

OPSAWG
Internet-Draft
Intended status: Informational
Expires: August 23, 2021

H. Song
Futurewei
F. Qin
China Mobile
P. Martinez-Julia
NICT
L. Ciavaglia
Nokia
A. Wang
China Telecom
February 19, 2021

Network Telemetry Framework
draft-ietf-opsawg-ntf-07

Abstract

Network telemetry is a technology for gaining network insight and facilitating efficient and automated network management. It encompasses various techniques for remote data generation, collection, correlation, and consumption. This document describes an architectural framework for network telemetry, motivated by challenges that are encountered as part of the operation of networks and by the requirements that ensue. Network telemetry, as necessitated by best industry practices, covers technologies and protocols that extend beyond conventional network Operations, Administration, and Management (OAM). The presented network telemetry framework promises flexibility, scalability, accuracy, coverage, and performance. In addition, it facilitates the implementation of automated control loops to address both today's and tomorrow's network operational needs. This document clarifies the terminologies and classifies the modules and components of a network telemetry system from several different perspectives. The framework and taxonomy help to set a common ground for the collection of related work and provide guidance for related technique and standard developments.

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 August 23, 2021.

Copyright Notice

Copyright (c) 2021 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
2.	Glossary	4
3.	Background	6
3.1.	Telemetry Data Coverage	7
3.2.	Use Cases	7
3.3.	Challenges	9
3.4.	Network Telemetry	10
4.	The Necessity of a Network Telemetry Framework	12
5.	Network Telemetry Framework	13
5.1.	Top Level Modules	14
5.1.1.	Management Plane Telemetry	17
5.1.2.	Control Plane Telemetry	17
5.1.3.	Forwarding Plane Telemetry	18
5.1.4.	External Data Telemetry	20
5.2.	Second Level Function Components	21
5.3.	Data Acquisition Mechanism and Type Abstraction	22
5.4.	Mapping Existing Mechanisms into the Framework	24
6.	Evolution of Network Telemetry Applications	25
7.	Security Considerations	26
8.	IANA Considerations	27
9.	Contributors	27
10.	Acknowledgments	28
11.	Informative References	28
Appendix A.	A Survey on Existing Network Telemetry Techniques	32

A.1.	Management Plane Telemetry	32
A.1.1.	Push Extensions for NETCONF	32
A.1.2.	gRPC Network Management Interface	32
A.2.	Control Plane Telemetry	33
A.2.1.	BGP Monitoring Protocol	33
A.3.	Data Plane Telemetry	33
A.3.1.	The Alternate Marking (AM) technology	33
A.3.2.	Dynamic Network Probe	34
A.3.3.	IP Flow Information Export (IPFIX) protocol	35
A.3.4.	In-Situ OAM	35
A.3.5.	Postcard Based Telemetry	35
A.4.	External Data and Event Telemetry	35
A.4.1.	Sources of External Events	36
A.4.2.	Connectors and Interfaces	37
	Authors' Addresses	37

[1.](#) Introduction

Network visibility is the ability of management tools to see the state and behavior of a network, which is essential for successful network operation. Network Telemetry revolves around network data that can help provide insights about the current state of the network, including network devices, forwarding, control, and management planes, and that can be generated and obtained through a variety of techniques, including but not limited to network instrumentation and measurements, and that can be processed for purposes ranging from service assurance to network security using a wide variety of techniques including machine learning, data analysis, and correlation. In this document, Network Telemetry refer to both the data itself (i.e., "Network Telemetry Data"), and the techniques and processes used to generate, export, collect, and consume that data for use by potentially automated management applications. Network telemetry extends beyond the conventional network Operations, Administration, and Management (OAM) techniques and expects to support better flexibility, scalability, accuracy, coverage, and performance.

However, the term of network telemetry lacks a solid and unambiguous definition. The scope and coverage of it cause confusion and misunderstandings. It is beneficial to clarify the concept and provide a clear architectural framework for network telemetry, so we can articulate the technical field, and better align the related techniques and standard works.

To fulfill such an undertaking, we first discuss some key characteristics of network telemetry which set a clear distinction from the conventional network OAM and show that some conventional OAM technologies can be considered a subset of the network telemetry

technologies. We then provide an architectural framework for network telemetry which includes four modules, each concerned with a different category of telemetry data and corresponding procedures. All the modules are internally structured in the same way, including components that allow to configure data sources with regards to what data to generate and how to make that available to client applications, components that instrument the underlying data sources, and components that perform the actual rendering, encoding, and exporting of the generated data. We show how the network telemetry framework can benefit the current and future network operations. Based on the distinction of modules and function components, we can map the existing and emerging techniques and protocols into the framework. The framework can also simplify the tasks for designing, maintaining, and understanding a network telemetry system. At last, we outline the evolution stages of the network telemetry system and discuss the potential security concerns.

The purpose of the framework and taxonomy is to set a common ground for the collection of related work and provide guidance for future technique and standard developments. To the best of our knowledge, this document is the first such effort for network telemetry in industry standards organizations.

2. Glossary

Before further discussion, we list some key terminology and acronyms used in this documents. We make an intended differentiation between the terms of network telemetry and OAM. However, it should be understood that there is not a hard-line distinction between the two concepts. Rather, network telemetry is considered as the extension of OAM. It covers all the existing OAM protocols but puts more emphasis on the newer and emerging techniques and protocols concerning all aspects of network data from acquisition to consumption.

AI: Artificial Intelligence. In network domain, AI refers to the machine-learning based technologies for automated network operation and other tasks.

AM: Alternate Marking, a flow performance measurement method, specified in [[RFC8321](#)].

BMP: BGP Monitoring Protocol, specified in [[RFC7854](#)].

DNP: Dynamic Network Probe, referring to programmable in-network sensors for network monitoring and measurement.

DPI: Deep Packet Inspection, referring to the techniques that examines packet beyond packet L3/L4 headers.

gNMI: gRPC Network Management Interface, a network management protocol from OpenConfig Operator Working Group, mainly contributed by Google. See [[gnmi](#)] for details.

gRPC: gRPC Remote Procedure Call, a open source high performance RPC framework that gNMI is based on. See [[grpc](#)] for details.

IPFIX: IP Flow Information Export Protocol, specified in [[RFC7011](#)].

IOAM: In-situ OAM, a dataplane on-path telemetry technique.

NETCONF: Network Configuration Protocol, specified in [[RFC6241](#)].

NetFlow: A Cisco protocol for flow record collecting, described in [[RFC3594](#)].

