critical if this mechanism is used for fan-out of read/write commands to agents that the end-to-end data has data integrity.

Editor question: Should I2RS have the optional capability to support end-to-end data integrity?

3.3. Role-Based Access Control of I2RS data

I2RS protocol uses the I2RS Role (Identity + Access (Read, Write, or Read/Write) to control access to the I2RS data. The impact of I2RS role-based security on I2RS data models is that certain portions of an I2RS data models may require:

- confidentiality - which requires a) mutual authentication, b) encryption, and c) message integrity protection with its associated replay protection,
- Message integrity protection - which requires mutual authentication, message integrity with replay protection,
- mutual authentication only, or
- no authentication.

Therefore, creators of I2RS Information Models (IM) and I2RS Data Models (DM) may want to consider the following factors:

- Does the client using this data model care if the agent is valid?
- Does the agent responding to this data model care if the client is valid?
- Does the client-agent exchange require mutual authentication for all of the data model or some?
- Does the client/agent care what operations are done? (secure communications)
- Does the client and agent care about protection - either 1) confidentiality or 2) replay with integrity?

- Are there other security issues unique to this Informational Model (IM) or Data Model (DM)

3.4. Impact of Data Confidentiality inclusion/exclusion in the I2RS Protocol

Confidentiality of role implies the following:

- a requirement for confidentiality of I2RS routing tree scope (portion) in I2RS client-agent communication;
- I2RS client and I2RS agent mutually validate identities; and
- encryption is supported in the I2RS protocol.

Mutual validation of client and agent's identities means that both:

- The I2RS client knows the I2RS agent has a valid identity, and that the I2RS agent has agreed that the I2RS client has a valid identity; and
- The I2RS agent knows that the I2RS client has a valid Identity, and the the I2RS client has agreed that the I2RS agent has a valid identity.

I2RS WG has indicated some I2RS client-agent message exchanges will not need encrypt data to obscure the data. If this is so, then the I2RS designers must understand if their data will be encrypted or sent without encryption. Information Model (IM) and Data Model (DM) creators must discuss determine the following:

- I2RS Client to Agent: Is encryption a recommendation or requirement?
- If it is a recommendation, must the I2RS agent/client support encryption but only use it for certain roles (portions of the tree with read/write scope)? If there are multiple channels for transporting data, one role could be operating without encryption on one part of the tree, and another role could be operating with encryption on another part of the tree.
- Does the Informational Model (IM) and Data Model (DM) make assumptions that would allow security attacks using the unencrypted data?

3.5. Transport requirements

The architecture provides the ability to have multiple transport sessions providing protocol and data communication between the I2RS Agent and the I2RS client. The document does not try to specify the protocols for securing I2RS packets, but provides considersations in choosing a transport protocol. These transports can be TCP or secure (SCTP) or a TLS based. If we use TLS based transports, we can use TLS over UDP (DTLS) or SSL with with TLS plus extensions.

The following are questions to address regarding the transport:

- Do we have mandatory-to-implement transport protocols?
- Will the association of I2RS Roles with transport protocols need to be configured in the I2RS client and I2RS agent?
- Do we allow the I2RS agent/client to automatically establish transport sessions to publish statistics for notifications/ subscriptions?
- Is a publishing broker feasible or does that cause security issues?

4. Audit-able Data streams

This section discusses data streams which have a security audit trail (see definitions) for the I2RS Client to I2RS Agent interactions. The I2RS Discussion group suggested that audit data streams are:

- a tracing of changes sent to a separate streams, and
- a portion of the data selected by policy
- turned on/off via I2RS protocol

I2RS is not inventing a new audit protocol as many protocols (syslog) are available to be used. Verifying audit stream data is outside the I2RS protocol, but those designing the IM and DMs with audit stream capability need to provide the appropriate hooks such as: on/off action, data selection, and protocol (for example syslog) that the I2RS Agent (or I2RS routing system) sends the audit data upon.

Agent audit trail could be the logging of what variables written by which client (identified by client ID) on behalf of a reported application (identified by the ID of the application). The audit stream turned on by the I2RS Agent may need to pass both the client ID and the application ID to the audit stream.

Out of scope for this work is the ability to audit the application to I2RS-Client interfaces, or the I2RS Agent to I2RS routing system.

Editor: Questions still to be answered:

- Is support for audit stream a requirement for all I2RS agents or an option dependent on the role which is dependent on the IM/DM (info and data models)?

- How does the filtering of event data impact the audit process? For example if BGP event changes are only taken from 50 out of 300 BGP peers, does this stop any ability to audit the session? Or if the read filters only watch for key prefixes to be received on a specific set of interfaces, does this stop the ability to audit?

- How do you handle filtering of reads/notifications by I2RS policy and auditing? If the I2RS client asks to read a IM/DM tree portion via a Role but the that read data requested of I2RS Agent is filtered before sending to client, how is this handled in the auditing protocols?

