

NMOP Working Group
Internet-Draft
Intended status: Standards Track
Expires: 8 January 2026

T. Hu
CMCC
L. M. Contreras
Telefonica
Q. Wu
Huawei
N. Davis
Ciena
C. Feng
7 July 2025

A YANG Data Model for Network Incident Management
draft-ietf-nmop-network-incident-yang-05

Abstract

A network incident refers to an unexpected interruption of a network service, degradation of a network service quality, or sub-health of a network service. Different data sources including alarms, metrics, and other anomaly information can be aggregated into a few of network incidents through data correlation analysis and the service impact analysis.

This document defines a YANG Module for the network incident lifecycle management. This YANG module is meant to provide a standard way to report, diagnose, and help resolve network incidents for the sake of network service health and probable cause analysis.

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 8 January 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.	Conventions and Definitions
3.	Sample Use Cases
3.1.	Incident-Based Trouble Tickets Dispatching
3.2.	Incident Derivation from L3VPN Services Unavailability
3.3.	Multi-layer Fault Demarcation
4.	Network Incident Management Architecture
4.1.	Interworking with Alarm Management
4.2.	Interworking with SAIN
4.3.	Relationship with RFC8969
4.4.	Relationship with Trace Context
5.	Functional Interface Requirements between the Client and the Server
5.1.	Incident Identification
5.2.	Incident Diagnosis
5.3.	Incident Resolution
6.	Incident Data Model Concepts
6.1.	Identifying the Incident Instance
6.2.	The Incident Lifecycle
6.2.1.	Network Incident Instance Lifecycle
6.2.2.	Operator Incident Lifecycle
7.	Incident Data Model Design
7.1.	Overview
7.2.	Incident Notifications
7.3.	Incident Acknowledge
7.4.	Incident Diagnose
7.5.	Incident Resolution
7.6.	RPC Failure
8.	Network Incident Management YANG Module
9.	Security Considerations
10.	IANA Considerations
10.1.	The "IETF XML" Registry
10.2.	The "YANG Module Names" Registry
Acknowledgements	
References	
	Normative References
	Informative References
Appendix A. Examples	
Appendix B. Changes between Revisions	
Contributors	
Authors' Addresses	

1. Introduction

[RFC8969] defines a framework for Automating Service and Network Management with YANG [RFC7950] to full life cycle network management. A set of YANG data models have already been developed in IETF for network performance monitoring and fault monitoring, e.g., a YANG data model for alarm management [RFC8632] defines a standard interface for alarm management. A data model for Network and VPN Service Performance Monitoring [RFC9375] defines a standard interface for network performance management. In addition, distributed tracing mechanism defined in [W3C-Trace-Context] can be used to analyze and debug operations, such as configuration transactions, across multiple distributed systems.

However, these YANG data models for network maintenance are based on specific data source information and manage alarms and performance metrics data separately at different layers in various different management systems. In addition, the frequency and quantity of alarms and performance metrics data reported to Operating Support System (OSS) are increased dramatically (in many cases multiple orders of magnitude) with the growth of service types and complexity and greatly overwhelm OSS platforms [TMF724A]; with existing known dependency relation between metric, alarm and events at each layer (e.g., packet layer or optical layer), it is possible to compress

series of alarms (see Section 3.5.3 of [RFC8632]) into fewer network incidents and there are many solutions in the market today that essentially do this to some degree. However, conventional solutions such as data compression are time-consuming and labor-intensive, usually rely on maintenance engineers' experience for data analysis, which, in many cases, result in low processing efficiency, inaccurate probable cause identification and duplicated tickets. It is also difficult to assess the impact of alarms, performance metrics and other anomaly data on network services without known relation across layers of the network topology data or the relation with other network topology data.

To address these challenges, a network wide incident-centric solution is specified to establish the dependency relation with both network service and network topology at various different layers, which not only can be used at a specific layer in one domain but also can be used to span across layers for multi-layer network troubleshooting.

A network incident refers to an undesired occurrence such as an unexpected interruption of a network service, degradation of a network service quality, or sub-health of a network service [I-D.ietf-nmop-terminology][TMF724A]. Different data sources including alarms, metrics, and other anomaly information can be aggregated into one or a few of network incidents irrespective of the layer through correlation analysis and the service impact analysis. For example, if the protocol-related interface fails to work properly, large amount of alarms may be reported to upper layer management system. Although a lot of network services may be affected by the interface, only one aggregated network incident pertaining to the abnormal interface will be reported. A network incident may also be raised through the analysis of some network performance metrics, for example, as described in SAIN [RFC9417], network services can be decomposed to several sub-services, specific metrics are monitored for each sub- service. Therefore symptoms will occur if services/sub-services are unhealthy (after analyzing metrics), in addition, these symptoms may give rise to a network incident when it causes degradation of the network services.

In addition, Artificial Intelligence (AI) and Machine Learning (ML) are key technologies in the processing of large amounts of data with complex data correlations (see Section 6.1 of [I-D.irtf-nmrg-ai-challenges]). For example, Neural Network Algorithm or Hierarchy Aggregation Algorithm can be used to replace manual alarm data correlation. Through online and offline self-learning, these algorithms can be continuously optimized to improve the efficiency of fault diagnosis.

This document defines a YANG data model for network incident lifecycle management, which improves troubleshooting efficiency, and improves network automation [RFC8969] with RPC operations in this YANG module.

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.

The following terms are defined in [RFC8632], [RFC9543],[I-D.ietf-nmop-terminology] and are not redefined here:

- * alarm
- * resource

- * event
- * problem
- * incident
- * anomaly
- * cause
- * SLA (Service Level Agreement)
- * SLO (Service Level Objective)

The following terms are defined in this document:

Service impact analysis: A process that uses algorithmic techniques (e.g., machine learning, automated reasoning, conformance checking, graph traversal, among others) to evaluate whether the network service has been impacted by the network incident and map network incident to one or a set of network service, which can reduce massive fault/alarms reporting, speed up troubleshooting, and assure network service performance and availability.

Incident management: Lifecycle management of network incidents, including network incident identification, reporting, acknowledgement, diagnosis, and resolution. Different from the traditional fault management, it take various different data sources including alarms, metrics, and other anomaly information and aggregate them into one or a few network incidents irrespective of layer through correlation analysis and the service impact analysis. One fault on the network device can cause multiple network incidents, e.g., multiple service offerings that are dependent on that device will go down and others may suffer increased latency as redundant routes become more congested. A network incident might impact one or a set of network services. The network incident can also been seen as customer incident [TMF724A] when service SLA [RFC9543] associated with one specific network service and network incident has been affected. How customer incident is translated from the network incident is beyond scope of this document.

Incident management system: An entity which implements network incident management. It includes (but not limited to) incident server and incident client.

Incident server: An entity which is responsible for detecting and reporting one network incident, performing network incident diagnosis, resolution and prediction, etc.

Incident client: An entity which can manage network incidents. For example, it can receive network incident notifications, query the information of network incidents, instruct an incident management server to diagnose, help resolve, etc.

Incident handler: An entity which can receive network incident notification, store and query the information of network incidents for data analysis. It has no control on incident server.

Probable root cause: A factor is considered the probable root cause of a problem if removing it fully resolves an ongoing incident -- specifically regarding network or service impairments and their related consequential failures and symptoms -- and prevents the problem from recurring. Causality is defined as the implementation of an inductive process at the network or

processing element level [ITU-T-G-7710] and as a doubt removal process at the Incident management system level [ITU-T-X-733]. The inductive process is based on the distinction made between primary failures (i.e., a failure that directly indicates the fault location and initiate a repair action, such as a broken cable or a misconnection) and secondary failures (i.e., a consequential failure, such as an upper level service that has gone down) [ITU-T-G-7710] Since one Fault may give rise to another Fault or Problem, a probable root cause is commonly meant to describe the original event or combination of circumstances that is the foundation of all related Faults.

Conversely, a causal factor is a contributing action that influences the outcome of an incident or event but is not the probably cause.

3. Sample Use Cases

3.1. Incident-Based Trouble Tickets Dispatching

Usually, the dispatching of trouble tickets in a network is mostly based on alarms data analysis and needs to involve operators' maintenance engineers. These operators' maintenance engineers are responsible to monitor and detect and correlate some alarms, e.g., that alarms at both endpoints of a specific tunnel or at both optical and IP layers which are associated with the same network fault. Therefore, they can correlate these alarms to the same trouble ticket, which is in the low automation. If there are more alarms, then the human costs for network maintenance are increased accordingly.

Some operators preconfigure accept-lists and adopt some coarse granularity data correlation rules for the alarm management. This approach seems to improve fault management automation. However, some trouble tickets might be missed if the filtering conditions are too restrictive. If the filtering conditions are not restrictive, it might end up with multiple trouble tickets being dispatched to the same network fault. It is hard to achieve a perfect balance between the network management automation and duplicated trouble tickets under the conventional working situations.

With the help of the network incident management, massive alarms can be aggregated into a few network incidents based on service impact analysis, so the number of trouble tickets will be reduced. At the same time, the efficiency of network troubleshooting can be largely improved, which address the pain point of traditional trouble ticket dispatching.

3.2. Incident Derivation from L3VPN Services Unavailability

The Service Attachment Points (SAPs) defined in [RFC9408] represent the network reference points where network services can be delivered or are being delivered to customers.