Network Telemetry: The process and instrumentation for acquiring and utilizing network data remotely for network monitoring and operation. A general term for a large set of network visibility techniques and protocols, concerning aspects like data generation, collection, correlation, and consumption. Network telemetry addresses the current network operation issues and enables smooth evolution toward future intent-driven autonomous networks.

NMS: Network Management System, referring to applications that allow network administrators manage a network.

OAM: Operations, Administration, and Maintenance. A group of network management functions that provide network fault indication, fault localization, performance information, and data and diagnosis functions. Most conventional network monitoring techniques and protocols belong to network OAM.

PBT: Postcard-Based Telemetry, a dataplane on-path telemetry technique.

SMIV2 Structure of Management Information Version 2, specified in [[RFC2578](#)].

SNMP: Simple Network Management Protocol. Version 1 and 2 are specified in [[RFC1157](#)] and [[RFC3416](#)], respectively.

YANG: The abbreviation of "Yet Another Next Generation". YANG is a data modeling language for the definition of data sent over

network management protocols such as the NETCONF and RESTCONF. YANG is defined in [[RFC6020](#)].

YANG ECA A YANG model for Event-Condition-Action policies, defined in [[I-D.wwx-netmod-event-yang](#)].

YANG PUSH: A method to subscribe pushed data from remote YANG datastore on network devices. Details are specified in [[RFC8641](#)] and [[RFC8639](#)].

3. Background

The term "big data" is used to describe the extremely large volume of data sets that can be analyzed computationally to reveal patterns, trends, and associations. Networks are undoubtedly a source of big data because of their scale and the volume of network traffic they forward. It is easy to see that network operations can benefit from network big data.

Today one can access advanced big data analytics capability through a plethora of commercial and open source platforms (e.g., Apache Hadoop), tools (e.g., Apache Spark), and techniques (e.g., machine learning). Thanks to the advance of computing and storage technologies, network big data analytics gives network operators an opportunity to gain network insights and move towards network autonomy. Some operators start to explore the application of Artificial Intelligence (AI) to make sense of network data. Software tools can use the network data to detect and react on network faults, anomalies, and policy violations, as well as predicting future events. In turn, the network policy updates for planning, intrusion prevention, optimization, and self-healing may be applied.

It is conceivable that an autonomic network [[RFC7575](#)] is the logical next step for network evolution following Software Defined Network (SDN), aiming to reduce (or even eliminate) human labor, make more efficient use of network resources, and provide better services more aligned with customer requirements. Intent-based Networking (IBN) [[I-D.irtf-nmrg-ibn-concepts-definitions](#)] requires network visibility and telemetry data in order to ensure that the network is behaving as intended. Although it takes time to reach the ultimate goal, the journey has started nevertheless.

However, while the data processing capability is improved and applications are hungry for more data, the networks lag behind in extracting and translating network data into useful and actionable information in efficient ways. The system bottleneck is shifting from data consumption to data supply. Both the number of network nodes and the traffic bandwidth keep increasing at a fast pace. The

network configuration and policy change at smaller time slots than before. More subtle events and fine-grained data through all network planes need to be captured and exported in real time. In a nutshell, it is a challenge to get enough high-quality data out of the network in a manner that is efficient, timely, and flexible. Therefore, we need to survey the existing technologies and protocols and identify any potential gaps.

In the remainder of this section, first we clarify the scope of network data (i.e., telemetry data) concerned in the context. Then, we discuss several key use cases for today's and future network operations. Next, we show why the current network OAM techniques and protocols are insufficient for these use cases. The discussion underlines the need of new methods, techniques, and protocols which we assign under the umbrella term - Network Telemetry.

3.1. Telemetry Data Coverage

Any information that can be extracted from networks (including data plane, control plane, and management plane) and used to gain visibility or as basis for actions is considered telemetry data. It includes statistics, event records and logs, snapshots of state, configuration data, etc. It also covers the outputs of any active and passive measurements [[RFC7799](#)]. Specially, raw data can be processed in-network before being sent to a data consumer. Such processed data is also considered telemetry data. A classification of telemetry data is provided in [Section 5](#).

3.2. Use Cases

The following set of use cases is essential for network operations. While the list is by no means exhaustive, it is enough to highlight the requirements for data velocity, variety, volume, and veracity in networks.

- o Security: Network intrusion detection and prevention systems need to monitor network traffic and activities and act upon anomalies. Given increasingly sophisticated attack vector coupled with increasingly severe consequences of security breaches, new tools and techniques need to be developed, relying on wider and deeper visibility into networks.
- o Policy and Intent Compliance: Network policies are the rules that constraint the services for network access, provide service differentiation, or enforce specific treatment on the traffic. For example, a service function chain is a policy that requires the selected flows to pass through a set of ordered network functions. Intent, as defined in

[[I-D.irtf-nmrg-ibn-concepts-definitions](#)], is a set of operational goal that a network should meet and outcomes that a network is supposed to deliver, defined in a declarative manner without specifying how to achieve or implement them. An intent requires a complex translation and mapping process before being applied on networks. While a policy or an intent is enforced, the compliance needs to be verified and monitored continuously, relying on visibility that is provided through network telemetry data, and any violation needs to be reported immediately.

- o **SLA Compliance:** A Service-Level Agreement (SLA) defines the level of service a user expects from a network operator, which include the metrics for the service measurement and remedy/penalty procedures when the service level misses the agreement. Users need to check if they get the service as promised and network operators need to evaluate how they can deliver the services that can meet the SLA based on realtime network telemetry data, including data from network measurements.
- o **Root Cause Analysis:** Any network failure can be the effect of a sequence of chained events. Troubleshooting and recovery require quick identification of the root cause of any observable issues. However, the root cause is not always straightforward to identify, especially when the failure is sporadic and the number of event messages, both related and unrelated to the same cause, is overwhelming. While machine learning technologies can be used for root cause analysis, it up to the network to sense and provide the relevant data to feed into machine learning applications.
- o **Network Optimization:** This covers all short-term and long-term network optimization techniques, including load balancing, Traffic Engineering (TE), and network planning. Network operators are motivated to optimize their network utilization and differentiate services for better Return On Investment (ROI) or lower Capital Expenditures (CAPEX). The first step is to know the real-time network conditions before applying policies for traffic manipulation. In some cases, micro-bursts need to be detected in a very short time-frame so that fine-grained traffic control can be applied to avoid network congestion. Long-term planning of network capacity and topology requires analysis of real-world network telemetry data that is obtained over long periods of time.
- o **Event Tracking and Prediction:** The visibility into traffic path and performance is critical for services and applications that rely on healthy network operation. Numerous related network events are of interest to network operators. For example, Network operators want to learn where and why packets are dropped for an application flow. They also want to be warned of issues in

advance so proactive actions can be taken to avoid catastrophic consequences.

