

Workgroup: NETCONF
Internet-Draft:
draft-rogaglia-netconf-trace-ctx-extension-00
Published: 12 October 2022
Intended Status: Standards Track
Expires: 15 April 2023
Authors: R. Gagliano K. Larsson J. Lindblad
 Cisco Systems Deutsche Telekom AG Cisco Systems
NETCONF Extension to support Trace Context propagation

Abstract

This document defines how to propagate trace context information across the Network Configuration Protocol (NETCONF), that enables distributed tracing scenarios. It is an adaption of the HTTP-based W3C specification.

About This Document

This note is to be removed before publishing as an RFC.

The latest revision of this draft can be found at [TBD](#). Status information for this document may be found at <https://datatracker.ietf.org/doc/draft-rogaglia-netconf-trace-ctx-extension/>.

Discussion of this document takes place on the NETCONF Working Group mailing list (<mailto:netconf@ietf.org>), which is archived at <https://mailarchive.ietf.org/arch/browse/netmod/>. Subscribe at <https://www.ietf.org/mailman/listinfo/netconf/>.

Source for this draft and an issue tracker can be found at <https://github.com/TBD>.

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 <https://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 15 April 2023.

Copyright Notice

Copyright (c) 2022 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 (<https://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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

- [1. Introduction](#)
 - [1.1. Implementation example: OpenTelemetry](#)
 - [1.2. Use Cases](#)
 - [1.2.1. Provisioning root cause analysis](#)
 - [1.2.2. System performance profiling](#)
 - [1.3. Terminology](#)
- [2. NETCONF Extension](#)
- [3. Security Considerations](#)
- [4. IANA Considerations](#)
- [5. Acknowledgments](#)
- [6. References](#)
 - [6.1. Normative References](#)
 - [6.2. Informative References](#)
- [Appendix A. TO DO List \(to be deleted by RFC Editor\)](#)
- [Appendix B. XML Attributes vs RPCs input augmentations discussion \(to be deleted by RFC Editor\)](#)
- [Authors' Addresses](#)

1. Introduction

Network automation and management systems commonly consist of multiple sub-systems and together with the network devices they manage, they effectively form a distributed system. Distributed tracing is a methodology implemented by tracing tools to follow, analyze and debug operations, such as configuration transactions, across multiple distributed systems. An operation is uniquely identified by a trace-id and through a trace context, carries some metadata about the operation. Propagating this "trace context" between systems enables forming a coherent view of the entire operation as carried out by all involved systems.

Each of the components in this example (Orchestrator, Controller and Network Elements) is exporting M.E.L.T information to the collector using the OpenTelemetry Protocol (OTLP).

For every edit-config operation, the trace context is included. In particular, the same trace-id "1" (simplified encoding for documentation) is included in all related NETCONF messages, which enables the collector and any backend application to correlate all M.E.L.T messages related to this transaction in this distributed stack.

Another interesting attribute is the parent-id. We can see in this example that the parent-id between the orchestrator and the controller ("A") is different from the one between the controller and the network elements ("B"). This attribute will help the collector and the backend applications to build a connectivity graph to understand how M.E.L.T information exported from one component relates to the information exported from a different component.

With this additional metadata exchanged between the components and exposed to the M.E.L.T collector, there are important improvements to the monitor and troubleshooting operations for the full application stack.

1.2. Use Cases

1.2.1. Provisioning root cause analysis

When a provisioning activity fails, errors are typically propagated northbound, however this information may be difficult to troubleshoot and typically, operators are required to navigate logs across all the different components.

With the support for trace context propagation as described in this document for NETCONF, the telemetry collector will be able to search every trace, event, metric, or log in connection to that trace-id and perform a root cause analysis.

1.2.2. System performance profiling

When operating a distributed system such as the one shown in Figure 2, operators are expected to benchmark what are the Key Performance Indicators (KPIs) for the most common tasks. For example, what is the typical delay when provisioning a VPN service across different controllers and devices.