SLOs [RFC9543] can be used to characterize the ability of a particular set of nodes to communicate according to certain measurable expectations [I-D.ietf-ippm-pam]. For example, an SLA might state that any given SLO applies to at least a certain percentage of packets, allowing for a certain level of packet loss and exceeding packet delay threshold to take place. For example, an SLA might establish a multi-tiered SLO of end-to-end latency as follows:

- * Not to exceed 30 ms for any packet.
- * Not to exceed 25 ms for 99.999% of packets.

* Not to exceed 20 ms for 99% of packets.

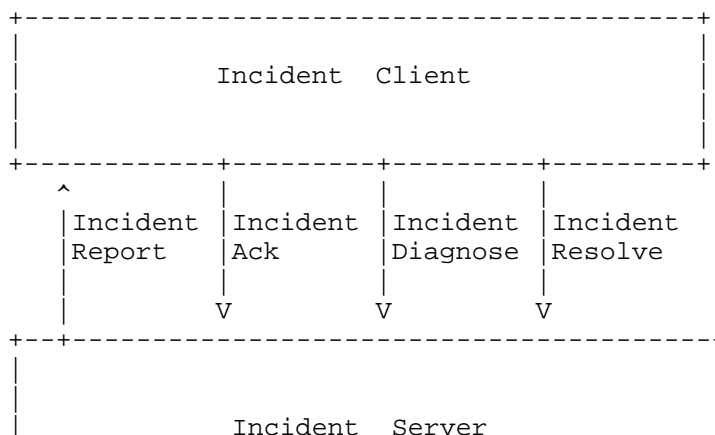
This SLA information can be bound with one or more SAPs defined in [RFC9408], so that the service orchestration layer can use these interfaces to commit the delivery of a service on specific point-to-point service topology or point to multi-point topology. When specific levels of a threshold of an SLO is violated, a specific network incident or customer incident [TMF724A], associated with, let's say L3VPN service will be derived.

3.3. Multi-layer Fault Demarcation

When a fault occurs in a network that contains both packet-layer devices and optical-layer devices, it may cause correlative faults in both layers, i.e., packet layer and optical layer. Specifically, faults propagation could be classified into three typical types. First, faults occurring at a packet-layer device might further cause fault (e.g., Wavelength Division Multiplexing (WDM) client fault) at an optical-layer device. Second, faults occurring at an optical-layer device might further cause fault (e.g., Layer 3 link down) at a packet-layer device. Third, faults occurring at the inter-layer link between a packet-layer device and an optical-layer device might further cause faults at both devices. Multiple operation teams are usually needed to first analyze huge amount of alarms (triggered by the above mentioned faults) from single network layer (either packet layer or optical layer) independently, then cooperate to locate the probable root cause through manually analyzing multi-layer topology data and service data, thus fault demarcation becomes more complex and time-consuming in multi-layer scenario than in single-layer scenario.

With the help of network incident management, the management systems first automatically analyze probable root cause of the alarms at each layer and report corresponding network incidents to the multi-layer, multi-domain management system, then such management system comprehensively analyzes the topology relationship and service relationship between the probable root causes of both layers. The inner relationship among the alarms will be identified and finally the probable root cause will be located among multiple layers. By cooperating with the integrated Optical time-domain reflectometer (OTDR) within the network device, we can determine the target optical exchange station before site visits. Therefore, the overall fault demarcation process is simplified and automated, the analyze result could be reported and visualized in time. In this case, operation teams only have to confirm the analyzing result and dispatch site engineers to perform relative maintenance actions (e.g., splice fiber) based on the probable root cause.

4. Network Incident Management Architecture



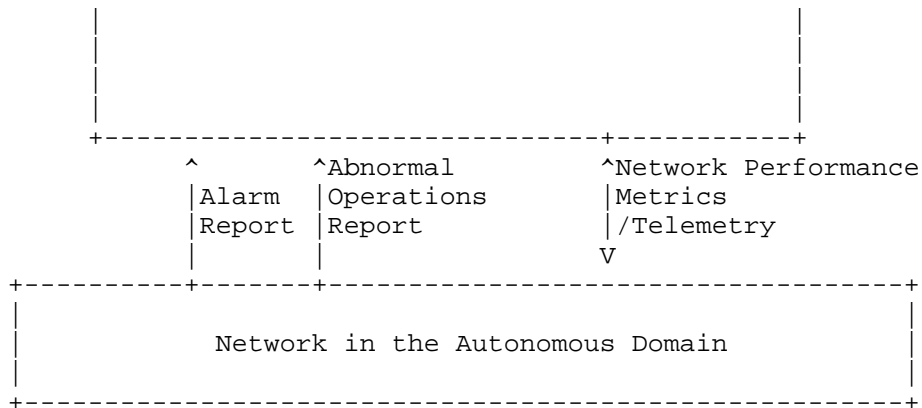


Figure 1: Network Incident Management Architecture

Figure 1 illustrates the network incident management architecture. Two key components for the incident management are incident client and incident server.


Incident server can be deployed in network operation platforms, network analytic platforms, controllers [RFC8969] and provides functionalities such as network incident identification, report, diagnosis, resolution, or querying for network incident lifecycle management.

Incident client can be deployed either in the same network operation platforms, network analytic platforms, controllers as the incident server within a single domain, or at the upper layer network operation platforms, network analytic platforms or controllers (i.e., multi-domain controllers), to invoke the functionalities provided by incident server to meet the business requirements of fault management. The entire network incident lifecycle management can be independent from or not under control of the network OSS or other business system of operators.

A typical workflow of network incident management is as follows:

- * Some alarm report or abnormal operations, network performance metrics are reported from the network. Incident server receives these alarms/abnormal operations/metrics and try to analyze the correlation of them, e.g., generate a symptom if some metrics are evaluated as unhealthy, the probable root cause can be detected based on the correlation analysis. If a network incident is identified, the "incident report" notification will be reported to the incident client. The impact of network services will be further analyzed and will update the network incident if the network service is impacted.
- * Incident client receives the network incident from "incident report" notification raised by incident server, and acknowledge it with "incident ack" rpc operation. Client may further invoke the "incident diagnose" rpc to diagnose this network incident to find the probable root causes.
- * If the probable root causes have been found, the incident client can resolve this network incident by invoking the 'incident resolve' rpc operation, dispatching a ticket or using other network functions (routing calculation, configuration, etc.)

4.1. Interworking with Alarm Management



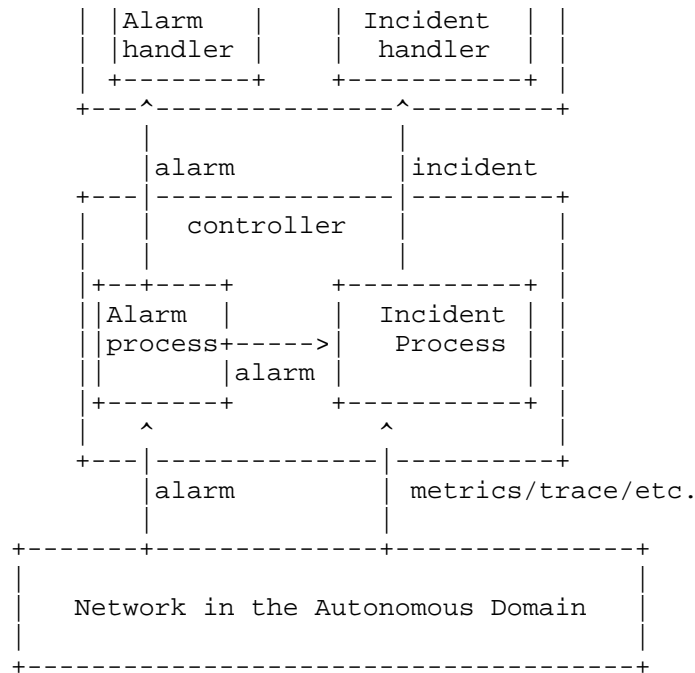


Figure 2: Interworking with Alarm Management

A YANG model for the alarm management [RFC8632] defines a standard interface to manage the lifecycle of alarms. Alarms represent the undesirable state of network resources [I-D.ietf-nmop-terminology], alarm data model also defines the probable root causes and impacted services fields, but there may lack sufficient information to determine them in lower layer system (mainly in devices level), so alarms do not always tell the status of network services or necessarily point to the probable root causes of problems. As described in [RFC8632], alarm management act as a starting point for high-level fault management. While network incident management often works at the network level, so it is possible to have enough information to perform correlation and service impact analysis. Alarms can work as one of data sources of network incident management and may be aggregated into few network incidents by correlation analysis, network service impact and probable root causes may be determined during incident process.

Network Incident also contains some related alarms, if needed users can query the information of alarms by alarm management interface [RFC8632]. In some cases, e.g., cutover scenario, incident server may use alarm management interface [RFC8632] to shelve some alarms.

Alarm management may keep the original process, alarms are reported from network to network controller or network analytic platform and then reported to upper layer system (e.g., the alarm handler within the OSS).

Similarly, the network incident is reported from the network to the network controller or network analytic platform and then reported to the upper layer system (e.g., incident handler within the OSS). Upper layer system may store these network incidents and provide the information for fault analysis (e.g., deeper customer incident analysis based on network incident).

Different from alarm management, incident process within the controller comprising both incident client and incident sever functionalities provides not only network incident reporting but also diagnosis and resolution functions, it's possible to support self-healing and may be helpful for single-domain closed-loop control.

Incident management is not a substitute for alarm management. Instead, they can work together to implement fault management.

4.2. Interworking with SAIN

SAIN [RFC9417] defines an architecture of network service assurance.

