

Internet Engineering Task Force (IETF)
Request for Comments: 9940
Category: Informational
ISSN: 2070-1721

N. Davis, Ed.
Ciena
A. Farrel, Ed.
Old Dog Consulting
T. Graf
Swisscom
Q. Wu
C. Yu
Huawei
April 2026

Some Key Terms for Network Fault and Problem Management

Abstract

This document sets out some terms that are fundamental to a common understanding of network fault and problem management within the IETF.

The purpose of this document is to bring clarity to discussions and other work related to network fault and problem management -- in particular, to YANG data models and management protocols that report, make visible, or manage network faults and problems.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are candidates for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9940>.

Copyright Notice

Copyright (c) 2026 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. Usage of Terms
3. Terminology
 - 3.1. Context Terminology

3.2.	Core Terms
3.3.	Other Terms
4.	Workflow Explanations
5.	Security Considerations
6.	Privacy Considerations
7.	IANA Considerations
8.	Informative References
	Acknowledgments
	Authors' Addresses

1. Introduction

Successful operation of large networks depends on effective network management. This requires a virtuous circle of network control, network observability, network analytics, network assurance, and back to network control. Network fault and problem management [RFC6632] is an important aspect of network management and control solutions. It deals with the detection, reporting, inspection, isolation, correlation, and management of events within the network. The intention of this document is to focus on those events that have a negative effect on the network's ability to forward traffic according to expected behaviors that may reduce the network's ability to deliver services. Such events may also impact the ability to control and operate the network. The document also considers other faults that reduce the quality or reliability of the delivered service. The concept of fault and problem management extends to include actions taken to determine the causes of problems and to work toward recovery of expected network behavior.

A number of work efforts within the IETF seek to provide components of a fault management system, such as YANG data models or management protocols. It is important that a common terminology be used so that there is a clear understanding of how the elements of the management and control solutions fit together and how faults and problems will be handled.

This document sets out some terms that are fundamental to a common understanding of network fault and problem management. While "faults" and "problems" are concepts that apply at all levels of technology in the Internet, the scope of this document is restricted to the network layer and below; hence, this document is specifically about "network fault and problem management." The concept of "incidents" is also touched on in this document, where an incident results from one or more problems and is the disruption of a network service.

Note that some useful terms are defined in [RFC3877] and [RFC8632]. The definitions in this document are informed by those documents, but they are not dependent on that prior work.

2. Usage of Terms

The terms defined in this document are intended for consistent use within the IETF in the scope of network fault and problem management. Where similar concepts are described in other bodies, an attempt has been made to harmonize with those other descriptions, but care is needed where terms are not used consistently between bodies or where terms are applied outside the network layer. If other bodies find the terminology defined in this document useful, they are free to use it.

The purpose of this document is to define the following terms for use in other documents. Other terms are defined to enable those definitions and may also be used by other documents, although that is not the principal purpose of their definitions here.

- * Event
- * State
- * Fault
- * Problem
- * Symptom
- * Cause
- * Alert
- * Alarm

When other documents make use of the terms as defined in this document, it is suggested here that such uses should use capitalization of the terms as in this document to help distinguish them from colloquial uses and should include an early section listing the terms inherited from this document with a citation.

3. Terminology

This section contains key terms. It is split into three subsections.

- * Section 3.1 contains terms that help set the context for network fault and problem management systems.
- * Section 3.2 includes specific and detailed core terms that will be used in other documents that describe elements of the network fault and problem management systems.
- * Section 3.3 provides three further terms that may be helpful.

3.1. Context Terminology

This section includes some terminology that helps describe the context for the rest of this work. The terms may be viewed as a cascaded sequence of processes, starting with Network Telemetry and building to Network Observability. The definitions are deliberately kept relatively terse. Further documents may expand on these terms without loss of specificity. Such contextualization (if any) should be highlighted clearly in those documents.

Network Telemetry: This is defined in [RFC9232] and describes the process of collecting operational network data categorized according to the network plane (e.g., Layer 3, Layer 2, and Layer 1) from which it was derived. Data collected through the Network Telemetry process does not contain any data related to service definitions (i.e., "intent" per Section 3.1 of [RFC9315]).

Network Monitoring: This is the process of keeping a continuous record of functions related to a network topology. It involves tracking various aspects such as traffic patterns, device health, performance metrics, and overall network behavior. This approach differentiates Network Monitoring from resource or device monitoring, which focuses on individual resources or components (Section 3.2).

