

NETCONF Working Group
Internet-Draft
Updates: 4253 (if approved)
Intended status: Standards Track
Expires: January 02, 2015

K. Watsen
Juniper Networks
July 2014

NETCONF Call Home using SSH
draft-ietf-netconf-reverse-ssh-06

Abstract

This document presents a technique for a NETCONF server to request that a NETCONF client initiates a SSH connection to the NETCONF server, a technique referred to as 'call home'. Call home is needed to support deployments where the NETCONF client is otherwise unable to initiate a SSH connection to the NETCONF server directly.

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 January 02, 2015.

Copyright Notice

Copyright (c) 2014 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. Requirements Terminology	2
2. Introduction	2
2.1. Applicability Statement	3
2.2. Update to RFC 4253	3
2.3. Draft Naming	3
3. Benefits to Device Management	3
4. Protocol	5
5. SSH Server Identification and Verification	5
6. Device Configuration	6
7. Security Considerations	7
8. IANA Considerations	8
9. Acknowledgements	8
10. References	8
10.1. Normative References	8
10.2. Informative References	9
Appendix A. Change Log	9
A.1. 05 to 06	9
A.2. 04 to 05	10
A.3. 03 to 04	10
A.4. 02 to 03	11
A.5. 01 to 02	11
A.6. 00 to 01	11

1. Requirements Terminology

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 RFC 2119 [RFC2119].

2. Introduction

This document presents a technique for a NETCONF server to request that a NETCONF [RFC6241] client initiates a SSH [RFC4251] connection to the NETCONF server, a technique referred to as 'call home'. Call home is needed to support deployments where the NETCONF client is otherwise unable to initiate a SSH connection to the NETCONF server directly.

2.1. Applicability Statement

The techniques described in this document are suitable for network management scenarios such as the ones described in section 3. However, these techniques SHOULD only be used for a NETCONF server to initiate a connection to a NETCONF client, as described in this document.

The reason for this restriction is that different protocols have different security assumptions. The NETCONF over SSH specification requires NETCONF clients and servers to verify the identity of the other party before starting the NETCONF protocol (section 6 of [RFC6242]). This contrasts with the base SSH protocol, which does not require programmatic verification of the other party (section 9.3.4 of [RFC4251] and section 4 of [RFC4252]). In such circumstances, allowing the SSH server to contact the SSH client would open new vulnerabilities. Therefore, any use of call home with SSH for purposes other than NETCONF will need a thorough, contextual security analysis.

2.2. Update to RFC 4253

This document updates the SSH Transport Layer Protocol [RFC4253] only by removing the restriction in Section 4 (Connection Setup) of [RFC4252] that the SSH Client must initiate the transport connection. Security implications related to this change are discussed in Security Considerations (Section 7).

2.3. Draft Naming

(this section should be removed if this draft becomes an RFC)

This draft's name includes the string "reverse-ssh", and yet currently nowhere in this draft is there any reference to reversing SSH. This apparent omission comes from the -05 edit of this draft, where "Reverse SSH" was changed to "Call Home" throughout. If this draft becomes an RFC, its name would no longer contain the obsolete "reverse-ssh" reference, thus self-correcting this inconsistency.

3. Benefits to Device Management

The SSH protocol is nearly ubiquitous for device management, as it is the transport for the command-line applications 'ssh', 'scp', and 'sftp' and is the required transport for the NETCONF protocol [RFC6241]. However, all these SSH-based protocols expect the network element to be the SSH server.

NETCONF over SSH Call Home enables the network element to consistently be the SSH server regardless of which peer initiates the underlying TCP connection. Maintaining the role of SSH server is both necessary and desirable. It is necessary because SSH channels and subsystems can only be opened on the SSH server. It is desirable because it conveniently leverages infrastructure that may be deployed for host-key verification and user authentication.