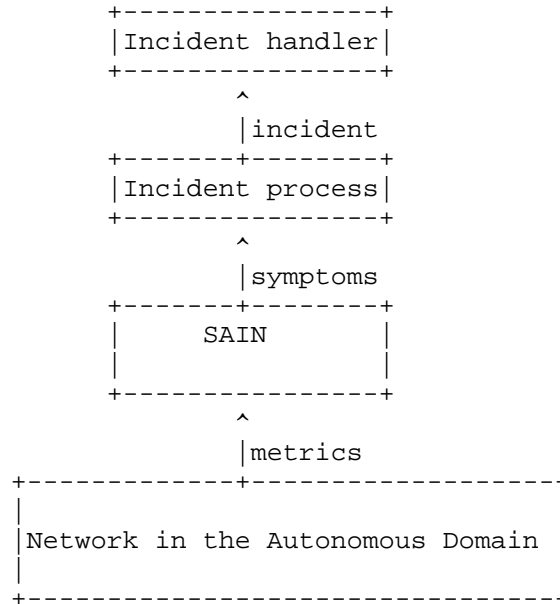


Figure 3: Interworking with SAIN

A network service can be decomposed into some sub-services, and some metrics can be monitored for sub-services. For example, a tunnel service can be decomposed into some peer tunnel interface sub-services and IP connectivity sub-service. If some metrics are evaluated to indicate unhealthy for specific sub-service, some symptoms will be present. Incident process comprising both incident client and incident server functionalities may identify the network incident based on symptoms, and then report it to incident handler within the Operation Support System (OSS). So, SAIN can be one way to identify network incident, services, sub-services and metrics can be preconfigured via APIs defined by service assurance YANG model [RFC9418] and the network incident will be reported if symptoms match certain condition or characteristic considered as an indication of a problem or potential problem.

4.3. Relationship with RFC8969

[RFC8969] defines a framework for network automation using YANG, this framework breaks down YANG modules into three layers, service layer, network layer and device layer, and contains service deployment, service optimization/assurance, and service diagnosis. Network incident works at the network layer and aggregates alarms, metrics and other information from device layer, it's helpful to provide service assurance. And the network incident diagnosis may be one way of service diagnosis.

4.4. Relationship with Trace Context

W3C defines a common trace context [W3C-Trace-Context] for distributed system tracing, [I-D.ietf-netconf-trace-ctx-extension] defines a netconf extension for [W3C-Trace-Context] and [I-D.ietf-netconf-configuration-tracing] defines a mechanism for configuration tracing. If some errors occur when services are deploying, it's very easy to identify these errors by distributed system tracing, and a network incident should be reported.

5. Functional Interface Requirements between the Client and the Server

5.1. Incident Identification

As depicted in Figure 4, multiple alarms, metrics, or hybrid can be aggregated into a network incident after analysis.




Figure 4: Incident Identification

The network incident management server MUST be capable of identifying network incidents. Multiple alarms, metrics and other information are reported to incident server, and the server must analyze it and find out the correlations of them, if the correlation match the network incident rules, network incident will be identified, and reported to the client. If the network incident is repeated many times, the problem needs to be raised. Service impact analysis SHOULD be performed if a network incident is identified, and the content of network incident SHOULD be updated if impacted network services are detected.

AI/ML may be used to identify the network incident. Expert system and online learning can help AI to identify the correlation of alarms, metrics and other information by time-base correlation algorithm, topology-based correlation algorithm, etc. For example, if the interface is down, then many protocol alarms will be reported, AI will think these alarms have some correlations. These new correlations will be put into the knowledge base, and the network incident will be identified faster according to knowledge base next time.




Figure 5: Example 1 of Network Incident Identification

As described in Figure 5, vpn a is deployed from PE1 to PE2, if a interface of P1 is going down, many alarms are triggered, such as interface down, igp down, and igp peer abnormal from P2.

These alarms are aggregated and analyzed by the controller/incident server, and then the network incident 'vpn unavailable' is triggered by the controller/incident server. If the network incident 'vpn unavailable' is repeated, the problem can be raised.

Note that incident server within the controller can rely on data correlation technology such as service impact analysis and data analytic component to evaluate the real effect on the relevant service and understand whether lower level or device level network anomaly, e.g., igp down, has impact on the service.




Figure 6: Example 2 of Network Incident Identification

As described in Figure 6, controller collect the network metrics from network elements, it finds the packet loss of P1 and the path delay of P2 exceed the thresholds, a network incident 'VPN A degradation' may be triggered after the service impact analysis.

5.2. Incident Diagnosis

After a network incident is reported to the network incident client, the incident client MAY diagnose the incident to determine the probable root cause. Some diagnosis operations may affect the running network services. The incident client can choose not to perform that diagnosis operation after determining the impact is not trivial. The incident server can also perform self-diagnosis. However, the self-diagnosis MUST not affect the running network services. Possible diagnosis methods include link reachability detection, link quality detection, alarm/log analysis, and short-term fine-grained monitoring of network quality metrics, etc.

5.3. Incident Resolution

After the probable root cause is diagnosed, the incident client MAY resolve the network incident. The incident client MAY choose resolve the network incident by invoking other functions, such as routing calculation function, configuration function, dispatching a ticket or asking the server to resolve it. Generally, the incident client would attempt to directly resolve the probable cause. If the probable root cause cannot be resolved, an alternative solution SHOULD be required. For example, if a network incident caused by a physical component failure, it cannot be automatically resolved, the standby link can be used to bypass the faulty component.

Incident server will monitor the status of the network incident, if the faults are fixed, the incident server will update the status of network incident to 'cleared', and report the updated network incident to the client.

Network incident resolution may affect the running network services. The client can choose not to perform those operations after determining the impact is not trivial.

6. Incident Data Model Concepts

6.1. Identifying the Incident Instance

An incident id is used as an identifier of an incident instance, if an incident instance is identified, a new incident ID is created. The incident id MUST be unique in the whole system.

6.2. The Incident Lifecycle

The network incident model clearly separates network incident instance lifecycle from operator incident lifecycle.

- o Network incident instance lifecycle: The network incident instrumentation that control network incident raised, updated and cleared.

- o Operator incident lifecycle: Operators acting upon the network incident with rpcs like acknowledged, diagnosed and resolved.

6.2.1. Network Incident Instance Lifecycle

From a network incident instance perspective, a network incident can

have the following lifecycle: 'raised', 'updated', 'cleared'. When a network incident instance is first generated, the status is 'raised'. If the status changes after the network incident instance is generated, (for example, self-diagnosis, diagnosis command issued by the client, or any other condition causes the status to change but does not reach the 'cleared' level.) , the status changes to 'updated'. When a network incident is successfully resolved, the status changes to 'cleared'.

6.2.2. Operator Incident Lifecycle

Operators can act upon network incident with network incident rpcs. From an operator perspective, the lifecycle of a network incident instance includes 'acknowledged', 'diagnosed', and 'resolved'.

When a network incident instance is generated, the operator SHOULD acknowledge the network incident with 'incident-acknowledge' rpc. And then the operator attempts to diagnose the network incident with 'incident-diagnose' rpc (for example, find out the probable root cause and affected components). Diagnosis is not mandatory. If the probable root cause and affected components are known when the network incident is generated, diagnosis is not required. After locating the probable root cause and affected components, operator can try to resolve the network incident by invoking 'incident-resolve' rpc.

7. Incident Data Model Design

7.1. Overview

There is one YANG module in the "ietf-incident" model, which defines technology independent abstraction of network incident construct for alarm, log, performance metrics, etc. The information reported in the network incident include probable root cause, priority, impact, suggestion, etc.

At the top of "ietf-incident" module is the Network Incident. Network incident is represented as a list and indexed by "incident-id". Each Network Incident is associated with a network service instance, domain and sources. Under sources, there is one or more sources. Each source corresponds to node defined in the network topology model and network resource in the network device, e.g., interface. In addition, "ietf-incident" supports one general notification to report network incident state changes and three rpcs to manage the network incidents.

```
module: ietf-incident
  +--ro incidents
    +--ro incident* [name type incident-id]
      +--ro incident-no      uint64
      +--ro name             string
      +--ro type             identityref
      +--ro incident-id?    string
      +--ro service-instance* string
      +--ro domain           identityref
      +--ro priority         incident-priority
      +--ro status?         enumeration
      +--ro ack-status?     enumeration
      +--ro category        identityref
      +--ro detail?         string
      +--ro resolve-advice? String
      +--ro sources
      ...
      +--ro probable-causes
      ...
      +--ro probable-events
```

```

...
+--ro events
...
+--ro raise-time? yang:date-and-time
+--ro occur-time? yang:date-and-time
+--ro clear-time? yang:date-and-time
+--ro ack-time? yang:date-and-time
+--ro last-updated? yang:date-and-time
rpcs:
+---x incident-acknowledge
...
+---x incident-diagnose
...
+---x incident-resolve
notifications:
+---n incident-notification
+--ro incident-no?
-> /inc:incidents/inc:incident/inc:incident-no
...
+--ro time? yang:date-and-time

```

7.2. Incident Notifications

```

notifications:
+---n incident-notification
+--ro incident-no?
-> /inc:incidents/inc:incident/inc:incident-no
+--ro name? string
+--ro type? identityref
+--ro incident-id? string
+--ro service-instance* string
+--ro domain? identityref
+--ro priority? int:incident-priority
+--ro status? enumeration
+--ro ack-status? enumeration
+--ro category? identityref
+--ro detail? string
+--ro resolve-advice? string
+--ro sources
| +--ro source* [node-ref]
| | +--ro node-ref leafref
| | +--ro network-ref? -> /nw:networks/network/network-id
| | +--ro resource* [name]
| | +--ro name al:resource
+--ro probable-causes
| +--ro probable-cause* [node-ref]
| | +--ro node-ref leafref
| | +--ro network-ref? -> /nw:networks/network/network-id
| | +--ro resource* [name]
| | | +--ro name al:resource
| | | +--ro cause-name? string
| | | +--ro detail? string
| | +--ro cause-name? string
| | +--ro detail? string
+--ro probable-events
| +--ro probable-event* [type event-id]
| | +--ro type -> ../../../../events/event/type
| | +--ro event-id leafref
+--ro events
| +--ro event* [type event-id]
| | +--ro type enumeration
| | +--ro event-id string
| | +--ro (event-type-info)?
| | | +--:(alarm)
| | | +--ro alarm

```

```

|           |      +--ro resource? leafref
|           |      +--ro alarm-type-id? leafref
|           |      +--ro alarm-type-qualifier? leafref
+--ro time? yang:date-and-time

```

A general notification, incident-notification, is provided here. When a network incident instance is identified, the notification will be sent. After a notification is generated, if the incident server performs self diagnosis or the incident client uses the interfaces provided by the incident server to deliver diagnosis and resolution actions, the notification update behavior is triggered, for example, the probable root cause objects and affected objects are updated. When a network incident is successfully resolved, the status of the network incident would be set to 'cleared'.