Network Analytics: This is the process of deriving analytical insights from operational network data. A process could be executed by a piece of software, a system, or a human that analyzes operational data and outputs new analytical data related to the operational data -- for example, a symptom.

Network Observability: This is the process of enabling network behavioral assessment through analysis of observed operational network data (logs, alarms, traces, etc.) with the aim of detecting symptoms of network behavior, and identifying anomalies and their causes. Network Observability begins with information gathered using Network Monitoring tools. That information may be further enriched with other operational data. The expected

outcome of the observability processes is identification and analysis of deviations in observed state versus the expected state of a network.

Thus, there is a cascaded sequence where the following relationships apply:

- * Network Telemetry is the process of collecting operational data from a network.
- * Network Monitoring is the process of creating/keeping a record of data gathered in Network Telemetry.
- * Network Analytics is the process of deriving insight through the data recorded in Network Monitoring.
- * Network Observability is the process of enabling behavioral assessment of a network through Network Analytics.

3.2. Core Terms

The terms in this section are presented in an order that is intended to flow such that it is possible to gain understanding reading top to bottom. The figures and explanations in Section 4 may aid understanding the terms set out here.

Resource: A Resource is an element of a network system.

- * Resource is a recursive concept so that a Resource may be a collection of other Resources (for example, a network node comprises a collection of network interfaces).

Characteristic: A Characteristic is an observable or measurable aspect or behavior associated with a Resource.

- * A Characteristic may be considered to be built on facts (see "Value", below) and the contexts and descriptors that identify and give meaning to the facts.
- * The term "Metric" (see "metric" in [RFC9417]) is another word for a measurable Characteristic, which may also be thought of as analogous to a "variable".

Value: A Value is a measure of a Characteristic associated with a Resource. It may be in the form of a categorization (e.g., high or low), an integer (e.g., a count or gauge), or a reading of a continuous variable (e.g., an analog measurement), etc.

Change: In the context of Network Monitoring, a Change is the variation in the Value of a Characteristic associated with a Resource. A Change may arise over a period of time.

- * Not all Changes are noteworthy (i.e., they do not have Relevance).
- * Perception of Change depends upon Detection, the sampling rate/accuracy/detail, and perspective.
- * It may be helpful to qualify this as "Value Change" because the English word "change" is often heavily used.

Event: An Event is the variation in Value of a Characteristic of a Resource at a distinct moment in time (i.e., the period is negligible).

- * Compared with a Change, which may be over a period of time, an

Event happens at a distinct moment in time. Thus, an Event may be the observation of a Change.

Condition: A Condition is an interpretation of the Values of a set of one or more Characteristics of a Resource (with respect to working order or some other aspect relevant to the Resource purpose/application) -- for example, "low available memory". Thus, it is the output of a function applied to a set of one or more variables.

State: A State is a particular Condition that a Resource has (i.e., it is in a State) at a specific time. For example, a router may report the total amount of memory it has and how much is free. These are the Values of two Characteristics of a Resource. These Values can be interpreted to determine the Condition of the Resource, and that may determine the State of the router, such as shortage of memory.

- * While a State may be observed at a specific moment in time, it is actually determined by summarizing measurement over time in a process sometimes called State compression.
- * It may be helpful to qualify this as "Resource State" to make clear the distinction between this and other uses of "state" such as "protocol state".
- * This term may be contrasted with "operational state" as used in [RFC8342]. For example, the state of a link might be up/down/degraded, but the operational state of the link would include a collection of Values of Characteristics of the link.

Detect (hence Detected, Detection): To Detect is to notice the presence of something (State, Change, Event, activity, etc.)

- * Also to notice a Change (from the perspective of an observer such as a monitoring system).

Relevance: Relevance is the consideration of an Event, State, or Value (through the application of policy, relative to a specific perspective or intent, and in relation to other Events, States, and Values) to determine whether it is of note to the system that controls or manages the network. Note, for example, that not all Changes are Relevant.

- * This term may also be used as "Relevant Event", "Relevant State", or "Relevant Value".

Occurrence: An Occurrence is a Relevant Event or a particular Relevant Change.

- * An Occurrence may be an aggregation or abstraction of multiple fine-grained Events or Changes.
- * An Occurrence may occur at any macro or micro scale because Resources are a recursive concept. An Occurrence may be perceived, depending on the scope of observation (i.e., according to the level of Resource recursion that is examined). That is, Occurrences, themselves, are a recursive concept.