Call home is useful for both initial deployment and on-going device management and may be used to enable any of the following scenarios:

- o The network element may proactively call home after being powered on for the first time to register itself with its management system.
- o The network element may access the network in a way that dynamically assigns it an IP address and it doesn't register its assigned IP address to a mapping service.
- o The network element may be configured in "stealth mode" and thus doesn't have any open ports for the management system to connect to.
- o The network element may be deployed behind a firewall that doesn't allow SSH access to the internal network.
- o The network element may be deployed behind a firewall that implements network address translation (NAT) for all internal network IP addresses, thus complicating the ability for a management system to connect to it.
- o The operator may prefer to have network elements initiate management connections believing it is easier to secure one open-port in the data center than to have an open port on each network element in the network.

One key benefit of using SSH as the transport protocol is its ability to multiplex an unspecified number of independently flow-controlled TCP sessions [RFC4254]. This is valuable as the network element only needs to be configured to initiate a single call home connection to a management system, regardless the number of NETCONF channels the management system wants to open.

4. Protocol

The NETCONF server's perspective (e.g., the network element)

- o The NETCONF server initiates a TCP connection to the NETCONF client on the IANA-assigned SSH for NETCONF Call Home port YYYY.
- o The TCP connection is accepted and a TCP session is established.
- o Using this TCP connection, the NETCONF server immediately starts the SSH server protocol. That is, the next message sent on the TCP stream is SSH's Protocol Version Exchange message (section 4.2, [RFC4253]).
- o The SSH connection is established.

The NETCONF client's perspective (e.g., the management system)

- o The NETCONF client listens for TCP connections on the IANA-assigned NETCONF over SSH Call Home port YYYY.
- o The NETCONF client accepts an incoming TCP connection and a TCP session is established.
- o Using this TCP connection, the NETCONF client immediately starts the SSH Client protocol, starting with sending the SSH's Protocol Version Exchange message (section 4.2, [RFC4253]).
- o The SSH connection is established.

5. SSH Server Identification and Verification

When the management system accepts a new incoming TCP connection on the NETCONF over SSH Call Home port, it starts the SSH client protocol. As the SSH client, it MUST authenticate the SSH server, by both identifying the network element and verifying its SSH host key.

Due to call home having the network element initiate the TCP connection, the management system MAY identify the remote peer using the source IP address of the TCP connection. However, identifying the remote peer using the source IP address of the TCP connection is NOT RECOMMENDED as it can only work in networks that use known static addresses.

To support network elements having dynamically-assigned IP addresses, or deployed behind gateways that translate their IP addresses (e.g., NAT), the management system MAY identify the device using its SSH host key. For instance, a fingerprint of the network element's host

key could itself be used as an identifier since each device has a statistically unique host key. However, identifying the remote peer using its host key directly is NOT RECOMMENDED as it requires the host key to be manually verified the first time the network element connects and anytime its host key changes thereafter.

Yet another option for identifying the network element is for its host key to encode the network element's identity, such as if the host key were a certificate. This option enables the host key to change over time, so long as it continues to encode the same identity, but brings the next issue of how the management system can verify the network element's host key is authentic.

The security of SSH is anchored in the ability for the SSH client to verify the SSH server's host key. Typically this is done by comparing the host key presented by the SSH server with one that was previously configured on the SSH client, looking it up in a local database using the identity of the SSH client as the lookup key. Nothing changes regarding this requirement due to the direction reversal of the underlying TCP connection. To ensure security, the management system MUST verify the network element's SSH host key each time a SSH session is established.

However, configuring distinct host keys on the management system doesn't scale well, which is an important consideration to a network management system. A more scalable strategy for the management system is for the network element's manufacturer to sign the network-element's host key with a common trusted key, such as a certificate authority. Then, when the network-element is deployed, the management system only needs to trust a single certificate, which vouches for the authenticity of the various network element host keys.