3.3. Challenges

For a long time, network operators have relied upon SNMP [[RFC3416](#)], Command-Line Interface (CLI), or Syslog to monitor the network. Some other OAM techniques as described in [[RFC7276](#)] are also used to facilitate network troubleshooting. These conventional techniques are not sufficient to support the above use cases for the following reasons:

- o Most use cases need to continuously monitor the network and dynamically refine the data collection in real-time. The poll-based low-frequency data collection is ill-suited for these applications. Subscription-based streaming data directly pushed from the data source (e.g., the forwarding chip) is preferred to provide enough data quantity and precision at scale.
- o Comprehensive data is needed from packet processing engine to traffic manager, from line cards to main control board, from user flows to control protocol packets, from device configurations to operations, and from physical layer to application layer. Conventional OAM only covers a narrow range of data (e.g., SNMP only handles data from the Management Information Base (MIB)). Traditional network devices cannot provide all the necessary probes. More open and programmable network devices are therefore needed.
- o Many application scenarios need to correlate network-wide data from multiple sources (i.e., from distributed network devices, different components of a network device, or different network planes). A piecemeal solution is often lacking the capability to consolidate the data from multiple sources. The composition of a complete solution, as partly proposed by Autonomic Resource Control Architecture(ARCA) [[I-D.pedro-nmrg-anticipated-adaptation](#)], will be empowered and guided by a comprehensive framework.
- o Some of the conventional OAM techniques (e.g., CLI and Syslog) lack a formal data model. The unstructured data hinder the tool automation and application extensibility. Standardized data models are essential to support the programmable networks.
- o Although some conventional OAM techniques support data push (e.g., SNMP Trap [[RFC2981](#)][[RFC3877](#)], Syslog, and sFlow), the pushed data are limited to only predefined management plane warnings (e.g., SNMP Trap) or sampled user packets (e.g., sFlow). Network

operators require the data with arbitrary source, granularity, and precision which are beyond the capability of the existing techniques.

- o The conventional passive measurement techniques can either consume excessive network resources and render excessive redundant data, or lead to inaccurate results; on the other hand, the conventional active measurement techniques can interfere with the user traffic and their results are indirect. Techniques that can collect direct and on-demand data from user traffic are more favorable.

These challenges were addressed by newer standards and techniques (e.g., IPFIX/Netflow, PSAMP, IOAM, and YANG-Push) and more are emerging. These standards and techniques need to be recognized and accommodated in a new framework.

3.4. Network Telemetry

Network telemetry has emerged as a mainstream technical term to refer to the network data collection and consumption techniques. Several network telemetry techniques and protocols (e.g., IPFIX [[RFC7011](#)] and gRPC [[grpc](#)]) have been widely deployed. Network telemetry allows separate entities to acquire data from network devices so that data can be visualized and analyzed to support network monitoring and operation. Network telemetry covers the conventional network OAM and has a wider scope. It is expected that network telemetry can provide the necessary network insight for autonomous networks and address the shortcomings of conventional OAM techniques.

Network telemetry usually assumes machines as data consumers rather than human operators. Hence, the network telemetry can directly trigger the automated network operation, while in contrast some conventional OAM tools are designed and used to help human operators to monitor and diagnose the networks and guide manual network operations. Such a proposition leads to very different techniques.

Although new network telemetry techniques are emerging and subject to continuous evolution, several characteristics of network telemetry have been well accepted. Note that network telemetry is intended to be an umbrella term covering a wide spectrum of techniques, so the following characteristics are not expected to be held by every specific technique.

- o Push and Streaming: Instead of polling data from network devices, telemetry collectors subscribe to streaming data pushed from data sources in network devices.

- o Volume and Velocity: The telemetry data is intended to be consumed by machines rather than by human being. Therefore, the data volume can be huge and the processing is optimized for the needs of automation in realtime.
- o Normalization and Unification: Telemetry aims to address the overall network automation needs. Efforts are made to normalize the data representation and unify the protocols, so to simplify data analysis and provide integrated analysis across heterogeneous devices and data sources across a network.
- o Model-based: The telemetry data is modeled in advance which allows applications to configure and consume data with ease.
- o Data Fusion: The data for a single application can come from multiple data sources (e.g., cross-domain, cross-device, and cross-layer) and needs to be correlated to take effect.
- o Dynamic and Interactive: Since the network telemetry means to be used in a closed control loop for network automation, it needs to run continuously and adapt to the dynamic and interactive queries from the network operation controller.

In addition, an ideal network telemetry solution may also have the following features or properties:

- o In-Network Customization: The data that is generated can be customized in network at run-time to cater to the specific need of applications. This needs the support of a programmable data plane which allows probes with custom functions to be deployed at flexible locations.
- o In-Network Data Aggregation and Correlation: Network devices and aggregation points can work out which events and what data needs to be stored, reported, or discarded thus reducing the load on the central collection and processing points while still ensuring that the right information is ready to be processed in a timely way.
- o In-Network Processing: Sometimes it is not necessary or feasible to gather all information to a central point to be processed and acted upon. It is possible for the data processing to be done in network, allowing reactive actions to be taken locally.
- o Direct Data Plane Export: The data originated from the data plane forwarding chips can be directly exported to the data consumer for efficiency, especially when the data bandwidth is large and the real-time processing is required.

- o In-band Data Collection: In addition to the passive and active data collection approaches, the new hybrid approach allows to directly collect data for any target flow on its entire forwarding path [[I-D.song-opsawg-ifit-framework](#)].

It is worth noting that a network telemetry system should not be intrusive to normal network operations by avoiding the pitfall of the "observer effect". That is, it should not change the network behavior and affect the forwarding performance. Otherwise, the whole purpose of network telemetry is compromised.

Although in many cases a system for network telemetry involves a remote data collecting and consuming entity, it is important to understand that there are no inherent assumptions about how a system should be architected. Telemetry data producers and consumers can work in distributed or peer-to-peer fashions rather than assuming a centralized data consuming entity. In such cases, a network node can be the direct consumer of telemetry data from other nodes.