Fault: A Fault is an Occurrence (i.e., an Event or a Change) that is not desired/required (as it may be indicative of a current or future undesired State). Thus, a Fault happens at a moment in time. A Fault can potentially be associated with a Cause. See [RFC8632] for a more detailed discussion of network faults.

- * Note that there is a distinction between a Fault and a Problem

that depends on context. For example, in a connectivity service where redundancy is present, a link down is a Problem, but from the perspective of managing the network resources, a link down is a Fault. Likewise, for example, in a router with two power supplies, if the backup power supply fails leaving the primary unprotected, this is a Problem.

Problem: A Problem is a State that is undesirable and that may require remedial action. A Problem cannot necessarily be associated with a Cause. The resolution of a Problem does not necessarily act on the thing that has the Problem.

- * Note that there is a historic aspect to the concept of a Problem. The current State may be operational, but there could have been a Fault that is unexplained, and the fact of that unexplained recent Fault is a Problem.
- * Note that while a Problem is unresolved it may continue to require attention. A record of resolved Problems may be maintained in a log.
- * Note that there may be a State that is considered to be a Problem from several perspectives. For example, consider a "loss of light" State that may cause multiple services to fail. In this example, a new State (the light recovers) may cause the Problem to be resolved from one perspective (the services are operational once more) but may leave the Problem as unresolved (because the loss of light has not been explained). Further, in this example, there could be another development (the reason for the temporary loss of light is traced to a microbend in the fiber that is repaired) resulting in that unresolved Problem now being resolved. But, in this example, this still leaves a further Problem unresolved (a microbend occurred, and that Problem is not resolved until it is understood how it occurred and a remedy is put in place to prevent recurrence).

Cause: A Cause is the Events (Detected or otherwise) that gave rise to a Fault/Problem.

Incident: Also referred to as "Network Incident". An Incident is an undesired Occurrence such as an unexpected interruption of a network service, degradation of the quality of a network service, or the below-target performance of a network service. An Incident results from one or more Problems, and a Problem may give rise to or contribute to one or more Incidents. Greater discussion of Network Incident relationships, including Customer Incidents and Incident management, can be found in [Net-Incident-Mgmt-YANG].

Symptom: A Symptom is an observable Value, Change, State, Event, or Condition considered as an indication of a Problem or potential Problem.

Anomaly: Also referred to as "Network Anomaly". An Anomaly is an unusual or unexpected Event or pattern in network data in the forwarding plane, control plane, or management plane that deviates from the normal, expected behavior. See [Net-Anomaly-Arch] for more details.

Alert: An Alert is an indication of a Fault.

Alarm: As specified in [RFC8632], an Alarm signifies an undesirable State in a Resource that requires corrective action. From a management point of view, an Alarm can be seen as a State in its own right and the transition to this State may result in an Alert being issued. The receipt of this Alert may give rise to a continuous indication (to a human operator) highlighting the

potential or actual presence of a Problem.

3.3. Other Terms

Three other terms may be helpful:

Intermittent: A State that is not continuous but that keeps recurring in some time frame.

Transient: A State that is not continuous and that occurs once in some time frame.

Recurrent: A Problem that is actively resolved but returns.

4. Workflow Explanations

This section aims to add information about the relationship between the terms defined in Section 3.2 in the context of network fault and problem management. The text and figures here are for explanation and are not normative for the definition of terms.

The relationship between Resources and Characteristics is shown in Figure 1. Note that there is a 1:n relationship between a Network system and Resources and between Resources and Characteristics: For clarity, this is not shown in the figure.

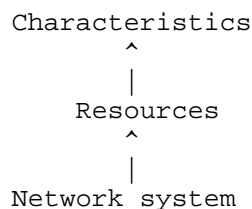


Figure 1: Resources and Characteristics

The Value of a Characteristic of a Resource may change over time. Specific Changes in Value may be noticed at a specific time (as digital Changes), Detected, and treated as Events. This is shown on the left-hand side of Figure 2.

The center of Figure 2 shows how the Value of a Characteristic may change over time. The Value may be Detected at specific times or periodically and give rise to Conditions that are States (and consequently State Changes).

In practice, the Characteristic may vary in an analog manner over time as shown on the right-hand side of Figure 2. The Value can be read or reported (i.e., Detected) periodically leading to analog Values that may be deemed Relevant Values, or it may be evaluated over time as shown in Figure 6.