5. Impact of Traceability

The draft [I-D.clarke-i2rs-traceability] provides an IM for the following use cases:

- Automated event correlations, trend analysis, and anomaly detection

- trace log storage

- improved accounting of routing system transactions

- Standardized structure data format for writing common tools

- real-time monitoring and troubleshooting

- enhanced network audit, management and forensic analysis capabilities

The operational guidance in the traceability IM includes creation of an I2RS log that is stored in a temporary storage, rotated, and retrieved via syslog, I2RS "snap-shot" available as one bulk snapshot or subscription, and in a I2RS publish-subscribe stream.

The security issues of the traceability log data sent to syslog are equivalent to the auditable data stream security issues covered in the previous section. The one-bulk snapshot data model and publish/

subscription model contain the same issues considered in the basic read functions described above. The traceability log issues beyond this are implementation or transport protocol issues regarding scale.

6. Deployment issues

This section provides consideration for the deployment issues around stacked I2RS clients. This section only has questions for now, and will be added to in future drafts.

6.1. Stacked I2RS Agent-Clients in Broker topologies

The [I-D.ietf-i2rs-architecture] describes a broker function that can be used in the topology server use case. The general concept for such a deployment would allows the following hierarchical scenario:

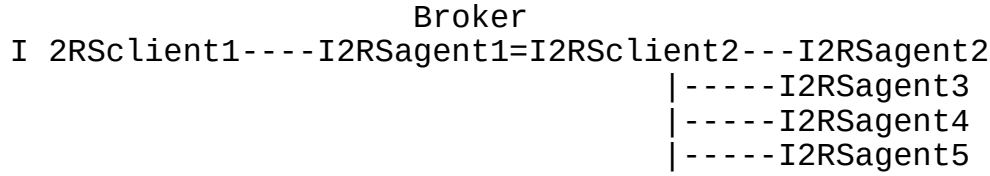

Figure 3

Editor: The implications of this deployment scenario will be added to this draft. For now we have the following questions:

- Does Stacked I2rs agent/client require end-to-end security?
- Does this scenario bring unique security issues?
- Is this scenario outside the I2RS venue? If
- If it is scope, do we need to alter the diagrams within the architecture document? If so, how would we re-write the diagrams.

7. Acknowledgement

The authors would like to thank Wes George, Ahmed Abro, Qin Wu, Eric Yu, Alia Atlas, and Jeff Haas for their wonderful contributions to our discussion discussion.

8. IANA Considerations

This draft includes no request to IANA.

9. Security Considerations

This is a document about security architecture beyond the consideration for I2RS. Additional security definitions will be added in this section.

10. Informative References

[I-D.clarke-i2rs-traceability]
Clarke, J., Salgueiro, G., and C. Pignataro, "Interface to the Routing System (I2RS) Traceability: Framework and Information Model", draft-clarke-i2rs-traceability-02 (work in progress), June 2014.

[I-D.hares-i2rs-info-model-policy]
Hares, S. and W. Wu, "An Information Model for Network policy", draft-hares-i2rs-info-model-policy-02 (work in progress), March 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-03 (work in progress), May 2014.

[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-03 (work in progress), June 2014.

[I-D.ietf-i2rs-rib-info-model]
Bahadur, N., Folkes, R., Kini, S., and J. Medved, "Routing Information Base Info Model", draft-ietf-i2rs-rib-info-model-03 (work in progress), May 2014.

[I-D.ji-i2rs-usecases-ccne-service]
Ji, X., Zhuang, S., Huang, T., and S. Hares, "I2RS Use Cases for Control of Forwarding Path by Central Control Network Element (CCNE)", draft-ji-i2rs-usecases-ccne-service-01 (work in progress), February 2014.

[I-D.keyupate-i2rs-bgp-usecases]
Patel, K., Fernando, R., Gredler, H., Amante, S., White, R., and S. Hares, "Use Cases for an Interface to BGP Protocol", draft-keyupate-i2rs-bgp-usecases-02 (work in progress), June 2014.

[I-D.white-i2rs-use-case]
White, R., Hares, S., and A. Retana, "Protocol Independent Use Cases for an Interface to the Routing System", draft-white-i2rs-use-case-05 (work in progress), June 2014.

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

[RFC4785] Blumenthal, U. and P. Goel, "Pre-Shared Key (PSK) Ciphersuites with NULL Encryption for Transport Layer Security (TLS)", RFC 4785, January 2007.

[RFC4949] Shirey, R., "Internet Security Glossary, Version 2", RFC 4949, August 2007.

[RFC4960] Stewart, R., "Stream Control Transmission Protocol", RFC 4960, September 2007.

Authors' Addresses

Susan Hares
Huawei
7453 Hickory Hill
Saline, MI 48176
USA

Email: shares@ndzh.com

Scott Brim
Consultant

Email: scott.brim@gmail.com

Nancy Cam-Winget
Cisco

Email: ncamwing@cisco.com

Joel Halpern
Ericcson

Email: joel.halpern@ericsson.com

DaCheng Zhang
Huawei

Email: zhangdacheng@huawei.com

Qin Wu
Huawei

Email: bill.wu@huawei.com

Ahmed Abro
Cisco

Email: aabro@cisco.com

Salman Asadullah
Cisco

Email: sasad@cisco.com

Joel Halpern
Ericcson

Email: joel.halpern@ericsson.com

Eric Yu
Cisco

Email: eyu@cisco.com