4. The Necessity of a Network Telemetry Framework

Network data analytics and machine-learning technologies are applied for network operation automation, relying on abundant and coherent data from networks. Data acquisition that is limited to a single source and static in nature will in many cases not be sufficient to meet an application's telemetry data needs. As a result, multiple data sources, involving a variety of techniques and standards, will need to be integrated. It is desirable to have a framework that classifies and organizes different telemetry data source and types, defines different components of a network telemetry system and their interactions, and helps coordinate and integrate multiple telemetry approaches across layers. This allows flexible combinations of data for different applications, while normalizing and simplifying interfaces. In detail, such a framework would benefit application development for the following reasons:

- o Future networks, autonomous or otherwise, depend on holistic and comprehensive network visibility. All the use cases and applications are better to be supported uniformly and coherently under a single intelligent agent using an integrated, converged mechanism and common telemetry data representations wherever feasible. Therefore, the protocols and mechanisms should be consolidated into a minimum yet comprehensive set. A telemetry framework can help to normalize the technique developments.
- o Network visibility presents multiple viewpoints. For example, the device viewpoint takes the network infrastructure as the monitoring object from which the network topology and device

status can be acquired; the traffic viewpoint takes the flows or packets as the monitoring object from which the traffic quality and path can be acquired. An application may need to switch its viewpoint during operation. It may also need to correlate a service and its impact on user experience to acquire the comprehensive information.

- o Applications require network telemetry to be elastic in order to make efficient use of network resources and reduce the impact of processing related to network telemetry on network performance. For example, routine network monitoring should cover the entire network with a low data sampling rate. Only when issues arise or critical trends emerge should telemetry data source be modified and telemetry data rates boosted as needed.
- o Efficient data fusion is critical for applications to reduce the overall quantity of data and improve the accuracy of analysis.

A telemetry framework collects together all of the telemetry-related works from different sources and working groups within IETF. This makes it possible to assemble a comprehensive network telemetry system and to avoid repetitious or redundant work. The framework should cover the concepts and components from the standardization perspective. This document describes the modules which make up a network telemetry framework and decomposes the telemetry system into a set of distinct components that existing and future work can easily map to.

5. Network Telemetry Framework

The top level network telemetry framework partitions the network telemetry into four modules based on the telemetry data object source and represents their relationship. At the next level, the framework decomposes each module into separate components. Each of the modules follows the same underlying structure, with one component dedicated to the configuration of data subscriptions and data sources, a second component dedicated to encoding and exporting data, and a third component instrumenting the generation of telemetry related to the underlying resources. Throughout the framework, the same set of abstract data acquiring mechanisms and data types are applied. The two-level architecture with the uniform data abstraction helps accurately pinpoint a protocol or technique to its position in a network telemetry system or disaggregate a network telemetry system into manageable parts.

5.1. Top Level Modules

Telemetry can be applied on the forwarding plane, the control plane, and the management plane in a network, as well as other sources out of the network, as shown in Figure 1. Therefore, we categorize the network telemetry into four distinct modules with each having its own interface to Network Operation Applications.

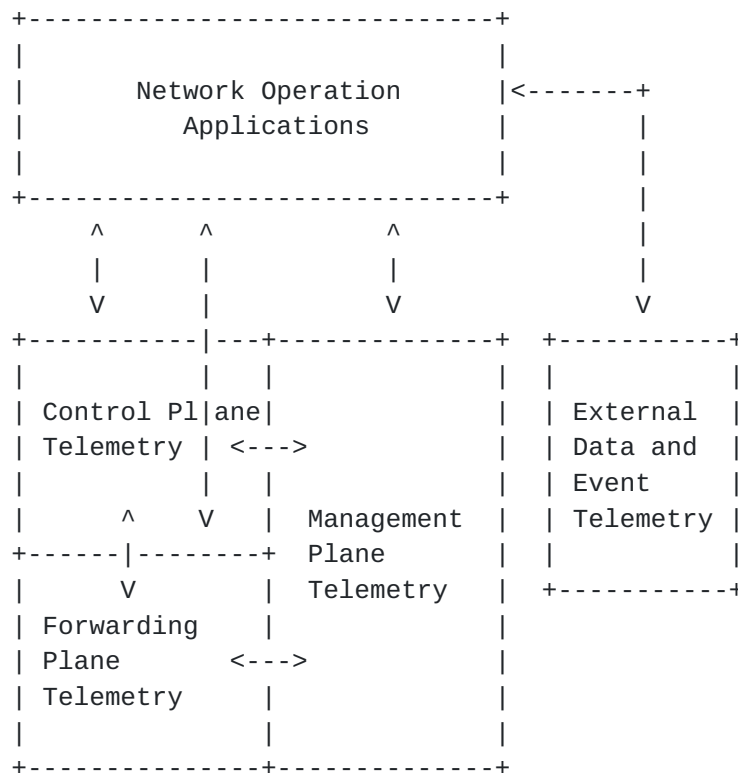


Figure 1: Modules in Layer Category of NTF

The rationale of this partition lies in the different telemetry data objects which result in different data source and export locations. Such differences have profound implications on in-network data programming and processing capability, data encoding and transport protocol, and required data bandwidth and latency.

We summarize the major differences of the four modules in the following table. They are compared from six angles:

- o Data Object
- o Data Export Location
- o Data Model

- o Data Encoding
- o Telemetry Protocol
- o Transport Method

Data Object is the target and source of each module. Because the data source varies, the location where data is mostly conveniently exported also varies. For example, forwarding plane data mainly originates from the fast path(e.g., forwarding chips) while control plane data mainly originates from the slow path (e.g., main control CPU). For convenience and efficiency, it is preferred to export the data from locations near the source. Because each location that can export data has different capability, the proper data model, encoding, and transport method cannot be kept the same. For example, the forwarding chip has high throughput but limited capacity for processing complex data and maintaining states, while the main control CPU is capable of complex data and state processing, but has limited bandwidth for high throughput data. As a result, the suitable telemetry protocol for each module can be different. Some representative techniques are shown in the corresponding table blocks to highlight the technical diversity of these modules. Note that the selected techniques just reflect the de-facto state of the art and are not exhaustive. The key point is that one cannot expect to use a universal protocol to cover all the network telemetry requirements.