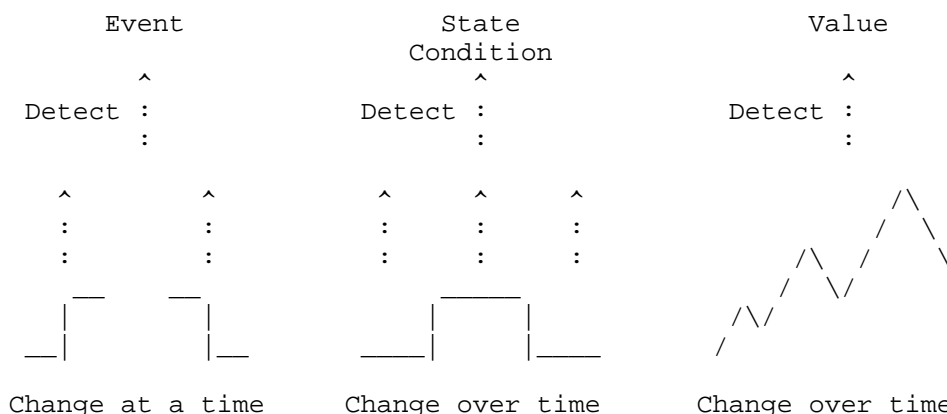


Figure 2: Characteristics and Changes

Figure 3 shows the workflow progress for Events. As noted above, an Event is a Change in the Value of a Characteristic at a time. The Event may be evaluated (considering policy, relative to a specific perspective, with a view to intent, and in relation to other Events, States, and Values) to determine if it is an Occurrence and possibly to indicate a Change of State. An Occurrence may be undesirable (a Fault), which might cause an Alert to be generated. Or, an Occurrence may be evidence of a Problem and could directly indicate a Cause. In some cases, an Alert may give rise to an Alarm highlighting the potential or actual presence of a Problem.

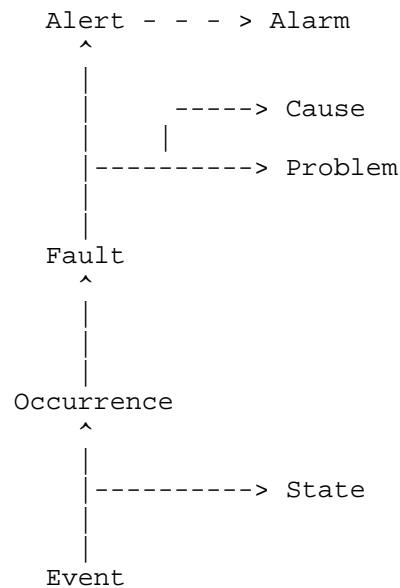


Figure 3: Event and Dependent Terms

Parallel to the workflow for Events, Figure 4 shows the workflow progress for States. As shown in Figure 2, Change noted at a particular time gives rise to State. The State may be deemed to have Relevance considering policy, relative to a specific perspective, with a view to intent, and in relation to other Events, States, and Values. A Relevant State may be deemed a Problem, or it may indicate a Problem or potential Problem.

Problems may be considered based on Symptoms and may map directly or indirectly to Causes. An Incident results from one or more Problems. An Alarm may be raised as the result of a Problem, and the transition to an alarmed State may give rise to an Alert.

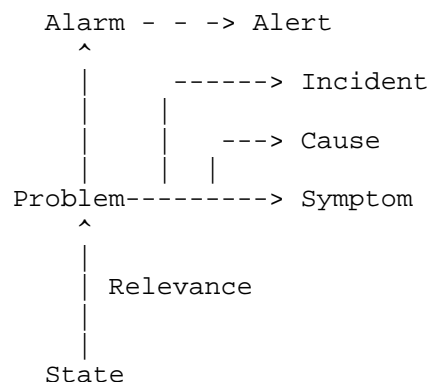


Figure 4: State and Dependent Terms

Figure 5 shows how Faults and Problems may be consolidated to determine the Causes. The arrows show how one item may give rise to another.

A Cause can be indicated by or determined from Faults, Problems, and Symptoms. It may be that one Cause points to another, and it can also be considered as a Symptom. The determination of Causes can consider multiple inputs. An Incident results from one or more Problems.