7.3. Incident Acknowledge

```

+---x incident-acknowledge
|   +---w input
|   |   +---w incident-no*
|   |   -> /inc:incidents/inc:incident/inc:incident-no

```

After an incident is generated, updated, or cleared, the operator needs to confirm the incident to ensure that the client knows the incident.

In some scenarios where automatic diagnosis and resolution are supported, the status of an incident may be updated multiple times or even automatically resolved. Therefore the incident-acknowledge rpc can confirm multiple incidents at a time.

7.4. Incident Diagnose

```

+---x incident-diagnose
|   +---w input
|   |   +---w incident-no*
|   |   -> /inc:incidents/inc:incident/inc:incident-no

```

After a network incident is generated, network incident diagnose rpc can be used to diagnose the network incident and locate the probable root causes. On demand Diagnosis can be performed on some detection tasks, such as bfd detection, flow detection, telemetry collection, short-term threshold alarm, configuration error check, or test packet injection.

After the on demand diagnosis is performed successfully, a separate network incident update notification will be triggered to report the latest status of the network incident asynchronously.

7.5. Incident Resolution

```

+---x incident-resolve
|   +---w input
|   |   +---w incident-no*
|   |   -> /inc:incidents/inc:incident/inc:incident-no

```

After the probable root causes and impacts are determined, incident-resolve rpc can be used to resolve the incident (if the server can resolve it). How to resolve an incident instance is out of the scope of this document.

Network incident resolve rpc allows multiple network incident instances to be resolved at a time. If a network incident instance is successfully resolved, a separate notification will be triggered to update the network incident status to 'cleared'. If the network incident content is changed during this process, a notification

update will be triggered.

7.6. RPC Failure

If the RPC fails, the RPC error response MUST indicate the reason for the failure. The structures defined in this document MUST encode specific errors and be inserted in the error response to indicate the reason for the failure.

The tree diagram [RFC8340] for structures are defined as follows:

```
structure incident-acknowledge-error-info:
  +-- incident-acknowledge-error-info
    +-- incident-no?    incident-ref
    +-- reason?        identityref
    +-- description?   string
structure incident-diagnose-error-info:
  +-- incident-diagnose-error-info
    +-- incident-no?    incident-ref
    +-- reason?        identityref
    +-- description?   string
structure incident-resolve-error-info:
  +-- incident-resolve-error-info
    +-- incident-no?    incident-ref
    +-- reason?        identityref
    +-- description?   string
```

Valid errors that can occur for each structure defined in this document are described as follows:

```
incident-acknowledge-error-info
-----
repeated-acknowledge

incident-diagnose-error-info
-----
probable-cause-unlocated
permission-denied
operation-timeout
resource-unavailable

incident-resolve-error-info
-----
probable-cause-unresolved
permission-denied
operation-timeout
resource-unavailable
```

8. Network Incident Management YANG Module

This module imports types defined in [RFC6991], [RFC8345], [RFC8632].

```
<CODE BEGINS> file "ietf-incident@2024-06-06.yang"
module ietf-incident {
  yang-version 1.1;
  namespace "urn:ietf:params:xml:ns:yang:ietf-incident";
  prefix inc;

  import ietf-yang-types {
    prefix yang;
    reference
      "RFC 6991: Common YANG Data Types";
  }
  import ietf-alarms {
    prefix al;
    reference
```



```

    "RFC 8632: A YANG Data Model for Alarm Management";
}
import ietf-network {
    prefix nw;
    reference
        "RFC 8345: A YANG Data Model for Network Topologies";
}
import ietf-yang-structure-ext {
    prefix sx;
}
organization
    "IETF NMOP Working Group";
contact
    "WG Web:    <https://datatracker.ietf.org/wg/nmop/>;
    WG List:    <mailto:nmop@ietf.org>

    Author:     Chong Feng
                <mailto:frank.fengchong@huawei.com>
    Author:     Tong Hu
                <mailto:hutong@cmhi.chinamobile.com>
    Author:     Luis Miguel Contreras Murillo
                <mailto:luismiguel.contrerasmurillo@telefonica.com>
    Author :    Qin Wu
                <mailto:bill.wu@huawei.com>
    Author:     Chaode Yu
                <mailto:yuchaode@huawei.com>
    Author:     Nigel Davis
                <mailto:ndavis@ciena.com>";

description
    "This module defines the interfaces for incident management
    lifecycle.

    This module is intended for the following use cases:
    * incident lifecycle management:
        - incident report: report incident instance to client
                        when an incident instance is detected.
        - incident acknowledge: acknowledge an incident instance.
        - incident diagnose: diagnose an incident instance.
        - incident resolve: resolve an incident instance.

    Copyright (c) 2024 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 Revised BSD License
    set forth in Section 4.c of the IETF Trust's Legal Provisions
    Relating to IETF Documents
    (https://trustee.ietf.org/license-info).

    This version of this YANG module is part of RFC XXXX
    (https://www.rfc-editor.org/info/rfcXXXX); ; 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. ";

revision 2024-06-06 {
    description
        "Merge incident yang with incident type yang
        and fix broken ref.";
    reference
        "RFC XXX: YANG module for network incident management.";
}

```

```

}

// Identities

identity incident-domain {
    description
        "The abstract identity to indicate the domain of
        an incident.";
}

identity single-domain {
    base incident-domain;
    description
        "single domain.";
}

identity access {
    base single-domain;
    description
        "access domain.";
}

identity ran {
    base access;
    description
        "radio access network domain.";
}

identity transport {
    base single-domain;
    description
        "transport domain.";
}

identity otn {
    base transport;
    description
        "optical transport network domain.";
}

identity ip {
    base single-domain;
    description
        "ip domain.";
}

identity ptn {
    base ip;
    description
        "packet transport network domain.";
}

identity cross-domain {
    base incident-domain;
    description
        "cross domain.";
}

identity incident-category {
    description
        "The abstract identity for incident category.";
}

identity device {
    base incident-category;
    description

```

```

        "device category.";
    }

identity power-environment {
    base device;
    description
        "power environment category.";
}

identity device-hardware {
    base device;
    description
        "hardware of device category.";
}

identity device-software {
    base device;
    description
        "software of device category";
}

identity line {
    base device-hardware;
    description
        "line card category.";
}

identity maintenance {
    base incident-category;
    description
        "maintenance category.";
}

identity network {
    base incident-category;
    description
        "network category.";
}

identity protocol {
    base incident-category;
    description
        "protocol category.";
}

identity overlay {
    base incident-category;
    description
        "overlay category";
}

identity vm {
    base incident-category;
    description
        "vm category.";
}

identity event-type {
    description
        "The abstract identity for Event type";
}

identity alarm {
    base event-type;
    description
        "alarm event type.";
}

```

```

}

identity notif {
    base event-type;
    description
        "Notification event type.";
}

identity log {
    base event-type;
    description
        "Log event type.";
}

identity KPI {
    base event-type;
    description
        "KPI event type.";
}

identity unknown {
    base event-type;
    description
        "Unknown event type.";
}

identity incident-class {
    description
        "The abstract identity for Incident category.";
}

identity problem {
    base incident-class;
    description
        "It indicates the class of the incident is a problem
        (i.e., cause of the incident) for example an interface
        fails to work.";
}

identity sla-violation {
    base incident-class;
    description
        "It indicates the class of the incident is a sla
        violation, for example high CPU rate may cause
        a fault in the future.";
}

identity acknowledge-error {
    description
        "Base identity for the problem found while attempting
        to fulfill an 'incident-acknowledge' RPC request.";
}

identity diagnose-error {
    description
        "Base identity for the problem found while attempting
        to fulfill an 'incident-diagnose' RPC request.";
}

identity resolve-error {
    description
        "Base identity for the problem found while attempting
        to fulfill an 'incident-resolve' RPC request.";
}

identity repeated-acknowledge {