Thanks to Application Performance Management (APM) systems, from these KPIs, an operator can detect a normal and abnormal behaviour of

the distributed system. Also, an operator can better plan any upgrades or enhancements in the platform.

With the support for context propagation as described in this document for NETCONF, much richer system-wide KPIs can be defined and used for troubleshooting as the metrics and traces propagated by the different components share a common context. Troubleshooting for abnormal behaviours can also be troubleshot from the system view down to the individual element.

1.3. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [[RFC2119](#)] [[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

The XML prefixes used in this document are mapped as follows:

```
*xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0",  
*xmlns:notif="urn:ietf:params:xml:ns:netconf:notification:1.0",  
*xmlns:yp="urn:ietf:params:xml:ns:yang:ietf-yang-patch" and  
*xmlns:ypatch="urn:ietf:params:xml:ns:yang:ietf-yang-patch".
```

2. NETCONF Extension

When performing NETCONF operations by sending NETCONF RPCs, a NETCONF client MAY include trace context information in the form of XML attributes. The [[W3C-Trace-Context](#)] defines two HTTP headers; `traceparent` and `tracestate` for this purpose. NETCONF clients that are taking advantage of this feature MUST add one `w3ctc:traceparent` attribute to the `nc:rpc` tag.

A NETCONF server that receives a trace context attribute in the form of a `w3ctc:traceparent` attribute SHOULD apply the mutation rules described in [[W3C-Trace-Context](#)]. A NETCONF server MAY add one `w3ctc:traceparent` attribute in the `nc:rpc-reply` response to the `nc:rpc` tag above. NETCONF servers MAY also add one `w3ctc:traceparent` attribute in notification and update message envelopes: `notif:notification`, `yp:push-update` and `yp:push-change-update`.

For example, a NETCONF client might send:

```
<rpc xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="1"
  xmlns:w3ctc="urn:ietf:params:xml:ns:netconf:w3ctc:1.0"
  w3ctc:traceparent=
    "00-4bf92f3577b34da6a3ce929d0e0e4736-00f067aa0ba902b7-01">
  <get-config/>
</rpc>
```

In all cases above where a client or server adds a `w3ctc:traceparent` attribute to a tag, that client or server MAY also add one `w3ctc:tracestate` attribute to the same tag.

The proper encoding and interpretation of the contents of the `w3ctc:traceparent` attribute is described in [\[W3C-Trace-Context\]](#) section 3.2 except 3.2.1. The proper encoding and interpretation of the contents in the `w3ctc:tracestate` attribute is described in [\[W3C-Trace-Context\]](#) section 3.3 except 3.3.1 and 3.3.1.1. A NETCONF tag can only have zero or one `w3ctc:tracestate` attributes, so its content MUST always be encoded as a single string. The `tracestate` field value is a list of list-members separated by commas (,). A list-member is a key/value pair separated by an equals sign (=). Spaces and horizontal tabs surrounding list-members are ignored. There is no limit to the number of list-members in a list.

For example, a NETCONF client might send:

```
<rpc xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="1"
  xmlns:w3ctc="urn:ietf:params:xml:ns:netconf:w3ctc:1.0"
  w3ctc:tracestate="rojo=00f067aa0ba902b7,congo=t61rcWkgMzE"
  w3ctc:traceparent=
    "00-4bf92f3577b34da6a3ce929d0e0e4736-00f067aa0ba902b7-01">
  <get-config/>
</rpc>
```

As in all XML documents, the order between the attributes in an XML tag has no significance. Clients and servers MUST be prepared to handle the attributes no matter in which order they appear. The `tracestate` value MAY contain double quotes in its payload. If so, they MUST be encoded according to XML rules, for example:

```
<rpc xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="1"
  xmlns:w3ctc="urn:ietf:params:xml:ns:netconf:w3ctc:1.0"
  w3ctc:traceparent=
    "00-4bf92f3577b34da6a3ce929d0e0e4736-00f067aa0ba902b7-01"
  w3ctc:tracestate=
    "value-with-quotes=&quot;Quoted string&quot;;other-value=123">
  <get-config/>
</rpc>
```

TBD Errors

3. Security Considerations

TODO Security

4. IANA Considerations

This document registers the following capability identifier URN in the 'Network Configuration Protocol (NETCONF) Capability URNs' registry:

urn:ietf:params:netconf:capability:w3ctc:1.0

This document registers one XML namespace URN in the 'IETF XML registry', following the format defined in [[RFC3688](https://tools.ietf.org/html/rfc3688)] (<https://tools.ietf.org/html/rfc3688>).