Module	Control Plane	Management Plane	Forwarding Plane	External Data
Object	control protocol & signaling, RIB, ACL	config. & operation state, MIB	flow & packet QoS, traffic stat., buffer & queue stat.	terminal, social & environ-mental
Export Location	main control CPU, linecard CPU or fwding chip	main control CPU	fwding chip or linecard CPU; main control CPU unlikely	various
Data Model	YANG, custom	MIB, syslog, YANG, custom	template, YANG, custom	YANG
Data Encoding	GPB, JSON, XML, plain	GPB, JSON, XML	plain	GPB, JSON, XML, plain
Protocol	gRPC, NETCONF, IPFIX, mirror	gRPC, NETCONF	IPFIX, mirror	gRPC
Transport	HTTP, TCP, UDP	HTTP, TCP	UDP	HTTP, TCP, UDP

Figure 2: Comparison of the Data Object Modules

Note that the interaction with the applications that consume network telemetry data can be indirect. Some in-device data transfer is possible. For example, in the management plane telemetry, the management plane may need to acquire data from the data plane. Some of the operational states can only be derived from data plane data sources such as the interface status and statistics. For another example, obtaining control plane telemetry data may require the ability access the Forwarding Information Base (FIB) of the data plane.

On the other hand, an application may involve more than one plane and interact with multiple planes simultaneously. For example, an SLA compliance application may require both the data plane telemetry and the control plane telemetry.

The requirements and challenges for each module are summarized as follows (note that the requirements may pertain across all telemetry modules; however, we emphasize those that are most pronounced for a particular plane).

5.1.1. Management Plane Telemetry

The management plane of network elements interacts with the Network Management System (NMS), and provides information such as performance data, network logging data, network warning and defects data, and network statistics and state data. The management plane includes many protocols, including some that are considered "legacy", such as SNMP and syslog. Regardless the protocol, management plane telemetry must address the following requirements:

- o Convenient Data Subscription: An application should have the freedom to choose the data export means such as the data types and the export frequency.
- o Structured Data: For automatic network operation, machines will replace human for network data comprehension. The schema languages such as YANG can efficiently describe structured data and normalize data encoding and transformation.
- o High Speed Data Transport: In order to keep up with the velocity of information, a server needs to be able to send large amounts of data at high frequency. Compact encoding formats are needed to compress the data and improve the data transport efficiency. The subscription mode, by replacing the query mode, reduces the interactions between clients and servers and helps to improve the server's efficiency.

5.1.2. Control Plane Telemetry

The control plane telemetry refers to the health condition monitoring of different network control protocols covering Layer 2 to Layer 7. Keeping track of the running status of these protocols is beneficial for detecting, localizing, and even predicting various network issues, as well as network optimization, in real-time and in fine granularity. Some particular challenges and issues faced by the control plane telemetry are as follows:

- o One challenging problem for the control plane telemetry is how to correlate the End-to-End (E2E) Key Performance Indicators (KPI) to a specific layer's KPIs. For example, an IPTV user may describe his User Experience (UE) by the video fluency and definition. Then in case of an unusually poor UE KPI or a service disconnection, it is non-trivial to delimit and pinpoint the issue

in the responsible protocol layer (e.g., the Transport Layer or the Network Layer), the responsible protocol (e.g., ISIS or BGP at the Network Layer), and finally the responsible device(s) with specific reasons.

- o Traditional OAM-based approaches for control plane KPI measurement include PING (L3), Tracert (L3), Y.1731 (L2), and so on. One common issue behind these methods is that they only measure the KPIs instead of reflecting the actual running status of these protocols, making them less effective or efficient for control plane troubleshooting and network optimization.
- o An example of the control plane telemetry is the BGP monitoring protocol (BMP), it is currently used to monitoring the BGP routes and enables rich applications, such as BGP peer analysis, AS analysis, prefix analysis, security analysis, and so on. However, the monitoring of other layers, protocols and the cross-layer, cross-protocol KPI correlations are still in their infancy (e.g., the IGP monitoring is missing), which require further research.

5.1.3. Forwarding Plane Telemetry

An effective forwarding plane telemetry system relies on the data that the network device can expose. The quality, quantity, and timeliness of data must meet some stringent requirements. This raises some challenges to the network data plane devices where the first hand data originates.

- o A data plane device's main function is user traffic processing and forwarding. While supporting network visibility is important, the telemetry is just an auxiliary function, and it should not impede normal traffic processing and forwarding (i.e., the performance is not lowered and the behavior is not altered due to the telemetry functions).
- o Network operation applications require end-to-end visibility across various sources, which can result in a huge volume of data. However, the sheer data quantity should not exhaust the network bandwidth, regardless of the data delivery approach (i.e., whether through in-band or out-of-band channels).
- o The data plane devices must provide timely data with the minimum possible delay. Long processing, transport, storage, and analysis delay can impact the effectiveness of the control loop and even render the data useless.
- o The data should be structured and labeled, and easy for applications to parse and consume. At the same time, the data

types needed by applications can vary significantly. The data plane devices need to provide enough flexibility and programmability to support the precise data provision for applications.

- o The data plane telemetry should support incremental deployment and work even though some devices are unaware of the system. This challenge is highly relevant to the standards and legacy networks.

Although not specific to the forwarding plane, these challenges are more difficult to the forwarding plane because of the limited resource and flexibility. The data plane programmability is essential to support network telemetry. Newer data plane forwarding chips are equipped with advanced telemetry features and provide flexibility to support customized telemetry functions.

Technique Taxonomy: concerning about how one instruments the telemetry, there can be multiple possible dimensions to classify the forwarding plane telemetry techniques.

- o Active, Passive, and Hybrid: This dimension concerns about the end-to-end measurement. Active and passive methods (as well as the hybrid types) are well documented in [\[RFC7799\]](#). Passive methods include TCPDUMP, IPFIX [\[RFC7011\]](#), sflow, and traffic mirroring. These methods usually have low data coverage. The bandwidth cost is very high in order to improve the data coverage. On the other hand, active methods include Ping, OWAMP [\[RFC4656\]](#), TWAMP [\[RFC5357\]](#), and Cisco's SLA Protocol [\[RFC6812\]](#). These methods are intrusive and only provide indirect network measurement results. Hybrid methods, including in-situ OAM [\[I-D.ietf-ippm-ioam-data\]](#), Alternate-Marking (AM) [\[RFC8321\]](#), and Multipoint Alternate Marking [\[I-D.fioccola-ippm-multipoint-alt-mark\]](#), provide a well-balanced and more flexible approach. However, these methods are also more complex to implement.
- o In-Band and Out-of-Band: The telemetry data, before being exported to some collector, can be carried in user packets. Such methods are considered in-band (e.g., in-situ OAM [\[I-D.ietf-ippm-ioam-data\]](#)). If the telemetry data is directly exported to some collector without modifying the user packets, such methods are considered out-of-band (e.g., postcard-based INT). It is possible to have hybrid methods. For example, only the telemetry instruction or partial data is carried by user packets (e.g., AM [\[RFC8321\]](#)).
- o E2E and In-Network: Some E2E methods start from and end at the network end hosts (e.g., Ping). The other methods work in

networks and are transparent to end hosts. However, if needed, in-network methods can be easily extended into end hosts.