```

```

    base acknowledge-error;
    description
        "The incident referred to has already been acknowledged.";
}

identity probable-cause-unlocated {
    base diagnose-error;
    description
        "Fails to locate the probable causes when performing the
        diagnosis operation. The detailed reason MUST be included
        in the 'description'.";
}

identity probable-cause-unresolved {
    base resolve-error;
    description
        "Fails to resolve the probable causes when performing the
        resolution operation. The detailed reason MUST be included
        in the 'description'";
}

identity permission-denied {
    base diagnose-error;
    base resolve-error;
    description
        "The permission required for performing specific
        detection/resolution task is not granted.";
}

identity operation-timeout {
    base diagnose-error;
    base resolve-error;
    description
        "The diagnosis/resolution time exceeds the preset time.";
}

identity resource-unavailable {
    base diagnose-error;
    base resolve-error;
    description
        "The resource is unavailable to perform
        the diagnosis/resolution operation.";
}

identity cause-name {
    description
        "Base identity for the cause name.";
}

// Typedefs

typedef incident-priority {
    type enumeration {
        enum critical {
            description
                "the incident MUST be handled immediately.";
        }
        enum high {
            description
                "the incident should be handled as soon as
                possible.";
        }
        enum medium {
            description
                "network services are not affected, or the
                services are slightly affected, but corrective

```

```

        measures need to be taken.";
    }
    enum low {
        description
            "potential or imminent service-affecting
            incidents are detected, but services are
            not affected currently.";
    }
}
description
    "define the priority of incident.";
}

typedef incident-ref {
    type leafref {
        path "/inc:incidents/inc:incident/inc:incident-no";
    }
    description
        "reference a network incident.";
}

// Groupings

grouping probable-cause-info {
    description
        "The information of probable cause.";
    leaf cause-name {
        type identityref {
            base cause-name;
        }
        description
            "the name of cause.";
    }
    leaf detail {
        type string;
        description
            "the detail information of the cause.";
    }
}

grouping resources-info {
    description
        "the grouping which defines the network
        resources of a node.";
    uses nw:node-ref;
    list resource {
        key "name";
        description
            "the resources of a network node.";
        leaf name {
            type al:resource;
            description
                "network resource name.";
        }
    }
}

grouping incident-time-info {
    description
        "the grouping defines incident time information.";
    leaf raise-time {
        type yang:date-and-time;
        description
            "the time when an incident instance is raised.";
    }
    leaf occur-time {

```

```

    type yang:date-and-time;
    description
        "the time when an incident instance occurs.
        It's the occur time of the first event during
        incident detection.";
}
leaf clear-time {
    type yang:date-and-time;
    description
        "the time when an incident instance is
        resolved.";
}
leaf ack-time {
    type yang:date-and-time;
    description
        "the time when an incident instance is
        acknowledged.";
}
leaf last-updated {
    type yang:date-and-time;
    description
        "the latest time when an incident instance is
        updated";
}
}

grouping incident-info {
    description
        "the grouping defines the information of an
        incident.";
    leaf name {
        type string;
        mandatory true;
        description
            "the name of an incident.";
    }
    leaf type {
        type identityref {
            base incident-class;
        }
        mandatory true;
        description
            "The type of an incident.";
    }
    leaf incident-id {
        type string;
        description
            "The unique qualifier of an incident instance type.
            This leaf is used when the 'type' leaf cannot
            uniquely identify the incident instance type. Normally,
            this is not the case, and this leaf is the empty string.";
    }
    leaf-list service-instance {
        type string;
        description
            "the related network service instances of
            the incident instance.";
    }
    leaf domain {
        type identityref {
            base incident-domain;
        }
        mandatory true;
        description
            "the domain of an incident.";
    }
}

```

```

leaf priority {
  type incident-priority;
  mandatory true;
  description
    "the priority of an incident instance.";
}
leaf status {
  type enumeration {
    enum raised {
      description
        "an incident instance is raised.";
    }
    enum updated {
      description
        "the information of an incident instance
          is updated.";
    }
    enum cleared {
      description
        "an incident is cleared.";
    }
  }
  default "raised";
  description
    "The status of an incident instance.";
}
leaf ack-status {
  type enumeration {
    enum acknowledged {
      description
        "The incident has been acknowledged by user.";
    }
    enum unacknowledged {
      description
        "The incident hasn't been acknowledged.";
    }
  }
  default "unacknowledged";
  description
    "the acknowledge status of an incident.";
}
leaf category {
  type identityref {
    base incident-category;
  }
  mandatory true;
  description
    "The category of an incident.";
}
leaf detail {
  type string;
  description
    "detail information of this incident.";
}
leaf resolve-advice {
  type string;
  description
    "The advice to resolve this incident.";
}
container sources {
  description
    "The source components.";
  list source {
    key "node-ref";
    min-elements 1;
    description

```



```

        "The source components of incident.";
    uses resources-info;
}
}
container probable-causes {
    description
        "The probable cause objects.";
    list probable-cause {
        key "node-ref";
        description
            "the probable causes of incident.";
        uses resources-info {
            augment "resource" {
                description
                    "augment probable cause information.";
                //if probable cause object is a resource of a node
                uses probable-cause-info;
            }
        }
        //if probable cause object is a node
        uses probable-cause-info;
    }
}
container probable-events {
    description
        "the probable cause related events of the incident.";
    list probable-event {
        key "type event-id";
        description
            "the probable cause related event of the incident.";
        leaf type {
            type leafref {
                path "../.../.../events/event/type";
            }
            description
                "the event type.";
        }
        leaf event-id {
            type leafref {
                path "../.../.../events/event[type = current()/../type]"
                    + "/event-id";
            }
            description
                "the event identifier, such as uuid,
                sequence number, etc.";
        }
    }
}
container events {
    description
        "related events.";
    list event {
        key "type event-id";
        description
            "related events.";
        leaf type {
            type identityref {
                base event-type;
            }
            description
                "event type.";
        }
        leaf event-id {
            type string;
            description
                "the event identifier, such as uuid,

```

```

        sequence number, etc.";
    }
    choice event-type-info {
        description
            "event type information.";
        case alarm {
            when "derived-from-or-self(type, 'alarm')" {
                description
                    "Only applies when type is alarm.";
            }
            container alarm {
                description
                    "alarm type event.";
                leaf resource {
                    type leafref {
                        path "/al:alarms/al:alarm-list/al:alarm"
                            + "/al:resource";
                    }
                    description
                        "network resource.";
                    reference
                        "RFC 8632: A YANG Data Model for Alarm
                            Management";
                }
                leaf alarm-type-id {
                    type leafref {
                        path "/al:alarms/al:alarm-list/al:alarm"
                            + "[al:resource = current()/../resource]"
                            + "/al:alarm-type-id";
                    }
                    description
                        "alarm type id";
                    reference
                        "RFC 8632: A YANG Data Model for Alarm
                            Management";
                }
                leaf alarm-type-qualifier {
                    type leafref {
                        path "/al:alarms/al:alarm-list/al:alarm"
                            + "[al:resource = current()/../resource]"
                            + "[al:alarm-type-id = current()/.."
                            + "/alarm-type-id]/al:alarm-type-qualifier";
                    }
                    description
                        "alarm type qualitifier";
                    reference
                        "RFC 8632: A YANG Data Model for Alarm
                            Management";
                }
            }
        }
    }
}

// RPCs

rpc incident-acknowledge {
    description
        "This rpc can be used to acknowledge the specified
            incidents.";
    input {
        leaf-list incident-no {
            type incident-ref;
            description

```

```

        "the identifier of an incident instance.";
    }
}

rpc incident-diagnose {
    description
        "This rpc can be used to diagnose the specified
        incidents. The result of diagnosis will be reported
        by incident notification.";
    input {
        leaf-list incident-no {
            type incident-ref;
            description
                "the identifier of an incident instance.";
        }
    }
}

rpc incident-resolve {
    description
        "This rpc can be used to resolve the specified
        incidents. The result of resolution will be reported
        by incident notification.";
    input {
        leaf-list incident-no {
            type incident-ref;
            description
                "the identifier of an incident instance.";
        }
    }
}

sx:structure incident-acknowledge-error-info {
    container incident-acknowledge-error-info {
        description
            "This structure data MAY be inserted in the RPC error
            response to indicate the reason for the
            incident acknowledge failure.";
        leaf incident-no {
            type incident-ref;
            description
                "Indicates the incident identifier that
                fails the operation.";
        }
        leaf reason {
            type identityref {
                base acknowledge-error;
            }
            description
                "Indicates the reason why the operation is failed.";
        }
        leaf description {
            type string;
            description
                "Indicates the detailed description about the failure.";
        }
    }
}

sx:structure incident-diagnose-error-info {
    container incident-diagnose-error-info {
        description
            "This structure data MAY be inserted in the RPC error
            response to indicate the reason for the
            incident diagnose failure.";
        leaf incident-no {

```

```

        type incident-ref;
        description
            "Indicates the incident identifier that
            fails the operation.";
    }
    leaf reason {
        type identityref {
            base diagnose-error;
        }
        description
            "Indicates the reason why the operation is failed.";
    }
    leaf description {
        type string;
        description
            "Indicates the detailed description about the failure.";
    }
}
}
sx:structure incident-resolve-error-info {
    container incident-resolve-error-info {
        description
            "This structure data MAY be inserted in the RPC error
            response to indicate the reason for the
            incident resolution failure.";
        leaf incident-no {
            type incident-ref;
            description
                "Indicates the incident identifier that
                fails the operation.";
        }
        leaf reason {
            type identityref {
                base resolve-error;
            }
            description
                "Indicates the reason why the operation is failed.";
        }
        leaf description {
            type string;
            description
                "Indicates the detailed description about the failure.";
        }
    }
}
}

```

// Notifications

```

notification incident-notification {
    description
        "incident notification. It will be triggered when
        the incident is raised, updated or cleared.";
    leaf incident-no {
        type incident-ref;
        description
            "the identifier of an incident instance.";
    }
    uses incident-info;
    leaf time {
        type yang:date-and-time;
        description
            "occur time of an incident instance.";
    }
}
}

```

// Data definitions

```

container incidents {
  config false;
  description
    "the information of incidents.";
  list incident {
    key "name type incident-id";
    description
      "the information of incident.";
    leaf incident-no {
      type uint64;
      mandatory true;
      description
        "The unique sequence number of the incident instance.";
    }
    uses incident-info;
    uses incident-time-info;
  }
}
}
<CODE ENDS>

```

9. Security Considerations

The YANG modules specified in this document define a schema for data that is designed to be accessed via network management protocol such as NETCONF [RFC6241] or RESTCONF [RFC8040]. The lowest NETCONF layer is the secure transport layer, and the mandatory-to-implement secure transport is Secure Shell (SSH) [RFC6242]. The lowest RESTCONF layer is HTTPS, and the mandatory-to-implement secure transport is TLS [RFC8446].

