

NMOP
Internet-Draft
Intended status: Standards Track
Expires: 23 April 2026

TG. Graf
AE. Elhassany
Swisscom
AHF. Huang Feng
INSA-Lyon
BC. Claise
Everything OPS
20 October 2025

YANG Message Keys for Message Broker Integration
draft-netana-nmop-yang-message-broker-message-key-01

Abstract

This document specifies a mechanism to define a unique message key for a YANG to message broker integration and a topic addressing scheme based on YANG-Push subscription type and a YANG index defined in this document. This enables YANG data consumption of a subset of subscribed YANG data, either per specific YANG node, identifier or telemetry message type, by indexing and organizing in Message Broker topics, indexing the information the right way.

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 23 April 2026.

Copyright Notice

Copyright (c) 2025 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	2
2. Conventions and Definitions	4
2.1. Terminology	4
3. Solution Design	6
3.1. YANG Message Keys and Indexes	6
3.1.1. YANG Message Broker Producer	7
3.1.2. YANG Message Broker Consumer	8
3.2. YANG-Push Message Broker Topic Naming	8
3.2.1. YANG Message Broker Producer	9
3.2.2. YANG Message Broker Consumer	9
4. Message Broker Implementations	9
4.1. Apache Kafka	9
4.2. Apache Pulsar	9
5. IANA Considerations	10
6. Security Considerations	10
7. Operational Considerations	10
8. References	10
8.1. Normative References	10
8.2. Informative References	11
Acknowledgements	12
Contributors	13
Authors' Addresses	13

1. Introduction

Nowadays network operators are using machine and human readable YANG [RFC7950] to model their configurations and monitor YANG operational data from their networks according to [Mar24].

Most network analytic use cases require real-time data and deliver near real-time analytical, actionable insights. This imposes high scalability, resilience and low overhead in the data processing pipeline. Accessing the right data for the right use case with minimal overhead and in the shortest period of time is therefore crucial.

Network operators organize their data in a Data Mesh [Deh22] according to [Bod24] where a Message Broker such as Apache Kafka [Kaf11] or Apache Pulsar [Pul16] facilitates the exchange of messages among data processing components in topics and subjects. Typically,

data is being stored in Message Broker topics for several hours or days to facilitate resilience in the data processing chain and addressed in subjects depending on schema, enabling a data consumer to address and re-consume previously consumed data again if previously lost.

Dimensional data is structured information in a data store. It uses a model of dimension tables to organize business metrics and their descriptive context. This model, developed by Ralph Kimball [Kim96], simplifies data analysis and reporting by creating denormalized, easy-to-understand structures for quick querying. It is optimized for online analytical processing (OLAP) and data warehouses. YANG [RFC7950] as a data modelling language facilitates the modelling of dimensional data.

An Architecture for YANG-Push to Message Broker Integration [I-D.ietf-nmop-yang-message-broker-integration] specifies an architecture for integrating YANG-Push with Message Brokers for a Data Mesh architecture. How the notification messages at a YANG-Push Receiver are being transformed to the Message Broker is described in Section 4.5 of [I-D.ietf-nmop-yang-message-broker-integration] and to which message schema in Section 3 of [I-D.ietf-nmop-message-broker-telemetry-message], however how messages should be indexed best for dimensional YANG data is left unspecified.

Due to the missing dimensional indexing for Message Broker stored YANG data, all YANG data is stored in one single Topic, distributed round robin across multiple Partitions and each YANG schema id is a subject within that topic. Therefore, the entire Topic from all Partitions needs to be consumed first before data selection can be applied. This leads to avoidable data processing overhead which in turn impairs scalability and real-time capabilities, required for certain Network Analytics use cases.

YANG telmetry data can be used for several network analytic use cases. Importantly, depending on the use case, only a subset of the subscribed YANG data might be necessary (in time or space). For example, for specific use cases, it's more important to know the current network state, as opposed to have the full series of the state changes over time. In other use cases, instead of consuming data for all network nodes, only a specific network node or network node component requires the YANG monitoring and hence subscription.

This document defines how YANG messages [I-D.ietf-nmop-message-broker-telemetry-message] should be indexed and organized in Message Broker topics by leveraging the network node hostname, YANG datastore name and YANG Item Identifier for indexing; and YANG-Push subscription type and YANG schema name for a Message Broker topic naming scheme.

Network node hostname, YANG datastore name and subtree and xpath filters are part of "ietf-yang-push-telemetry-message" structured YANG data defined in Section 3 of [I-D.ietf-nmop-message-broker-telemetry-message]. YANG item identifier are derived from subtree and xpath filters respectively from their YANG schema tree.

2. Conventions and Definitions

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.