Since both the identification and verification issues are addressed using certificates, this draft RECOMMENDS network elements use a host key that can encode a unique identifier (e.g., its serial number) and be signed by a common trust anchor (e.g., a certificate authority). Examples of suitable public host keys are the X.509v3 keys defined in defined in [RFC6187] and the PGP keys defined in [RFC4253].

6. Device Configuration

How to configure a device to initiate a NETCONF over SSH Call Home connection is outside the scope of this document, as implementations can support this protocol using a proprietary configuration data model. That said, a YANG [RFC6020] model to configure NETCONF over SSH Call Home is specified in [draft-ietf-netconf-server-model].

7. Security Considerations

This RFC deviates from standard SSH protocol usage by allowing the SSH server to initiate the TCP connection. This conflicts with section 4 of the SSH Transport Layer Protocol RFC [RFC4253], which states "The client initiates the connection". However this statement is made without rationalization and it's not clear how it impacts the security of the protocol, so this section analyzes the security offered by having the client initiate the connection.

First, assuming the SSH server is not using a public host key algorithm that certifies its identity, the security of the protocol doesn't seem to be sensitive to which peer initiates the connection. That is, it is still the case that reliable distribution of host keys (or their fingerprints) should occur prior to first connection and that verification for subsequent connections happens by comparing the host keys in a locally cached database. It does not seem to matter if the SSH server's host name is derived from user-input or extracted from the TCP layer, potentially via a reverse-DNS lookup. Once the host name-to-key association is stored in a local database, no man-in-the-middle attack is possible due to the attacker being unable to guess the real SSH server's private key (Section 9.3.4 (Man-in-the-middle) of [RFC4251]).

That said, this RFC recommends implementations use a public host key algorithm that certifies the SSH server's identity. The identity can be any unique identifier, such as a device's serial number or a deployment-specific value. If this recommendation is followed, then no information from the TCP layer would be needed to lookup the device in a local database and therefore the directionality of the TCP layer is clearly inconsequential.

The SSH protocol negotiates which algorithms it will use during key exchange (Section 7.1 (Algorithm Negotiation) in [RFC4253]). The algorithm selected is essentially the first compatible algorithm listed by the SSH client that is also listed by the SSH server. For a network management application, there may be a need to advertise a large number of algorithms to be compatible with the various devices it manages. The SSH client SHOULD order its list of public host key algorithms such that all the certifiable public host key algorithms are listed first. Additionally, when possible, SSH servers SHOULD only list certifiable public host key algorithms. Note that since the SSH server would have to be configured to know which IP address it is to connect to, it is expected that it will also be configured to know which host key algorithm to use for the particular application, and hence only needs to list just that one public host key algorithm.

This RFC suggests implementations can use a device's serial number as a form of identity. A potential concern with using a serial number is that the SSH protocol passes the SSH server's host-key in the clear and many times serial numbers encode revealing information about the device, such as what kind of device it is and when it was manufactured. While there is little security in trying to hide this information from an attacker, it is understood that some deployments may want to keep this information private. If this is a concern, deployments SHOULD use an alternate unique identifier, if even just the hash of the device's serial number.

An attacker could DoS the application by having it perform computationally expensive operations, before deducing that the attacker doesn't possess a valid key. This is no different than any secured service and all common precautions apply (e.g., blacklisting the source address after a set number of unsuccessful login attempts).

8. IANA Considerations

This document requests that IANA assigns a TCP port number in the "Registered Port Numbers" range with the service name "netconf-ssh-ch". This port will be the default port for NETCONF over SSH Call Home protocol and will be used when the NETCONF server is to initiate a connection to a NETCONF client using SSH. Below is the registration template following the rules in [RFC6335].

Service Name:	netconf-ssh-ch
Transport Protocol(s):	TCP
Assignee:	IESG <iesg@ietf.org>
Contact:	IETF Chair <chair@ietf.org>
Description:	NETCONF over SSH Call Home
Reference:	RFC XXXX
Port Number:	YYYY

9. Acknowledgements

The author would like to thank for following for lively discussions on list and in the halls (ordered by last name): Andy Bierman, Martin Bjorklund, Mehmet Ersue, Wes Hardaker, Stephen Hanna, David Harrington, Jeffrey Hutzelman, Radek Krejci, Alan Luchuk, Mouse, Russ Mundy, Tom Petch, Peter Saint-Andre, Joe Touch, Sean Turner, Bert Wijnen.