The Network Configuration Access Control Model (NACM) [RFC8341] provides the means to restrict access for particular NETCONF or RESTCONF users to a preconfigured subset of all available NETCONF or RESTCONF protocol operations and content.

Some of the readable data nodes in this YANG module may be considered sensitive or vulnerable in some network environments. It is thus important to control read access (e.g., via get, get-config, or notification) to these data nodes. These are the subtrees and data nodes and their sensitivity/vulnerability:

'/incidents/incident': This list specifies the network incident entries. Unauthorized read access of this list can allow intruders to access network incident information and potentially get a picture of the broken state of the network. Intruders may exploit the vulnerabilities of the network to lead to further negative impact on the network. Care must be taken to ensure that this list are accessed only by authorized users.

Some of the RPC operations in this YANG module may be considered sensitive or vulnerable in some network environments. It is thus important to control access to these operations. These are the operations and their sensitivity/vulnerability:

"incident-diagnose": This RPC operation performs network incident diagnosis and probable root cause locating. If a malicious or buggy client performs an unexpectedly large number of this operation, the result might be an excessive use of system resources [I-D.ietf-nmop-terminology] on the server side as well as network resources. Servers MUST ensure they have sufficient resources to fulfill this request; otherwise, they MUST reject the request using RPC errors defined in section 7.6.

"incident-resolve": This RPC operation is used to resolve the network

incident. If a malicious or buggy client performs an unexpectedly large number of this operation, the result might be an excessive use of system resources on the server side as well as network resources. Servers MUST ensure they have sufficient resources to fulfill this request; otherwise, they MUST reject the request without compromise on security of data-at-rest in the server.

10. IANA Considerations

10.1. The "IETF XML" Registry

This document requests IANA to register one XML namespace URN in the "ns" subregistry within the "IETF XML Registry" [RFC3688]:

URI: urn:ietf:params:xml:ns:yang:ietf-incident
Registrant Contact: The IESG.
XML: N/A, the requested URIs are XML namespaces.

10.2. The "YANG Module Names" Registry

This document requests IANA to register one module name in the 'YANG Module Names' registry, defined in [RFC6020].

Name: ietf-incident
Maintained by IANA? N
Namespace: urn:ietf:params:xml:ns:yang:ietf-incident
Prefix: inc
Reference: RFC XXXX
// RFC Ed.: replace XXXX and remove this comment

Acknowledgements

The authors would like to thank Mohamed Boucadair, Robert Wilton, Benoit Claise, Oscar Gonzalez de Dios, Adrian Farrel, Mahesh Jethanandani, Balazs Lengyel, Dhruv Dhody, Bo Wu, Qiufang Ma, Haomian Zheng, YuanYao, Wei Wang, Peng Liu, Zongpeng Du, Zhengqiang Li, Andrew Liu, Joe Clark, Roland Scott, Alex Huang Feng, Kai Gao, Jensen Zhang, Ziyang Xing, Mingshuang Jin, Aihua Guo, Zhidong Yin, Guoxiang Liu, Kaichun Wu for their valuable comments and great input to this work.

References

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/rfc/rfc2119>>.
- [RFC3688] Mealling, M., "The IETF XML Registry", BCP 81, RFC 3688, DOI 10.17487/RFC3688, January 2004, <<https://www.rfc-editor.org/rfc/rfc3688>>.
- [RFC6020] Bjorklund, M., Ed., "YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)", RFC 6020, DOI 10.17487/RFC6020, October 2010, <<https://www.rfc-editor.org/rfc/rfc6020>>.
- [RFC6241] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed., and A. Bierman, Ed., "Network Configuration Protocol (NETCONF)", RFC 6241, DOI 10.17487/RFC6241, June 2011, <<https://www.rfc-editor.org/rfc/rfc6241>>.
- [RFC6242] Wasserman, M., "Using the NETCONF Protocol over Secure Shell (SSH)", RFC 6242, DOI 10.17487/RFC6242, June 2011,

<<https://www.rfc-editor.org/rfc/rfc6242>>.

- [RFC6991] Schoenwaelder, J., Ed., "Common YANG Data Types", RFC 6991, DOI 10.17487/RFC6991, July 2013, <<https://www.rfc-editor.org/rfc/rfc6991>>.
- [RFC8040] Bierman, A., Bjorklund, M., and K. Watsen, "RESTCONF Protocol", RFC 8040, DOI 10.17487/RFC8040, January 2017, <<https://www.rfc-editor.org/rfc/rfc8040>>.
- [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/rfc/rfc8174>>.
- [RFC8341] Bierman, A. and M. Bjorklund, "Network Configuration Access Control Model", STD 91, RFC 8341, DOI 10.17487/RFC8341, March 2018, <<https://www.rfc-editor.org/rfc/rfc8341>>.
- [RFC8345] Clemm, A., Medved, J., Varga, R., Bahadur, N., Ananthakrishnan, H., and X. Liu, "A YANG Data Model for Network Topologies", RFC 8345, DOI 10.17487/RFC8345, March 2018, <<https://www.rfc-editor.org/rfc/rfc8345>>.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/rfc/rfc8446>>.
- [RFC8632] Vallin, S. and M. Bjorklund, "A YANG Data Model for Alarm Management", RFC 8632, DOI 10.17487/RFC8632, September 2019, <<https://www.rfc-editor.org/rfc/rfc8632>>.
- [RFC8969] Wu, Q., Ed., Boucadair, M., Ed., Lopez, D., Xie, C., and L. Geng, "A Framework for Automating Service and Network Management with YANG", RFC 8969, DOI 10.17487/RFC8969, January 2021, <<https://www.rfc-editor.org/rfc/rfc8969>>.

Informative References

- [BERT] "BERT (language model)", n.d., <[https://en.wikipedia.org/wiki/BERT_\(language_model\)](https://en.wikipedia.org/wiki/BERT_(language_model))>.
- [I-D.ietf-ippm-pam] Mirsky, G., Halpern, J. M., Min, X., Clemm, A., Strassner, J., and J. Francois, "Precision Availability Metrics for Services Governed by Service Level Objectives (SLOs)", Work in Progress, Internet-Draft, draft-ietf-ippm-pam-09, 1 December 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-ippm-pam-09>>.
- [I-D.ietf-netconf-configuration-tracing] Quilbeuf, J., Claise, B., Graf, T., Lopez, D., and S. Qiong, "External Trace ID for Configuration Tracing", Work in Progress, Internet-Draft, draft-ietf-netconf-configuration-tracing-05, 18 March 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-netconf-configuration-tracing-05>>.
- [I-D.ietf-netconf-trace-ctx-extension] Gagliano, R., Larsson, K., and J. Lindblad, "NETCONF Extension to support Trace Context propagation", Work in Progress, Internet-Draft, draft-ietf-netconf-trace-ctx-extension-04, 3 March 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-netconf-trace-ctx-extension-04>>.

[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-19, 18 June 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-nmop-terminology-19>>.

[I-D.irtf-nmrg-ai-challenges]

Franois, J., Clemm, A., Papadimitriou, D., Fernandes, S., and S. Schneider, "Research Challenges in Coupling Artificial Intelligence and Network Management", Work in Progress, Internet-Draft, draft-irtf-nmrg-ai-challenges-05, 18 March 2025, <<https://datatracker.ietf.org/doc/html/draft-irtf-nmrg-ai-challenges-05>>.

[ITU-T-G-7710]

"ITU-T G.7710/Y.1701 - Common equipment management function requirements", 2020, <<https://www.itu.int/rec/T-REC-G.7710>>.

[ITU-T-X-733]

"ITU-T X.733 - Information technology - Open Systems Interconnection - Systems Management - Alarm reporting function", 1999, <<https://www.itu.int/rec/T-REC-X.733/fr>>.

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

[RFC9375] Wu, B., Ed., Wu, Q., Ed., Boucadair, M., Ed., Gonzalez de Dios, O., and B. Wen, "A YANG Data Model for Network and VPN Service Performance Monitoring", RFC 9375, DOI 10.17487/RFC9375, April 2023, <<https://www.rfc-editor.org/rfc/rfc9375>>.

[RFC9408] Boucadair, M., Ed., Gonzalez de Dios, O., Barguil, S., Wu, Q., and V. Lopez, "A YANG Network Data Model for Service Attachment Points (SAPs)", RFC 9408, DOI 10.17487/RFC9408, June 2023, <<https://www.rfc-editor.org/rfc/rfc9408>>.

[RFC9417] Claise, B., Quilbeuf, J., Lopez, D., Voyer, D., and T. Arumugam, "Service Assurance for Intent-Based Networking Architecture", RFC 9417, DOI 10.17487/RFC9417, July 2023, <<https://www.rfc-editor.org/rfc/rfc9417>>.

