

Workgroup: Network Working Group
Internet-Draft:
draft-ietf-netmod-yang-schema-comparison-01
Published: 1 November 2020
Intended Status: Standards Track
Expires: 5 May 2021
Authors: R. Wilton
Cisco Systems, Inc.
YANG Schema Comparison

Abstract

This document specifies an algorithm for comparing two revisions of a YANG schema to determine the scope of changes, and a list of changes, between the revisions. The output of the algorithm can be used to help select an appropriate revision-label or YANG semantic version number for a new revision. This document defines a YANG extension that provides YANG annotations to help the tool accurately determine the scope of changes between two revisions.

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 5 May 2021.

Copyright Notice

Copyright (c) 2020 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 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. Terminology and Conventions](#)
- [3. Generic YANG schema tree comparison algorithm](#)
 - [3.1. YANG module revision scope extension annotations](#)
- [4. YANG module comparison algorithm](#)
- [5. YANG schema comparison algorithms](#)
 - [5.1. Standard YANG schema comparison algorithm](#)
 - [5.2. Filtered YANG schema comparison algorithm](#)
- [6. Comparison tooling](#)
- [7. Module Versioning Extension YANG Modules](#)
- [8. Contributors](#)
- [9. Security Considerations](#)
- [10. IANA Considerations](#)
 - [10.1. YANG Module Registrations](#)
- [11. References](#)
 - [11.1. Normative References](#)
 - [11.2. Informative References](#)
- [Author's Address](#)

1. Introduction

Warning, this is an early (-00) draft with the intention of scoping the outline of the solution, hopefully for the WG to back the direction of the solution. Refinement of the solution details is expected, if this approach is accepted by the WG.

This document defines a solution to Requirement 2.2 in [[I-D.ietf-netmod-yang-versioning-reqs](#)]. Complementary documents provide a complete solution to the YANG versioning requirements, with the overall relationship of the solution drafts described in [[I-D.ietf-netmod-yang-solutions](#)].

YANG module 'revision-labels' [[I-D.ietf-netmod-yang-module-versioning](#)] and the use of YANG semantic version numbers [[I-D.ietf-netmod-yang-semver](#)] can be used to help manage and report changes between revisions of individual YANG modules.

YANG packages [[I-D.ietf-netmod-yang-packages](#)] along with YANG semantic version numbers can be used to help manage and report changes between revisions of YANG schema.

[[I-D.ietf-netmod-yang-module-versioning](#)] and [[I-D.ietf-netmod-yang-packages](#)] define how to classify changes between two module or package revisions, respectively, as backwards compatible or non-backwards-compatible. [[I-D.ietf-netmod-yang-semver](#)] refines the

definition, to allow backwards compatible changes to be classified as 'minor changes' or 'editorial changes'.

'Revision-label's and YANG semantic version numbers, whilst being generally simple and helpful in the mainline revision history case, are not sufficient in all scenarios. For example, when comparing two revisions/versions on independent revision branches, without a direct ancestor relationship between the two revisions/versions. In this cases, an algorithmic comparison approach is beneficial.

In addition, the module revision history's 'nbc-changes' extension statement, and YANG semantic version numbers, effectively declare the worst case scenario. If any non-backwards-compatible changes are restricted to only parts of the module/schema that are not used by an operator, then the operator is able to upgrade, and effectively treat the differences between the two revisions/versions as backwards compatible because they are not materially impacted by the non-backwards-compatible changes.

Hence, this document defines algorithms that can be applied to revisions of YANG modules or versions of YANG schema (e.g., as represented by YANG packages), to determine the changes, and scope of changes between the revisions/versions.

For many YANG statements, programmatic tooling can determine whether the changes between the statements constitutes a backwards-compatible or non-backwards-compatible change. However, for some statements, it is not feasible for current tooling to determine whether the changes are backwards-compatible or not. For example, in the general case, tooling cannot determine whether the change in a YANG description statement causes a change in the semantics of a YANG data node. If the change is to fix a typo or spelling mistake then the change can be classified as an editorial backwards-compatible change. Conversely, if the change modifies the behavioral specification of the data node then the change would need to be classified as either a non editorial backwards-compatible change or a non-backwards-compatible change. Hence, extension statements are defined to annotate a YANG module with additional information to clarify the scope of changes in cases that cannot be determined by algorithmic comparison.