2.1. Terminology

The following terms are used as defined in [I-D.ietf-nmop-terminology]:

- * Network Telemetry
- * Network Analytics
- * Value
- * State
- * Change

The following terms are used as defined in [I-D.ietf-nmop-yang-message-broker-integration]:

- * Message Broker
- * YANG Message Broker Producer
- * YANG Message Broker Consumer

The following terms are used as defined in Apache Kafka [Kaf11] and Apache Pulsar [Pul16] Message Broker:

- * **Subject:** A named communication channel where a schema id is associated.
- * **Topic:** A communication channel for publishing and subscribing messages with one or more subjects.
- * **Topic Compaction:** The act of compressing messages in a topic to the latest state. As used with Apache Pulsar. Apache Kafka uses the term Log Compaction with identical meaning.
- * **Partition:** Messages in a topic are spread over hash buckets where a hash bucket refers to a partition.
- * **Message:** A piece of structured data sent between data processing components to facilitate communication in a distributed system
- * **Message Key:** Metadata associated with a message to facilitate deterministic hash bucketing.

The following terms are used as defined in Confluent Schema Registry Documentation [ConDoc18]:

- * **Schema:** A formalized, documented structure that defines the shape and content of the messages exchange.
- * **Schema ID:** A unique identifier of a schema associated to a Message Broker subject.

The following terms are used as defined in [RFC8641]:

- * Periodical
- * On-Change
- * Sync-On-Start
- * Xpath Filter
- * Subtree Filter

The following terms are used as defined in [I-D.ietf-netconf-notif-envelope]:

- * Notification
- * Hostname

The following terms are used as defined in [RFC8342]:

- * Datastore

The following terms are used as defined in [RFC7950]:

- * Schema Node Identifier

- * Schema Tree

The following terms are used as defined in [RFC9254]:

- * YANG Item Identifiers

This document defines the following term:

- * YANG Index: Is a subset of YANG Item Identifiers containing only schema node identifiers. Different to an absolute schema node identifier it includes the YANG module name and is therefore globally unique. When the schema node identifier points to a YANG list, then the key to that list is included.

3. Solution Design

In order to identify which YANG node identifier of a network node YANG datastore is produced in which Message Broker Topic and Partition for which Subject, YANG Message Keys and Indexes (Section 3.1) are being introduced.

In order to facilitate Message Broker Topic Compaction, a YANG-Push subscription type based topic naming scheme (Section 3.2) is proposed. This segregates statistical (Value), State and State change YANG metrics and facilitates a YANG Message Broker Consumer to use the Topic wild card consumption method to select based on YANG-Push subscription type.

3.1. YANG Message Keys and Indexes

A Message Broker MUST use a Message Key to index the message and a value to carry the Message content. If no Message Key is defined then the Messages are distributed in a round robin fashion across partitions. If a Message Key is defined, then the value of the Message Key is being used as input for the Message Broker Producer hash function to distribute across Partitions. Therefore, Message Keys facilitate Message ordering.

The Message Key not only used for Message indexing at the Message Producer but also at the Message Broker for topic compaction.

For YANG, the network node hostname, from which YANG datastore the YANG metrics are published from and the YANG index is used to generate the Message Key.

3.1.1. YANG Message Broker Producer

YANG data nodes are uniquely identifiable within the YANG schema tree. Section 6.5 of [RFC7950] defines with "absolute-schema-nodeid" how absolute YANG schema node identifiers are being crafted locally unique to the YANG module.

Section 3.3 of [RFC9254] defines how globally unique YANG Item Identifiers are defined as text strings.

Section 3.6 of [RFC8641] defines how YANG data nodes can be subscribed with subtree and xpath selection filters. A YANG-Push publisher publishes with "subscription-started" state notifications for each subscription which filter and filter type is being used to the YANG-Push receiver.

To calculate the YANG Index of the Message Key, the YANG item identifier needs to be extracted from the used YANG-Push subtree or xpath subscription filter. If the YANG item identifier is a YANG list as defined in Section 7.8 of [RFC7950] the YANG list key defined in Section 7.8.2 of [RFC7950] statement is suffixed with a "/" to the YANG Item Identifier.

For example, if the following xpath filter is being used, the YANG Item Identifier is "ietf-interface:interfaces/interface". Interface is a YANG list with name as key. Therefore, the YANG Index of the Message Key is "ietf-interface:interfaces/interface/name".

```
ietf-interface:interfaces/interface[type='ianaift:ethernetCsmacd']
```

Figure 1: YANG-Push ietf-interface Xpath Filter Example

For example, if the following subtree filter is being used, the YANG Item Identifier is "ietf-hardware:hardware/component/state". Therefore, the YANG Index of the Message Key is "ietf-hardware:hardware/component/state".

```
<get>
  <filter type="subtree">
    <hardware xmlns="urn:ietf:params:xml:ns:yang:ietf-hardware">
      <component>
        <state/>
      </component>
    </hardware>
  </filter>
</get>
```

Figure 2: YANG-Push ietf-hardware Subtree Filter Example

When the Message is being produced to the Message Broker, the Network node hostname and YANG datastore name is used from the structured YANG data defined in "ietf-yang-push-telemetry-message" Section 3 of [I-D.ietf-nmop-message-broker-telemetry-message] where the YANG Index is derived from subtree and xpath filters, respectively from their YANG schema tree.