[RFC9418] Claise, B., Quilbeuf, J., Lucente, P., Fasano, P., and T. Arumugam, "A YANG Data Model for Service Assurance", RFC 9418, DOI 10.17487/RFC9418, July 2023, <<https://www.rfc-editor.org/rfc/rfc9418>>.

[RFC9543] Farrel, A., Ed., Drake, J., Ed., Rokui, R., Homma, S., Makhijani, K., Contreras, L., and J. Tantsura, "A Framework for Network Slices in Networks Built from IETF Technologies", RFC 9543, DOI 10.17487/RFC9543, March 2024, <<https://www.rfc-editor.org/rfc/rfc9543>>.

[TMF724A] "Incident Management API Profile v1.0.0", 2023, <<https://www.tmforum.org/resources/standard/tmf724a-incident-management-api-profile-v1-0-0/>>.

[W3C-Trace-Context]

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

A.1 Network Incident management with specific network topology and the network service

In this example, we show a network incident that are associated with the node 'D1', the network topology 'L2-Topo' and the domain 'FAN'. The probable root cause is also analysed.

```
{
  "incident-no": 56433218,
  "incident-id": "line fault",
  "service-instance": ["optical-svc-A"],
  "domain": "FAN",
  "priority": "critical",
  "occur-time": "2026-03-10T04:01:12Z",
  "clear-time": "2026-03-10T06:01:12Z",
  "ack-time": "2026-03-10T05:01:12Z",
  "last-updated": "2026-03-10T05:31:12Z",
  "status": "unacknowledged-and-uncleared",
  "category": "Line",
  "source": [
    {
      "node-ref": "example:D1",
      "network-ref": "example:L2-topo",
      "resource": [
        {
          "name": "7985e01a-5aad-11ea-b214-286ed488cf99"
        }
      ]
    }
  ],
  "probable-causes": [
    {
      "name": "Feeder fiber great loss change",
      "detail-information": "The connector of the optical fiber is contaminated, Or the optical fiber is bent too much.",
      "probable-cause": {
        "network-ref": "example:L2-topo",
        "node-ref": "example:D1",
        "resource": [
          {
            "name": "7985e01a-5aad-11ea-b214-286ed488cf99",
            "cause-name": "ltp",
            "detail": "Frame=0, Slot=6, Subslot=65535, Port=7, ODF= ODF001, Level1Splitter= splitter0025"
          }
        ]
      }
    }
  ],
  "probable-event": [
    {
      "event-id": "8921834",
      "type": "alarm"
    }
  ],
  "events": [
    {
      "even-id": "8921832",
      "type": "alarm"
    },
    {
      "even-id": "8921833",
      "type": "alarm"
    }
  ]
}
```

```

    "even-id": "8921834",
    "type": "alarm"
  }
]
}

```

A.2 Correlation with trouble tickets

In this document, the objective of the incident management is to identify probable causes and reduce duplicated ticket amounts.

Traditionally, troubleshooting ticket is created upon critical alert is received, e.g., due to excessive BGP flaps on a particular device by the OSS system. Such troubleshooting ticket will trigger network incident management in the network controller. Therefore normally trouble shooting tickets and network incident are managed by the OSS and the network controller respectively. However Network troubleshooting is sometimes complicated and requires data gathering and analysis from many different tools from the controllers, therefore correlation between troubleshooting ticket and network incident becomes necessary.

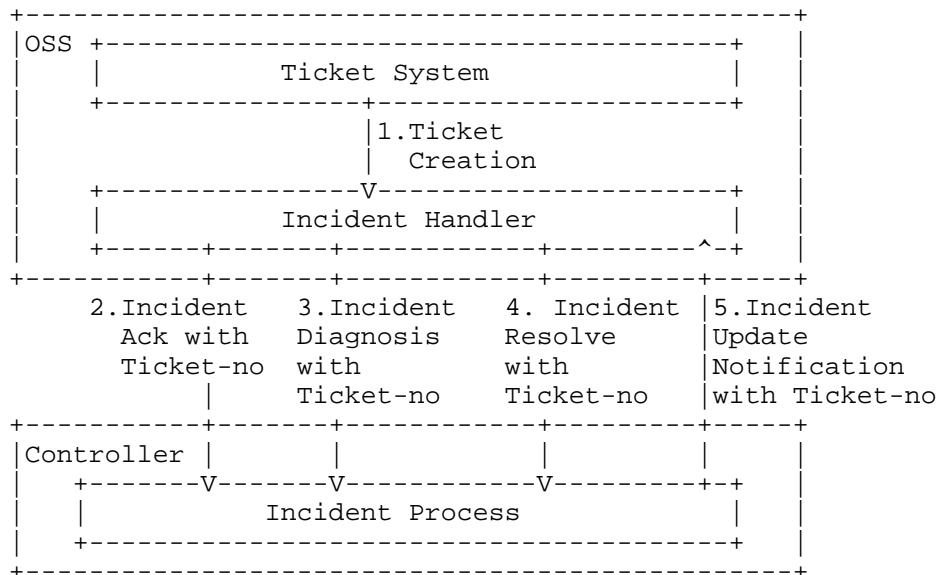


Figure 7: Correlation with troubleshooting tickets

In order to manage the correlation between network incidents and trouble tickets in the YANG data model, three rpcs to manage the network incidents and one notification to report on network incident state changes defined in "ietf-incident" module can be further extended to include "ticket-no" attribute so that such correlation can be carried in the incident update notification and report the upper layer OSS system. Such correlation can be used by the incident handler in the upper layer OSS system for further fault demarcation, e.g., identify whether the fault is on the user side or on the network side.

```

rpcs:
+---x incident-acknowledge
| +---w input
| | +---w incident-no* incident-ref
| | +---w ticket-no? string
+---x incident-diagnose
| +---w input
| | +---w incident-no* incident-ref
| | +---w ticket-no? string
+---ro output

```

```

| |   +---ro task-id? string
+---x incident-resolve
|   +---w input
|       +---w incident-no* incident-ref
|       +---w ticket-no? string

```

notifications:

```

+---n incident-notification
|   +---ro incident-no? incident-ref
|   +---ro ticket-no? string

```

```

+--
...

```

A.3 Intent based Networking with incident diagnosis task list

In this document, incident-diagnosis RPC defined in in "ietf-incident" module can be used to identify probable root causes; and an incident update notification can be triggered to report the diagnosis status if successful.

In some case, workflows may span a long duration or involve multiple step task. In such case, intent based networking can be used to support such multiple step task and provide more detailed network diagnosis information, the "incident-diagnosis" RPC can be further extended to support "task-id" attribute and other auxiliary attributes. "ietf-incident" module can be further extended to include "incident-diagnosis-task" list with the following diagnosis information:

The current status (e.g., created, diagnosing, diagnosed, finished) of each diagnosis task.

Task start time, end time, diagnosis result (succeeded, failed), failure description, etc.

probable root causes, probable events, repair recommendations, etc.

```

+--ro incident-diagnosis-tasks
|   +--ro incident-diagnosis-task* [task-id]
|   +--ro task-id? String
|   +--ro incident-no* incident-ref
|   +--ro ticket-no? string
|   +--ro start-time? yang:date-and-time
|   +--ro end-time? yang:date-and-time
|   +--ro task-state? enumeration
|   +--ro diagnosis-result? enumeration
|   +--ro diagnosis-result-description? String
|   +--ro probable-causes leafref //List <RootCause>
...
|   +--ro probable-events leafref //List <Event>
...
|   +-- ro repair-advice
|   +-- ro state enumeration // Incident states such as Creation, Update, Clear
...

```

Take multi-domain fault demarcation as an example, when both base station incident in the RAN network and Network Link incident in the IP network are received and base station incident from user side results from network incident in other domains, the OSS system is unable to find network side problem simply based on base station incident. Therefore incident diagnosis RPC will be invoked with IP address of Base station and incident start time as input and sent to the network controller. The network controller can use network diagnosis related intent based interface to find the corresponding network side port according to the base station IP address, and then further associated with transmission path (current path, historical

path) to the base station and current and historical network performance, network resources, and incident status data, to diagnose the probable root cause of the network incident and provide repair suggestions.