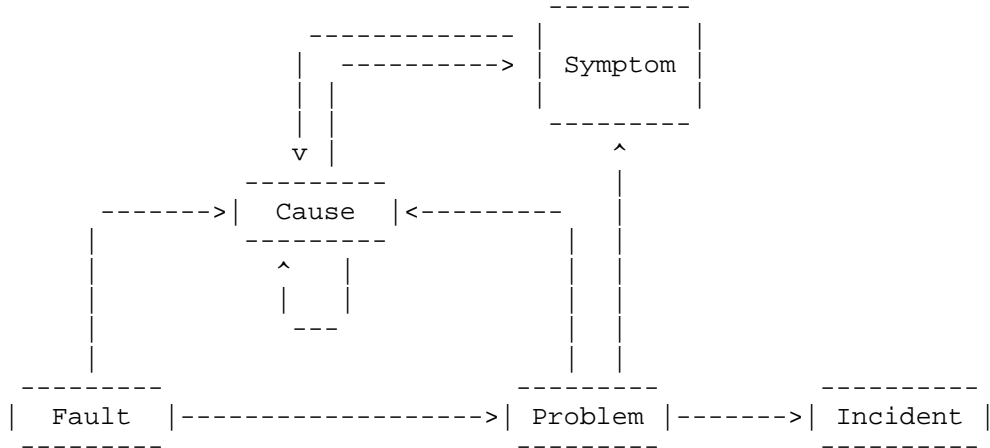


Figure 5: Consolidation of Symptoms and Causes

Figure 6 shows how thresholds are important in the consideration of analog Values and Events. The arrows in the figure show how one item may give rise to or utilize another. The use of threshold-driven Events and States (and the Alerts that they might give rise to) must be treated with caution to dampen any "flapping" (so that consistent States may be observed) and to avoid overwhelming management processes or systems. Analog Values may be read or notified from the Resource and could transition a threshold, be deemed Relevant Values, or be evaluated over time. Events may be counted, and the Count may cross a threshold or reach a Relevant Value.

The Threshold Process may be implementation specific and subject to policies. When a threshold is crossed and any other conditions are matched, an Event may be determined and treated like any other Event.

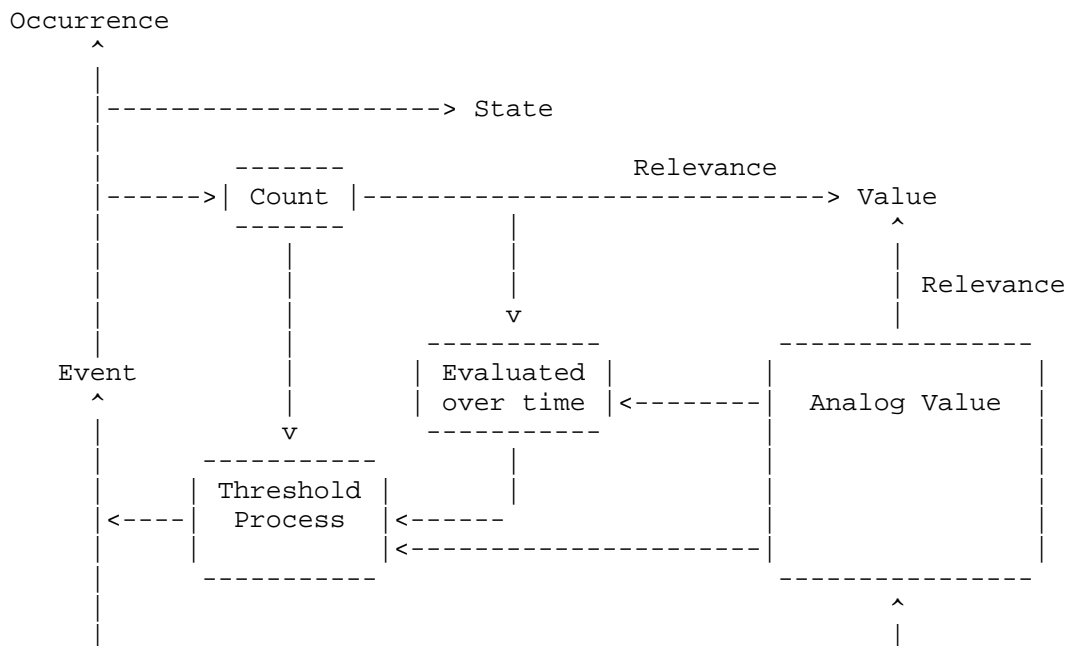




Figure 6: Counts, Thresholds, and Values

5. Security Considerations

This document specifies terminology and has no direct effect on the security of implementations or deployments. However, protocol solutions and management models need to be aware of several aspects:

- * The exposure of information pertaining to Faults and Problems may make available knowledge of the internal workings of a network (in particular, its vulnerabilities) that may be of use to an attacker.
- * Systems that generate management information (messages, notifications, etc.) when Faults occur may be attacked by causing them to generate so much information that the system that manages the network is swamped and unable to properly manage the network.
- * Reporting false information about Faults (or masking reports of Faults) may cause the system that manages the network to function incorrectly.