3.1.2. YANG Message Broker Consumer

The consumer hashes the Message Key and applies modulo with the number of partitions to determine the partition it needs to consume from to obtain Messages with desired Message Key.

3.2. YANG-Push Message Broker Topic Naming

YANG can be subscribed periodically, on-change or on-change with sync-on-start. Periodical subscriptions are used for obtaining statistical metrics. On-Change subscriptions are used for obtaining State Changes and on-change with sync-on-start for obtaining States.

Message Brokers topics are addressed with a unique name. Usually topics are named hierarchically similar to the DNS namespace where "." delimitates hierarchies.

This document defines "statistics", "states" and "state-changes" in the topic name as the first part to denote the types of data. Followed by "yang" to denote YANG data. Followed by the YANG module names subscribed, and followed by the YANG Schema Node Identifier where "/" is substituted by "_".

For example, if the "ietf-interface:interfaces/interface" xpath filter is being used, the Message Broker topic name would be as following. In the example the project name and environment (prod, dev, test etc.) is prefixed.

project.environment.statistics.yang.ietf-interfaces.interfaces_interface

Figure 3: YANG-Push ietf-interface Topic Name Example

3.2.1. YANG Message Broker Producer

For the Message Broker topic creation, the "periodic", "on-change" and "sync-on-start" contained data in "update-trigger" from "ietf-subscribed-notifications", YANG module defined in Section 4.1 of [RFC8641], subscription state notifications MUST be used to derive wherever subscribed YANG data is "statistics", "states" or "state-changes". The YANG Index MUST be derived from subtree and xpath filter data of subscription state notifications, respectively from their YANG schema tree.

3.2.2. YANG Message Broker Consumer

The consumer has the ability to consume with a wildcard denoted with "*" in the topic name to consume from more than one topic.

For example, if YANG states should be consumed and indexed in Time Series database or stream processor than below Topic Name could be used, and the YANG data could be ingested into tables according to topic names and indexed per Message Key. If Topic Compaction is enabled, only current state is consumed.

project.environment.states.yang.*

Figure 4: YANG-Push Wildcard Topic Name Example

4. Message Broker Implementations

Topic, Partitioning and Message Keying are generic concepts of Message Brokers. There are two known Message Broker implementations supporting all features described in this document.

4.1. Apache Kafka

Apache Kafka supports Message Keying, Partitioning and Log Compaction. The topic names are constrained to 249 character length and the following characters: "a-z", "A-Z", "0-9", ".", "_" and "-".

4.2. Apache Pulsar

Apache Pulsar supports Message Keying, Partitioning and Topic Compaction. The topic names allow all characters except: "/".

5. IANA Considerations

This document includes no request to IANA.

6. Security Considerations

This document should not affect the security of the Internet.

7. Operational Considerations

The YANG Message Broker Producer of a YANG-Push receiver should have three config knobs facilitate the features described in this document as optional:

- * Topic Distribution: Select between "topic" and "subject" distribution. Default is subject to remain backward compatibility to [I-D.ietf-nmop-yang-message-broker-integration].
- * Distribution Type: Select between "none" and "YANG-Push subscription type".
- * YANG Message Key: Select between "enable" and "disable".

To accommodate for potential data loss throughout the data processing pipeline, periodical update of the current State for State metrics is RECOMMENDED. This can be accommodated with YANG-Push as defined in [RFC8641] by complementing "on-change sync on start" subscriptions with periodical subscriptions. Alternatively, in YANG-Push Lite defined in Section 7.6 of [I-D.wilton-netconf-yang-push-lite] this simplified in one subscription.

8. References

8.1. 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>>.
- [RFC7950] Bjorklund, M., Ed., "The YANG 1.1 Data Modeling Language", RFC 7950, DOI 10.17487/RFC7950, August 2016, <<https://www.rfc-editor.org/info/rfc7950>>.
- [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>>.