Open issues are tracked at <https://github.com/netmod-wg/yang-ver-dt/issues>, tagged with 'schema-comparison'.

2. Terminology and Conventions

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.

This document uses terminology introduced in the YANG versioning requirements document [[I-D.ietf-netmod-yang-versioning-reqs](#)].

This document makes of the following terminology introduced in the YANG Packages [[I-D.ietf-netmod-yang-packages](#)]:

*YANG schema

In addition, this document defines the terminology:

*Change scope: Whether a change between two revisions is classified as non-backwards-compatible, backwards-compatible, or editorial.

3. Generic YANG schema tree comparison algorithm

The generic schema comparison algorithm works on any YANG schema. This could be a schema associated with an individual YANG module, or a YANG schema represented by a set of modules, e.g., specified by a YANG package.

The algorithm performs a recursive tree wise comparison of two revisions of a YANG schema, with the following behavior:

The comparison algorithm primarily acts on the parts of the schema defined by unique identifiers.

Each identifier is qualified with the name of the module that defines the identifier.

Identifiers in different namespaces (as defined in 6.2.1 or RFC 7950) are compared separately. E.g., 'features' are compared separately from 'identities'.

Within an identifier namespace, the identifiers are compared between the two schema revisions by qualified identifier name. The 'renamed-from' extension allow for a meaningful comparison where the name of the identifier has changed between revisions. The 'renamed-from' identifier parameter is only used when an identifier in the new schema revision cannot be found in the old schema revision.

YANG extensions, features, identities, typedefs are checked by comparing the properties defined by their YANG sub-statements between the two revisions.

YANG groupings, top-level data definition statements, rpcs, and notifications are checked by comparing the top level properties defined by their direct child YANG sub-statements, and also by recursively checking the data definition statements.

The rules specified in section 3 of [[I-D.ietf-netmod-yang-module-versioning](#)] determine whether the changes are backwards-compatible or non-backwards-compatible.

The rules specified in section 3.2 of [[I-D.ietf-netmod-yang-packages](#)] determine whether backwards-compatible changes are 'minor' or 'editorial'.

For YANG description, "must", and "when" statements, the "backwards-compatible" and "editorial" extension statements can be used to mark instances when the statements have changed in a backwards-compatible or editorial way. Since by default the comparison algorithm assumes that any changes in these statements are non-backwards-compatible. XXX, more info required here, since the revisions in the module history probably need to be available for this to work in the general branched revisions case.

Submodules are not relevant for schema comparison purposes, i.e. the comparison is performed after submodule resolution has been completed.

3.1. YANG module revision scope extension annotations

4. YANG module comparison algorithm

The schema comparison algorithm defined in [Section 3](#) can be used to compare the schema for individual modules, but with the following modifications:

Changes to the module's metadata information (i.e. module level description, contact, organization, reference) should be checked (as potential editorial changes).

The module's revision history should be ignored from the comparison.

Changes to augmentations and deviations should be sorted by path and compared.

5. YANG schema comparison algorithms

5.1. Standard YANG schema comparison algorithm

The standard method for comparing two YANG schema versions is to individually compare the module revisions for each module

implemented by the schema using the algorithm defined in [Section 4](#) and then aggregating the results together:

- *If all implemented modules in the schema have only changed in an editorial way then the schema is changed in an editorial way
- *If all implemented modules in the schema have only been changed in an editorial or backwards-compatible way then the schema is changed in a backwards-compatible way
- *Otherwise if any implemented module in the schema has been changed in a non-backwards-compatible way then the schema is changed in a non-backwards-compatible way.

The standard schema comparison method is the RECOMMENDED scheme to calculate the version number change for new versions of YANG packages, because it allows the package version to be calculated based on changes to implemented modules revision history (or YANG semantic version number if used to identify module revisions).

5.2. Filtered YANG schema comparison algorithm

Another method to compare YANG schema, that is less likely to report inconsequential differences, is to construct full schema trees for the two schema versions, directly apply a version of the comparison algorithm defined in [Section 3](#). This may be particularly useful when the schema represents a complete datastore schema for a server because it allows various filtered to the comparison algorithm to provide a more specific answer about what changes may impact a particular client.