URI: urn:ietf:params:xml:ns:netconf:w3ctc:1.0

Registrant Contact: The NETCONF WG of the IETF.

XML: N/A, the requested URI is an XML namespace.

5. Acknowledgments

TBD

6. References

6.1. Normative References

[[RFC2119](#)] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/

RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.

[RFC3688] Mealling, M., "The IETF XML Registry", BCP 81, RFC 3688, DOI 10.17487/RFC3688, January 2004, <<https://www.rfc-editor.org/info/rfc3688>>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

[W3C-Trace-Context] "W3C Recommendation on Trace Context", 23 November 2021, <<https://www.w3.org/TR/2021/REC-trace-context-1-20211123/>>.

6.2. Informative References

[I-D.lindblad-netconf-transaction-id]

Lindblad, J., "Transaction ID Mechanism for NETCONF", Work in Progress, Internet-Draft, draft-lindblad-netconf-transaction-id-02, 8 June 2022, <<https://www.ietf.org/archive/id/draft-lindblad-netconf-transaction-id-02.txt>>.

[OpenTelemetry] "OpenTelemetry Cloud Native Computing Foundation project", 29 August 2022, <<https://opentelemetry.io>>.

[RFC8309] Wu, Q., Liu, W., and A. Farrel, "Service Models Explained", RFC 8309, DOI 10.17487/RFC8309, January 2018, <<https://www.rfc-editor.org/info/rfc8309>>.

[W3C-Baggage] "W3C Propagation format for distributed context Baggage", 23 November 2021, <<https://www.w3.org/TR/baggage/#examples-of-http-headers>>.

Appendix A. TO DO List (to be deleted by RFC Editor)

*Manage versioning of the trace-context specification

*We intend to extend the trace-concext capability to RESTCONF in a future draft

*The W3C is working on a draft document to introduce the concept of "baggage" [W3C-Baggage] that we expect part of a future draft for NETCONF and RESTCONF

Appendix B. XML Attributes vs RPCs input augmentations discussion (to be deleted by RFC Editor)

There are arguments that can be raised regarding using XML Attribute or to augment NETCONF RPCs.

We studied Pros/Cons of each option and decided to propose XML attributes:

XML Attributes Pro:

- *Literal alignment with W3C specification
- *Same encoding for RESTCONF and NETCONF enabling code reuse
- *One specification for all current and future rpcs

XML Attributes Cons:

- *No YANG modeling, multiple values represented as a single string
- *Dependency on W3C for any extension or changes in the future as encoding will be dictated by string encoding

RPCs Input Augmentations Pro:

- *YANG model of every leaf
- *Re-use of YANG toolkits
- *Simple updates by augmentations on existing YANG module
- *Possibility to express deviations in case of partial support

RPCs Input Augmentations Cons:

- *Need to augment every rpc, including future rpcs would need to consider these augmentations, which is harder to maintain
- *There is no literal alignment with W3C standard. However, as mentioned before most of the time there will be modifications to the content
- *Would need updated RFP for each change at W3C, which will make adoption of new features slower

Authors' Addresses

Roque Gagliano
Cisco Systems
Avenue des Uttins 5
CH-1180 Rolle
Switzerland

Email: rogaglia@cisco.com

Kristian Larsson

Deutsche Telekom AG

Email: kl@dev.terastrm.net

Jan Lindblad
Cisco Systems

Email: jlindbla@cisco.com