- o Data Subject: Depending on the telemetry objective, the methods can be flow-based (e.g., in-situ OAM [[I-D.ietf-ippm-ioam-data](#)]), path-based (e.g., Traceroute), and node-based (e.g., IPFIX [[RFC7011](#)]). The various data objects can be packet, flow record, measurement, states, and signal.

5.1.4. External Data Telemetry

Events that occur outside the boundaries of the network system are another important source of network telemetry. Correlating both internal telemetry data and external events with the requirements of network systems, as presented in [[I-D.pedro-nmrg-anticipated-adaptation](#)], provides a strategic and functional advantage to management operations.

As with other sources of telemetry information, the data and events must meet strict requirements, especially in terms of timeliness, which is essential to properly incorporate external event information to management cycles. The specific challenges are described as follows:

- o The role of external event detector can be played by multiple elements, including hardware (e.g. physical sensors, such as seismometers) and software (e.g. Big Data sources that analyze streams of information, such as Twitter messages). Thus, the transmitted data must support different shapes but, at the same time, follow a common but extensible schema.
- o Since the main function of the external event detectors is to perform the notifications, their timeliness is assumed. However, once messages have been dispatched, they must be quickly collected and inserted into the control plane with variable priority, which will be high for important sources and/or important events and low for secondary ones.
- o The schema used by external detectors must be easily adopted by current and future devices and applications. Therefore, it must be easily mapped to current information models, such as in terms of YANG.

Organizing together both internal and external telemetry information will be key for the general exploitation of the management possibilities of current and future network systems, as reflected in the incorporation of cognitive capabilities to new hardware and software (virtual) elements.

5.2. Second Level Function Components

Reflecting the best current practice, the telemetry module at each plane is further partitioned into five distinct components:

- o Data Query, Analysis, and Storage: This component works at the application layer. It is a part of the network management system at the receiver side. On the one hand, it is responsible for issuing data requirements. The data of interest can be modeled data through configuration or custom data through programming. The data requirements can be queries for one-shot data or subscriptions for events or streaming data. On the other hand, it receives, stores, and processes the returned data from network devices. Data analysis can be interactive to initiate further data queries. This component can reside in either network devices or remote controllers. It can be centralized and distributed, and involve one or more instances.
- o Data Configuration and Subscription: This component deploys data queries on devices. It determines the protocol and channel for applications to acquire desired data. This component is also responsible for configuring the desired data that might not be directly available from data sources. The subscription data can be described by models, templates, or programs.
- o Data Encoding and Export: This component determines how telemetry data is delivered to the data analysis and storage component. The data encoding and the transport protocol may vary due to the data exporting location.
- o Data Generation and Processing: The requested data needs to be captured, processed, and formatted in network devices from raw data sources. This may involve in-network computing and processing on either the fast path or the slow path in network devices.
- o Data Object and Source: This component determines the monitoring object and original data source. The data source usually just provides raw data which needs further processing. A data source can be considered a probe. Some data sources can be dynamically installed, while others will be more static.

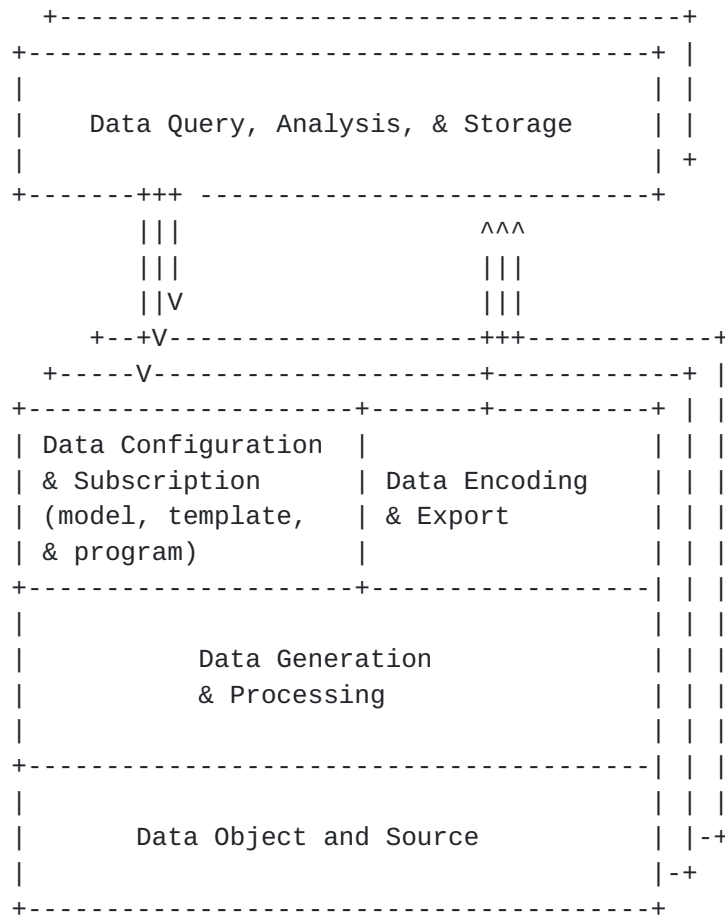


Figure 3: Components in the Network Telemetry Framework

5.3. Data Acquisition Mechanism and Type Abstraction

Broadly speaking, network data can be acquired through subscription (push) and query (poll). A subscription is a contract between publisher and subscriber. After initial setup, the subscribed data is automatically delivered to registered subscribers until the subscription expires. Subscription can be partitioned into two sub modes: the Publish-Subscription (Pub-Sub) mode and the Subscription-Publish (Sub-Pub) mode. In the Pub-Sub mode, a publisher publishes pre-defined data and any qualified subscribers can subscribe the data as-is. In the Sub-Pub mode, a subscriber initiates a data request and sends it to a publisher; the publisher will deliver the requested data when available. While for both modes, the subscribed data is pushed to the subscriber, the Sub-Pub mode allows subscribers to customize their subscriptions.

In contrast, query is used when a querier expects immediate and one-off feedback from network devices. The queried data may be directly extracted from some specific data source, or synthesized and processed from raw data. Query suits for interactive network telemetry applications.