The full schema tree can easily be constructed from a YANG package definition, or alternative YANG schema definition.

Controlled by input parameters to the comparison algorithm, the following parts of the schema trees can optionally be filtered during the comparison:

All "grouping" statements can be ignored (after all "use" statements have been processed when constructing the schema).

All module and submodule metadata information (i.e. module level description, contact, organization, reference) can be ignored.

The comparison can be restricted to the set of features that are of interest (different sets of features may apply to each schema versions).

The comparison can be restricted to the subset of data nodes, RPCs, notifications and actions, that are of interest (e.g., the

subset actually used by a particular client), providing a more meaningful result.

The comparison could filter out backwards-compatible 'editorial' changes.

In addition to reporting the overall scope of changes at the schema level, the algorithm output can also optionally generate a list of specific changes between the two schema, along with the classification of those individual changes.

6. Comparison tooling

'pyang' has some support for comparison two module revisions, but this is currently limited to a linear module history.

TODO, it would be helpful if there is reference tooling for schema comparison.

7. Module Versioning Extension YANG Modules

YANG module with extension statements for annotating NBC changes, revision label, status description, and importing by version.

```

<CODE BEGINS> file "ietf-yang-rev-annotations@2019-11-11.yang"

module ietf-yang-rev-annotations {
  yang-version 1.1;
  namespace "urn:ietf:params:xml:ns:yang:ietf-yang-rev-annotations";
  prefix rev-ext;

  organization
    "IETF NETMOD (Network Modeling) Working Group";
  contact
    "WG Web:  <https://datatracker.ietf.org/wg/netmod/>
    WG List:  <mailto:netmod@ietf.org>

    Author:   Robert Wilton
              <mailto:rwilton@cisco.com>";

  description
    "This YANG 1.1 module contains extensions to annotation to YANG
    module with additional metadata information on the nature of
    changes between two YANG module revisions.

    XXX, maybe these annotations could also be included in
    ietf-yang-revisions?

    Copyright (c) 2019 IETF Trust and the persons identified as
    authors of the code.  All rights reserved.

    Redistribution and use in source and binary forms, with or
    without modification, is permitted pursuant to, and subject
    to the license terms contained in, the Simplified BSD License
    set forth in Section 4.c of the IETF Trust's Legal Provisions
    Relating to IETF Documents
    (http://trustee.ietf.org/license-info).

    This version of this YANG module is part of RFC XXXX; see
    the RFC itself for full legal notices.

    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 (RFC 2119) (RFC 8174) when, and only when,
    they appear in all capitals, as shown here.";

  // RFC Ed.: update the date below with the date of RFC publication
  // and remove this note.
  // RFC Ed.: replace XXXX (inc above) with actual RFC number and
  // remove this note.

  revision 2019-11-11 {
    description

```



```

    "Initial version.";
reference
    "XXX: YANG Schema Comparison";
}

extension backwards-compatible {
    argument revision-date-or-label;
description
    "Identifies a revision (by revision-label, or revision date if
    a revision-label is not available) where a
    backwards-compatible change has occurred relative to the
    previous revision listed in the revision history.

    The format of the revision-label argument MUST conform to the
    pattern defined for the ietf-yang-revisions
    revision-date-or-label typedef.

    The following YANG statements MAY have zero or more
    'rev-ext:non-backwards-compatible' statements:
        description
        must
        when

    Each YANG statement MUST only a have a single
    non-backwards-compatible, backwards-compatible, or editorial
    extension statement for a particular revision-label, or
    corresponding revision-date.";

reference
    "XXX: YANG Schema Comparison;
    Section XXX, XXX";
}

extension editorial {
    argument revision-date-or-label;
description
    "Identifies a revision (by revision-label, or revision date if
    a revision-label is not available) where an editorial change
    has occurred relative to the previous revision listed in the
    revision history.

    The format of the revision-label argument MUST conform to the
    pattern defined for the ietf-yang-revisions
    revision-date-or-label typedef.

    The following YANG statements MAY have zero or more
    'rev-ext:non-backwards-compatible' statements:
        description

    Each YANG statement MUST only a have a single