10. References

10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels ", BCP 14, RFC 2119, March 1997.
- [RFC4250] Lehtinen, S. and C. Lonvick, "The Secure Shell (SSH) Protocol Assigned Numbers ", RFC 4250, December 2005.
- [RFC4251] Ylonen, T. and C. Lonvick, "The Secure Shell (SSH) Protocol Architecture ", RFC 4251, January 2006.
- [RFC4252] Ylonen, T. and C. Lonvick, "The Secure Shell (SSH) Authentication Protocol ", RFC 4252, January 2006.
- [RFC4253] Ylonen, T. and C. Lonvick, "The Secure Shell (SSH) Transport Layer Protocol ", RFC 4253, January 2006.
- [RFC4254] Ylonen, T. and C. Lonvick, "The Secure Shell (SSH) Connection Protocol ", RFC 4254, January 2006.
- [RFC6020] Bjorklund, M., "YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF) ", RFC 6020, October 2010.
- [RFC6187] Igoe, K. and D. Stebila, "X.509v3 Certificates for Secure Shell Authentication ", RFC 6187, March 2011.
- [RFC6241] Enns, R., Bjorklund, M., Schoenwaelder, J., and A. Bierman, "NETCONF Configuration Protocol", RFC 6241, June 2011.
- [RFC6242] Wasserman, M., "Using the NETCONF Protocol over Secure Shell (SSH)", RFC 6242, June 2011.
- [RFC6335] Cotton, M., Eggert, L., Touch, J., Westerlund, M., and S. Cheshire, "Internet Assigned Numbers Authority (IANA) Procedures for the Management of the Service Name and Transport Protocol Port Number Registry", RFC 6335, August 2011.

10.2. Informative References

- [draft-ietf-netconf-server-model]
Watsen, K. and J. Schoenwaelder, "A YANG Data Model for NETCONF Server Configuration", RFC 6242, June 2011.

Appendix A. Change Log

A.1. 05 to 06

Changed title to "NETCONF Call Home using SSH"

Revised the Abstract and Introduction to better explain what the document regards.

Changed "MUST" to "SHOULD" in the Applicability Statement.

Added a "Draft Naming" section explaining why, despite its name, reversing SSH is nowhere in the text

Added PGP keys as another kind of SSH host key encoding identity and signed by a trust anchor.

Revised the Device Considerations section to more clearly explain why a device configuration data model is out of scope, and hence an Informative reference.

Clarified Security Considerations section on use of serial numbers.

A.2. 04 to 05

Changed "Reverse SSH" to "Call Home"

Added references to Applicability Statement

A.3. 03 to 04

Changed title to "Reverse SSH for NETCONF Call Home" (changed again in -05)

Removed statement on how other SSH channels might be used for other protocols

Improved language on how the management system, as the SSH client, MUST authenticate the SSH server

Clarified that identifying the network element using source IP address is NOT RECOMMENDED

Clarified that identifying the NE using simple certificate comparison is NOT RECOMMENDED

Device Configuration section now more clearly states that the YANG model is out of scope

Change requested port name to "netconf-ssh-ch"

General edits for grammar, capitalization, and spellings

A.4. 02 to 03

Updated Device Configuration section to reference
[draft-ietf-netconf-server-model]

A.5. 01 to 02

Added Applicability Statement

Removed references to ZeroConf / ZeroTouch

Clarified the protocol section

Added a section for identification and verification

A.6. 00 to 01

Removed the hmac-* family of algorithms

Author's Address

Kent Watsen
Juniper Networks

EMail: kwatsen@juniper.net