There are four types of data from network devices that a telemetry data consumer can subscribe or query:

- o Simple Data: The data that are steadily available from some data store or static probes in network devices. such data can be specified by YANG model.
- o Complex Data: The data need to be synthesized or processed in network from raw data from one or more network devices. The data processing function can be statically or dynamically loaded into network devices.
- o Event-triggered Data: The data are conditionally acquired based on the occurrence of some events. It can be actively pushed through subscription or passively polled through query. There are many ways to model events, including using Finite State Machine (FSM) or Event Condition Action (ECA) [[I-D.wwx-netmod-event-yang](#)].
- o Streaming Data: The data are continuously generated. It can be time series or the dump of databases. The streaming data reflect realtime network states and metrics and require large bandwidth and processing power. The streaming data are always actively pushed to the subscribers.

The above data types are not mutually exclusive. Rather, they often overlap. For example, event-triggered data can be simple or complex, and streaming data can be simple, complex, or triggered by events. The relationships of these data types are illustrated in Figure 4.

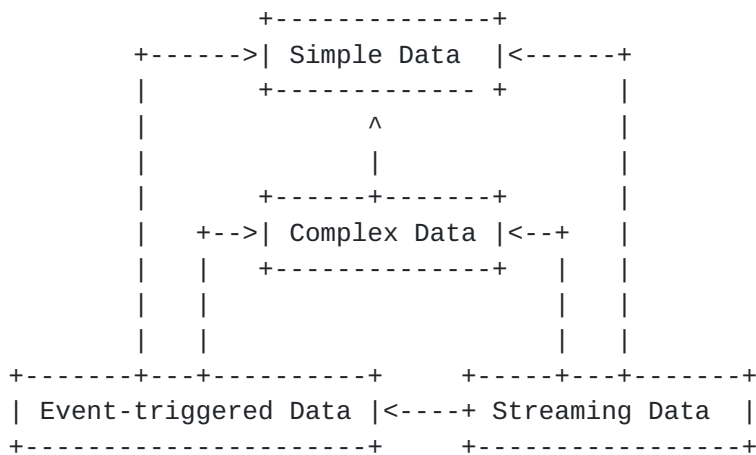


Figure 4: Data Type Relationship

Subscription usually deals with event-triggered data and streaming data, and query usually deals with simple data and complex data. But the other ways are also possible. The conventional OAM techniques are mostly about querying simple data. While these techniques are still useful, more advanced network telemetry techniques are designed mainly for event-triggered or streaming data subscription, and complex data query.

5.4. Mapping Existing Mechanisms into the Framework

The following two tables show how the existing mechanisms (mainly published in IETF and with the emphasis on the latest new technologies) are positioned in the framework. Given the vast body of existing work, we cannot provide an exhaustive list, so the mechanisms in the tables should be considered as just examples. Also, some comprehensive protocols and techniques may cover multiple aspects or modules of the framework, so a name in a block only emphasizes one particular characteristic of it. More details about some listed mechanisms can be found in [Appendix A](#).

The first table is based on the data acquisition mechanisms and data types.

	Query	Subscription
Simple Data	SNMP	YANG
Complex Data	DNP	YANG PUSH
Event-triggered Data	DNP	YANG PUSH
Streaming Data	N/A	gRPC

Figure 5: Existing Work Mapping I

The second table is based on the telemetry modules and components.

	Management Plane	Control Plane	Forwarding Plane
data config. & subscribe	gRPC, NETCONF, SMIPv2, YANG PUSH	NETCONF/YANG YANG PUSH	NETCONF/YANG, YANG PUSH
data gen. & process	DNP, YANG	DNP, YANG	IOAM, PSAMP PBT, AM, DNP
data export	gRPC, NETCONF YANG PUSH	BMP, NETCONF	IPFIX

Figure 6: Existing Work Mapping II

6. Evolution of Network Telemetry Applications

Network telemetry is a fast evolving technical area. As the network moves towards the automated operation, network telemetry applications undergo several stages of evolution which add new layer of requirements to the underlying network telemetry techniques. Each stage is built upon the techniques adopted by the previous stages plus some new requirements.

Stage 0 - Static Telemetry: The telemetry data source and type are determined at design time. The network operator can only configure how to use it with limited flexibility.

Stage 1 - Dynamic Telemetry: The custom telemetry data can be dynamically programmed or configured at runtime without interrupting the network operation, allowing a tradeoff among resource, performance, flexibility, and coverage. DNP is an effort towards this direction.

Stage 2 - Interactive Telemetry: The network operator can continuously customize and fine tune the telemetry data in real time to reflect the network operation's visibility requirements. Compared with Stage 1, the changes are frequent based on the real-time feedback. At this stage, some tasks can be automated, but human operators still need to sit in the middle to make decisions.

Stage 3 - Closed-loop Telemetry: The telemetry is free from the interference of human operators, except for generating the reports. The intelligent network operation engine automatically issues the telemetry data requests, analyzes the data, and updates the network operations in closed control loops.

Existing technologies are ready for stage 0 and stage 1. Individual stage 2 and stage 3 applications are also possible now. However, the future autonomic networks may need a comprehensive operation management system which works at stage 2 and stage 3 to cover all the network operation tasks. A well-defined network telemetry framework is the first step towards this direction.

7. Security Considerations

The complexity of network telemetry raises significant security implications. For example, telemetry data can be manipulated to exhaust various network resources at each plane as well as the data consumer; falsified or tampered data can mislead the decision making and paralyze networks; wrong configuration and programming for telemetry is equally harmful.

Given that this document has proposed a framework for network telemetry and the telemetry mechanisms discussed are more extensive (in both message frequency and traffic amount) than the conventional network OAM concepts, we must also reflect that various new security considerations may also arise. A number of techniques already exist for securing the forwarding plane, the control plane, and the management plane in a network, but it is important to consider if any new threat vectors are now being enabled via the use of network telemetry procedures and mechanisms.

Security considerations for networks that use telemetry methods may include:

- o Telemetry framework trust and policy model;
- o Role management and access control for enabling and disabling telemetry capabilities;
- o Protocol transport used telemetry data and inherent security capabilities;
- o Telemetry data stores, storage encryption and methods of access;
- o Tracking telemetry events and any abnormalities that might identify malicious attacks using telemetry interfaces.
- o Authentication and signing of telemetry data to make data more trustworthy.

Some of the security considerations highlighted above may be minimized or negated with policy management of network telemetry. In a network telemetry deployment it would be advantageous to separate telemetry capabilities into different classes of policies, i.e., Role Based Access Control and Event-Condition-Action policies. Also, potential conflicts between network telemetry mechanisms must be detected accurately and resolved quickly to avoid unnecessary network telemetry traffic propagation escalating into an unintended or intended denial of service attack.