6. Privacy Considerations

Network fault and problem management should preserve user privacy by not exposing user data or information about end-user activities.

Network Telemetry involves observing network traffic and collecting operational data from the network, while Network Monitoring is the process of keeping records of data gathered in Network Telemetry. Therefore, it is possible that the data observed and collected includes users' privacy information. Such information must be protected and controlled to avoid exposure to unauthorized parties. Particular care may need to be exercised over stores of such information that might be accessed at any time (including far into the future).

Additionally, a network operator will be concerned about keeping control of all information about Faults to protect their own privacy and the details of how they operate their network.

7. IANA Considerations

This document has no IANA actions.

8. Informative References

[Net-Anomaly-Arch]

Graf, T., Du, W., Francois, P., and A. Huang Feng, "A Framework for a Network Anomaly Detection Architecture", Work in Progress, Internet-Draft, draft-ietf-nmop-network-anomaly-architecture-06, 21 November 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-nmop-network-anomaly-architecture-06>>.

[Net-Incident-Mgmt-YANG]

Hu, T., Contreras, L. M., Wu, Q., Davis, N., and C. Feng, "A YANG Data Model for Network Incident Management", Work in Progress, Internet-Draft, draft-ietf-nmop-network-incident-yang-08, 13 February 2026, <<https://datatracker.ietf.org/doc/html/draft-ietf-nmop-network-incident-yang-08>>.

- [RFC3877] Chisholm, S. and D. Romascanu, "Alarm Management Information Base (MIB)", RFC 3877, DOI 10.17487/RFC3877, September 2004, <<https://www.rfc-editor.org/info/rfc3877>>.
- [RFC6632] Ersue, M., Ed. and B. Claise, "An Overview of the IETF Network Management Standards", RFC 6632, DOI 10.17487/RFC6632, June 2012, <<https://www.rfc-editor.org/info/rfc6632>>.
- [RFC8342] Bjorklund, M., Schoenwaelder, J., Shafer, P., Watsen, K., and R. Wilton, "Network Management Datastore Architecture (NMDA)", RFC 8342, DOI 10.17487/RFC8342, March 2018, <<https://www.rfc-editor.org/info/rfc8342>>.
- [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/info/rfc8632>>.
- [RFC9232] Song, H., Qin, F., Martinez-Julia, P., Ciavaglia, L., and A. Wang, "Network Telemetry Framework", RFC 9232, DOI 10.17487/RFC9232, May 2022, <<https://www.rfc-editor.org/info/rfc9232>>.
- [RFC9315] Clemm, A., Ciavaglia, L., Granville, L. Z., and J. Tantsura, "Intent-Based Networking - Concepts and Definitions", RFC 9315, DOI 10.17487/RFC9315, October 2022, <<https://www.rfc-editor.org/info/rfc9315>>.
- [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/info/rfc9417>>.

Acknowledgments

The authors would like to thank Med Boucadair, Wanting Du, Joe Clarke, Javier Antich, Benoit Claise, Christopher Janz, Sherif Mostafa, Kristian Larsson, Dirk Von Hugo, Carsten Bormann, Hilarie Orman, Stewart Bryant, Bo Wu, Paul Kyzivat, Jouni Korhonen, Reshad Rahman, Rob Wilton, Mahesh Jethanandani, Tim Bray, Paul Aitken, and Deb Cooley for their helpful comments.

Special thanks to the team that met at a side meeting at IETF 120 to discuss some of the thorny issues:

- * Benoit Claise
- * Watson Ladd
- * Brad Peters
- * Bo Wu
- * Georgios Karagiannis
- * Olga Havel
- * Vincenzo Riccobene
- * Yi Lin
- * Jie Dong
- * Aihua Guo
- * Thomas Graf
- * Qin Wu
- * Chaode Yu
- * Adrian Farrel

Authors' Addresses

Nigel Davis (editor)
Ciena
United Kingdom

Email: ndavis@ciena.com

Adrian Farrel (editor)
Old Dog Consulting
United Kingdom
Email: adrian@olddog.co.uk

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

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

Chaode Yu
Huawei
Email: yuchaode@huawei.com