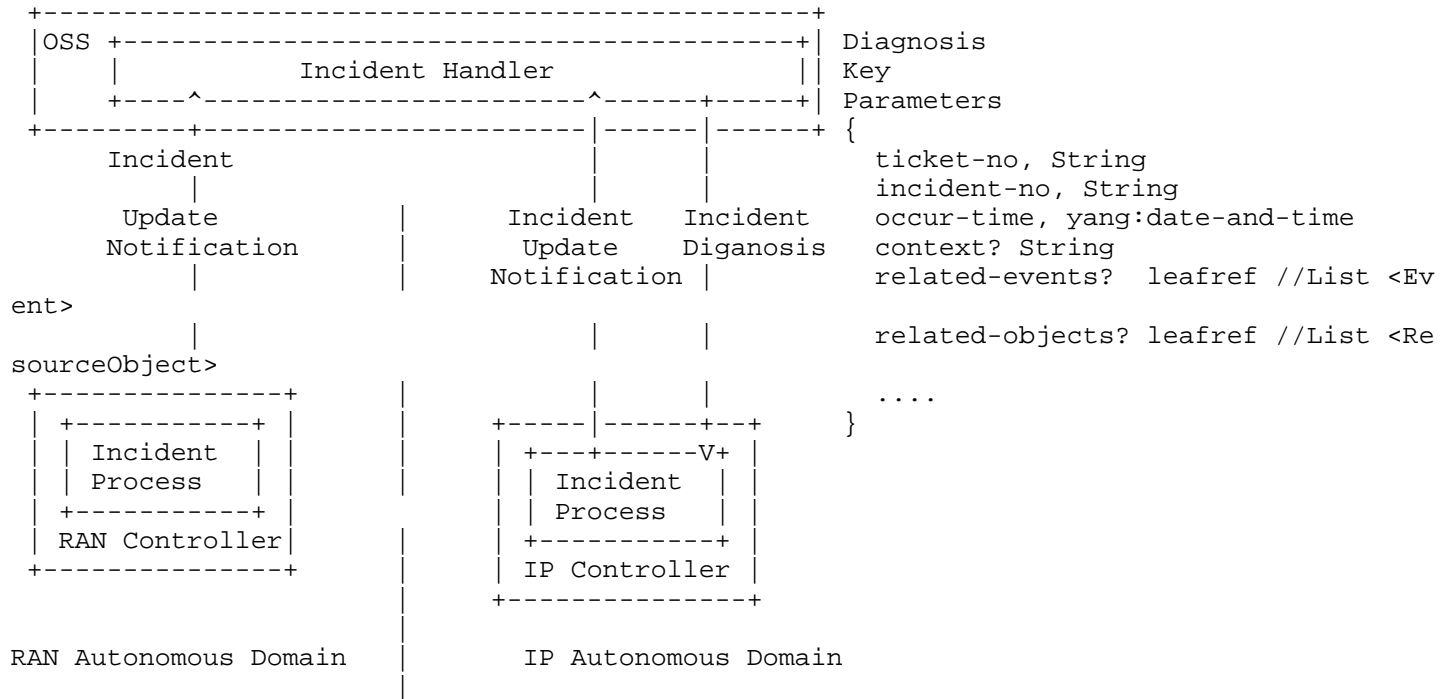


Figure 8: Multi-Domain Fault Demarcation

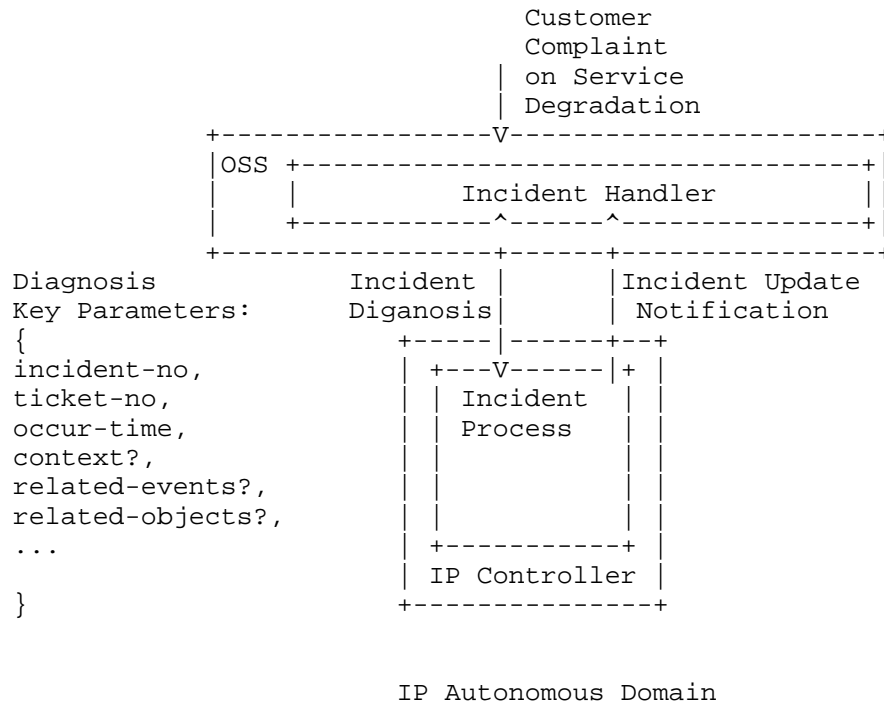


Figure 9: Service Complaint triggered Network Diagnosis

Similarly, in case of service degradation for a lease line service receiving from the customer, the OSS system can request network diagnosis at the network side conducted by the network controller. The network controller can use network diagnosis related intent based interface to find the corresponding network side port based on the dedicated line service, and then further associate the transmission path (current path, historical path) and current and historical

network performance, network resources, and incident status data to diagnose the probable root cause of the fault and provide repair suggestions.

Appendix B. Changes between Revisions

v04 - v05

- * Replace probable cause with probable root cause based on Adrian and Benoit's suggestion.
- * Address editorial comments raised by Aitken Paul.
- * YANG Model editorial changes based on Aitken Paul's comments.

v03 - v04

- * Remove constraint of using machine learning for service impact analysis and replace machine learning with algorithmic techniques.
- * Replace root cause with probable cause based on IETF 122 NMOP Session Discussion.
- * Add two ITU-T references for probable cause definition in the terminologies section.
- * Add Lionel Tailhardat from Orange as new contributors based on his input.
- * Add two new examples in the Appendix to explore correlation between troubleshooting ticket and incident management and intent based network diagnosis interaction.

v02 - v03

- * Cross-checking terminology across NMOP drafts based on Adrian's comments.
- * Align with the Terminology draft based on Thomas's comments.
- * Clarify the relation between the Network Incident, and Customer Incident.
- * Add service impact analysis term and its definition.
- * Clarify the relation between fault, problem, incident, service.
- * Other Editorial changes.

v01 - v02

- * Clarify the relation between fault, incident and problem.
- * Clarify the relation between fault management and incident management.
- * Add clarification text to make draft focus on network level incident management, not be tied with OSS or under the control of OSS.
- * Other Editorial changes.

v00 - v01

- * Clarify the relationship between incident-no and incident-id.

- * Fix Tree Diagram to align with YANG module code change.
- * Add json example in the appendix.
- * Add failure handling process for rpc error.
- * Clarify the relationship between events and cause.
- * Clarify synchronous nature of these RPCs.
- * Clarify the relationship between inter-layer and inter-domain.
- * Refer to terminology draft for terminology alignment.
- * Fix pyang compilation issue and yang lint issue.
- * Fix Broken ref by using node-ref defined in RFC8345.
- * Update YANG data model based on issues raised in issue tracker of the github.
- * Shorten the list of authors to 5 based on chairs' comment and move additional authors to top 3 contributors.
- * Merge ietf-incident-type.yang into ietf-incident.yang
- * Fix enumeration on leaf type
- * Clarify the scope in the abstract and introduction and make the scope focus on YANG data model
- * Provide text around figure 5 to clarify how the incident server know the real effect on the relevant services.
- * Other editorial changes.

v00 (draft-ietf-nmop-network-incident-yang)

- * Change draft name from draft-feng-opsawg-incident-management into draft-feng-nmop-netwrok-incident-yang
- * Change title into A YANG Data Model for Network Incident Management
- * open issues is tracked in <https://github.com/billwuqin/network-incident/issues>

v03 - v04 (draft-feng-opsawg-incident-management)

- * Update incident defintion based on TMF incident API profile specification.
- * Update use case on Multi-layer Fault Demarcation based on side meeting discussion and IETF 119 session discussion.
- * Update section 5.1 to explain how network incident is generated based on other factors.
- * Add one new use cases on Security Events noise reduction based on Situation Awareness.
- * Other Editorial changes.

v02 - v03 (draft-feng-opsawg-incident-management)

- * Add one new use cases on Incident Generation.

- * Add reference to Precision Availability Metric defined in IPPM PAM WG document.

v01 - v02

- * A few Editorial change to YANG data models in section 8.
- * Add some text to the model design overview.
- * Revise sample use cases section to focus on two key use cases.
- * Motivation and goal clarification in the introduction section.

v00 - v01 (draft-feng-opsawg-incident-management)

- * Modify the introduction.
- * Rename incident agent to incident server.
- * Add the interworking with alarm management.
- * Add the interworking with SAIN.
- * Add the relationship with RFC8969.
- * Add the relationship with observation timestamp and trace context.
- * Clarify the incident identification process.
- * Modify the work flow of incident diagnosis and resolution.
- * Remove identities and typedefs from ietf-incident YANG module, and create a new YANG module called ietf-incident-types.
- * Modify ietf-incident YANG module, for example, modify incident-diagnose rpc and incident-resolve rpc.

Contributors

Lionel Tailhardat
Orange
Email: lionel.tailhardat@orange.com

Thomas Graf
Swisscom
Switzerland
Email: thomas.graf@swisscom.com

Zhenqiang Li
CMCC
Email: li_zhenqiang@hotmail.com

Yanlei Zheng
China Unicom
Email: zhengyanlei@chinaunicom.cn

Yunbin Xu
CAICT
Email: xuyunbin@caict.ac.cn

Xing Zhao
CAICT
Email: zhaoxing@caict.ac.cn

Chaode Yu
Huawei
Email: yuchaode@huawei.com

Authors' Addresses

Tong Hu
CMCC
Building A01, 1600 Yuhangtang Road, Wuchang Street, Yuhang District
Hangzhou
311121
China
Email: hutong@cmhi.chinamobile.com

Luis M. Contreras
Telefonica
Madrid
Spain
Email: luismiguel.contrerasmurillo@telefonica.com

Qin Wu
Huawei
101 Software Avenue, Yuhua District
Nanjing
210012
China
Email: bill.wu@huawei.com

Nigel Davis
Ciena
Email: ndavis@ciena.com

Chong Feng
Email: fengchonglilly@gmail.com