```

```

        non-backwards-compatible, backwards-compatible, or editorial
        extension statement for a particular revision-label or
        corresponding revision-date.";

reference
  "XXXX: YANG Schema Comparison;
  Section XXX, XXX";
}

extension renamed-from {
  argument yang-identifier;
  description
    "Specifies a previous name for this identifier.

    This can be used when comparing schema to optimize handling
    for data nodes that have been renamed rather than naively
    treated them as data nodes that have been deleted and
    recreated.

    The argument 'yang-identifier' MUST take the form of a YANG
    identifier, as defined in section 6.2 of RFC 7950.

    Any YANG statement that takes a YANG identifier as its
    argument MAY have a single 'rev-ext:renamed-from'
    sub-statement.

    TODO, we should also facilitate identifiers being moved into
    other modules, e.g. by supporting a module-name qualified
    identifier.";

reference
  "XXXX: YANG Schema Comparison;
  Section XXX, XXX";
}
}

```

<CODE ENDS>

8. Contributors

This document grew out of the YANG module versioning design team that started after IETF 101. The following individuals are (or have been) members of the design team and have worked on the YANG versioning project:

*Balazs Lengyel

*Benoit Claise

*Bo Wu

*Ebben Aries

*Jason Sterne

*Joe Clarke

*Juergen Schoenwaelder

*Mahesh Jethanandani

*Michael Wang

*Qin Wu

*Reshad Rahman

*Rob Wilton

The ideas for a tooling based comparison of YANG module revisions was first described in [[I-D.clacla-netmod-yang-model-update](#)]. This document extends upon those initial ideas.

9. Security Considerations

The document does not define any new protocol or data model. There are no security impacts.

10. IANA Considerations

10.1. YANG Module Registrations

The following YANG module is requested to be registered in the "IANA Module Names" registry:

The ietf-yang-rev-annotations module:

Name: ietf-yang-rev-annotations

XML Namespace: urn:ietf:params:xml:ns:yang:ietf-yang-rev-annotations

Prefix: rev-ext

Reference: [RFCXXXX]

11. References

11.1. Normative References

[I-D.ietf-netmod-yang-module-versioning]

Wilton, R., Rahman, R., Lengyel, B., Clarke, J., Sterne, J., Claise, B., and K. D'Souza, "Updated YANG Module Revision Handling", Work in Progress, Internet-Draft, draft-ietf-netmod-yang-module-versioning-01, 10 July 2020, <<https://tools.ietf.org/html/draft-ietf-netmod-yang-module-versioning-01>>.

[I-D.ietf-netmod-yang-packages]

Wilton, R., Rahman, R., Clarke, J., Sterne, J., and W. Bo, "YANG Packages", Work in Progress, Internet-Draft, draft-ietf-netmod-yang-packages-00, 17 March 2020, <<https://tools.ietf.org/html/draft-ietf-netmod-yang-packages-00>>.

[I-D.ietf-netmod-yang-semver]

Claise, B., Clarke, J., Rahman, R., Wilton, R., Lengyel, B., Sterne, J., and K. D'Souza, "YANG Semantic Versioning", Work in Progress, Internet-Draft, draft-ietf-netmod-yang-semver-01, 13 July 2020, <<https://tools.ietf.org/html/draft-ietf-netmod-yang-semver-01>>.

[I-D.ietf-netmod-yang-solutions]

Wilton, R., "YANG Versioning Solution Overview", Work in Progress, Internet-Draft, draft-ietf-netmod-yang-solutions-00, 19 March 2020, <<https://tools.ietf.org/html/draft-ietf-netmod-yang-solutions-00>>.

[I-D.ietf-netmod-yang-versioning-reqs]

Clarke, J., "YANG Module Versioning Requirements", Work in Progress, Internet-Draft, draft-ietf-netmod-yang-versioning-reqs-03, 29 June 2020, <<https://tools.ietf.org/html/draft-ietf-netmod-yang-versioning-reqs-03>>.

[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>>.

[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>>.

11.2. Informative References

[I-D.clacla-netmod-yang-model-update]

Claise, B., Clarke, J., Lengyel, B., and K. D'Souza, "New YANG Module Update Procedure", Work in Progress, Internet-Draft, draft-clacla-netmod-yang-model-update-06,

2 July 2018, <<https://tools.ietf.org/html/draft-clacla-netmod-yang-model-update-06>>.

Author's Address

Robert Wilton
Cisco Systems, Inc.

Email: rwilton@cisco.com