- [RFC8342] Bjorklund, M., Schoenwaelder, J., Shafer, P., Watsen, K., and R. Wilton, "Network Management Datastore Architecture (NMDA)", RFC 8342, DOI 10.17487/RFC8342, March 2018, <<https://www.rfc-editor.org/info/rfc8342>>.
- [RFC8641] Clemm, A. and E. Voit, "Subscription to YANG Notifications for Datastore Updates", RFC 8641, DOI 10.17487/RFC8641, September 2019, <<https://www.rfc-editor.org/info/rfc8641>>.
- [RFC9254] Veillette, M., Ed., Petrov, I., Ed., Pelov, A., Bormann, C., and M. Richardson, "Encoding of Data Modeled with YANG in the Concise Binary Object Representation (CBOR)", RFC 9254, DOI 10.17487/RFC9254, July 2022, <<https://www.rfc-editor.org/info/rfc9254>>.
- [I-D.ietf-nmop-terminology]
Davis, N., Farrel, A., Graf, T., Wu, Q., and C. Yu, "Some Key Terms for Network Fault and Problem Management", Work in Progress, Internet-Draft, draft-ietf-nmop-terminology-23, 18 August 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-nmop-terminology-23>>.
- [I-D.ietf-nmop-yang-message-broker-integration]
Graf, T. and A. Elhassany, "An Architecture for YANG-Push to Message Broker Integration", Work in Progress, Internet-Draft, draft-ietf-nmop-yang-message-broker-integration-08, 7 July 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-nmop-yang-message-broker-integration-08>>.
- [I-D.ietf-nmop-message-broker-telemetry-message]
Elhassany, A. and T. Graf, "Extensible YANG Model for Network Telemetry Messages", Work in Progress, Internet-Draft, draft-ietf-nmop-message-broker-telemetry-message-02, 5 September 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-nmop-message-broker-telemetry-message-02>>.
- [I-D.ietf-netconf-notif-envelope]
Feng, A. H., Francois, P., Graf, T., and B. Claise, "Extensible YANG Model for YANG-Push Notifications", Work in Progress, Internet-Draft, draft-ietf-netconf-notif-envelope-02, 17 June 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-netconf-notif-envelope-02>>.

8.2. Informative References

- [I-D.wilton-netconf-yang-push-lite]
Wilton, R., Keller, H., Claise, B., Aries, E., Cumming, J., and T. Graf, "YANG Datastore Telemetry (YANG Push Lite)", Work in Progress, Internet-Draft, draft-wilton-netconf-yang-push-lite-01, 7 July 2025, <<https://datatracker.ietf.org/doc/html/draft-wilton-netconf-yang-push-lite-01>>.
- [Mar24] Martinez-Casanueva, I. D., Gonzalez-Sanchez, D., Bellido, L., Fernandez, D., and D. R. Lopez, "Toward Building a Semantic Network Inventory for Model-Driven Telemetry", IEEE, DOI 10.1109/MCOM.001.2200222, February 2024, <<https://arxiv.org/html/2402.06511v1>>.
- [Bod24] Bode, J., K端hl, N., Kreuzberger, D., and C. Holtmann, "Toward Avoiding the Data Mess: Industry Insights From Data Mesh Implementations", IEEE, DOI 10.1109/ACCESS.2024.3417291, January 2024, <<https://arxiv.org/html/2302.01713v4>>.
- [Deh22] Dehghani, Z., "Data Mesh", O'Reilly Media, ISBN 9781492092391, March 2022, <<https://www.oreilly.com/library/view/data-mesh/9781492092384/>>.
- [Kim96] Kimball, R. and M. Ross, "The Data Warehouse Toolkit", Wiley, ISBN 9781118530801, 1996, <<https://www.kimballgroup.com/data-warehouse-business-intelligence-resources/books/data-warehouse-dw-toolkit/>>.
- [Kaf11] Narkhede, N., "Apache Kafka", Apache Software Foundation, January 2011, <<https://kafka.apache.org/>>.
- [Pul16] Guo, S. and M. Merli, "Apache Pulsar", Apache Software Foundation, January 2016, <<https://pulsar.apache.org/>>.
- [ConDoc18] Yokota, R., "Confluent Schema Registry Documentation", Confluent Community and Apache Software Foundation, December 2018, <<https://docs.confluent.io/platform/current/schema-registry/>>.

Acknowledgements

Thanks to

Contributors

We like to thank Victor Lopez for the initial idea on the network controller use case. Ashley Woods, Sivakumar Sundaravadivel and Rafael Julio for the idea of grouping topics by YANG-Push subscription type and insisting that Topic Compaction is a key enabler for inventory metrics and YANG data consumer integration and should be supported day 1. And Nigel Davis for confirming that Topic Compaction simplifies indeed data processing system architecture.

Authors' Addresses

Thomas Graf
Swisscom
Binzring 17
CH-8045 Zurich
Switzerland
Email: thomas.graf@swisscom.com

Ahmed Elhassany
Swisscom
Binzring 17
CH-8045 Zurich
Switzerland
Email: ahmed.elhassany@swisscom.com

Alex Huang Feng
INSA-Lyon
Lyon
France
Email: alex.huang-feng@insa-lyon.fr

Benoît Claise
Everything OPS
Liege
Belgium
Email: benoit@everything-ops.net