Further study of the security issues will be required, and it is expected that the security mechanisms and protocols are developed and deployed along with a network telemetry system.

8. IANA Considerations

This document includes no request to IANA.

9. Contributors

The other contributors of this document are listed as follows.

- o Tianran Zhou
- o Zhenbin Li
- o Zhenqiang Li
- o Daniel King

- o Adrian Farrel
- o Alexander Clemm

10. Acknowledgments

We would like to thank Greg Mirsky, Randy Presuhn, Joe Clarke, Victor Liu, James Guichard, Uri Blumenthal, Giuseppe Fioccola, Yunan Gu, Parviz Yegani, Young Lee, Qin Wu, and many others who have provided helpful comments and suggestions to improve this document.

11. Informative References

- [gnmi] "gNMI - gRPC Network Management Interface",
<<https://github.com/openconfig/reference/tree/master/rpc/gnmi>>.
- [grpc] "gRPC, A high performance, open-source universal RPC framework", <<https://grpc.io>>.
- [I-D.fioccola-ippm-multipoint-alt-mark]
Fioccola, G., Cociglio, M., Sapio, A., and R. Sisto,
"Multipoint Alternate Marking method for passive and
hybrid performance monitoring", [draft-fioccola-ippm-multipoint-alt-mark-04](#) (work in progress), June 2018.
- [I-D.ietf-grow-bmp-adj-rib-out]
Evens, T., Bayraktar, S., Lucente, P., Mi, K., and S.
Zhuang, "Support for Adj-RIB-Out in BGP Monitoring
Protocol (BMP)", [draft-ietf-grow-bmp-adj-rib-out-07](#) (work
in progress), August 2019.
- [I-D.ietf-grow-bmp-local-rib]
Evens, T., Bayraktar, S., Bhardwaj, M., and P. Lucente,
"Support for Local RIB in BGP Monitoring Protocol (BMP)",
[draft-ietf-grow-bmp-local-rib-09](#) (work in progress),
January 2021.
- [I-D.ietf-ippm-ioam-data]
Brockners, F., Bhandari, S., and T. Mizrahi, "Data Fields
for In-situ OAM", [draft-ietf-ippm-ioam-data-11](#) (work in
progress), November 2020.
- [I-D.ietf-netconf-distributed-notif]
Zhou, T., Zheng, G., Voit, E., Graf, T., and P. Francois,
"Subscription to Distributed Notifications", [draft-ietf-netconf-distributed-notif-01](#) (work in progress), November 2020.

[I-D.ietf-netconf-udp-notif]

Zheng, G., Zhou, T., Graf, T., Francois, P., and P. Lucente, "UDP-based Transport for Configured Subscriptions", [draft-ietf-netconf-udp-notif-01](#) (work in progress), November 2020.

[I-D.irtf-nmrg-ibn-concepts-definitions]

Clemm, A., Ciavaglia, L., Granville, L., and J. Tantsura, "Intent-Based Networking - Concepts and Definitions", [draft-irtf-nmrg-ibn-concepts-definitions-02](#) (work in progress), September 2020.

[I-D.kumar-rtgwg-grpc-protocol]

Kumar, A., Kolhe, J., Ghemawat, S., and L. Ryan, "gRPC Protocol", [draft-kumar-rtgwg-grpc-protocol-00](#) (work in progress), July 2016.

[I-D.openconfig-rtgwg-gnmi-spec]

Shakir, R., Shaikh, A., Borman, P., Hines, M., Lebsack, C., and C. Morrow, "gRPC Network Management Interface (gNMI)", [draft-openconfig-rtgwg-gnmi-spec-01](#) (work in progress), March 2018.

[I-D.pedro-nmrg-anticipated-adaptation]

Martinez-Julia, P., "Exploiting External Event Detectors to Anticipate Resource Requirements for the Elastic Adaptation of SDN/NFV Systems", [draft-pedro-nmrg-anticipated-adaptation-02](#) (work in progress), June 2018.

[I-D.song-ippm-postcard-based-telemetry]

Song, H., Zhou, T., Li, Z., Mirsky, G., Shin, J., and K. Lee, "Postcard-based On-Path Flow Data Telemetry using Packet Marking", [draft-song-ippm-postcard-based-telemetry-08](#) (work in progress), October 2020.

[I-D.song-opsawg-dnp4iq]

Song, H. and J. Gong, "Requirements for Interactive Query with Dynamic Network Probes", [draft-song-opsawg-dnp4iq-01](#) (work in progress), June 2017.

[I-D.song-opsawg-ifu-framework]

Song, H., Qin, F., Chen, H., Jin, J., and J. Shin, "In-situ Flow Information Telemetry", [draft-song-opsawg-ifu-framework-13](#) (work in progress), October 2020.

[I-D.wwx-netmod-event-yang]

WU, Q., Bryskin, I., Birkholz, H., Liu, X., and B. Claise,
"A YANG Data model for ECA Policy Management", [draft-wwx-netmod-event-yang-10](#) (work in progress), November 2020.

[RFC1157] Case, J., Fedor, M., Schoffstall, M., and J. Davin,
"Simple Network Management Protocol (SNMP)", [RFC 1157](#),
DOI 10.17487/RFC1157, May 1990,
<<https://www.rfc-editor.org/info/rfc1157>>.

[RFC2578] McCloghrie, K., Ed., Perkins, D., Ed., and J.
Schoenwaelder, Ed., "Structure of Management Information
Version 2 (SMIv2)", STD 58, [RFC 2578](#),
DOI 10.17487/RFC2578, April 1999,
<<https://www.rfc-editor.org/info/rfc2578>>.

[RFC2981] Kavasseri, R., Ed., "Event MIB", [RFC 2981](#),
DOI 10.17487/RFC2981, October 2000,
<<https://www.rfc-editor.org/info/rfc2981>>.

[RFC3416] Presuhn, R., Ed., "Version 2 of the Protocol Operations
for the Simple Network Management Protocol (SNMP)",
STD 62, [RFC 3416](#), DOI 10.17487/RFC3416, December 2002,
<<https://www.rfc-editor.org/info/rfc3416>>.

[RFC3594] Duffy, P., "PacketCable Security Ticket Control Sub-Option
for the DHCP CableLabs Client Configuration (CCC) Option",
[RFC 3594](#), DOI 10.17487/RFC3594, September 2003,
<<https://www.rfc-editor.org/info/rfc3594>>.

[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>>.

[RFC4656] Shalunov, S., Teitelbaum, B., Karp, A., Boote, J., and M.
Zekauskas, "A One-way Active Measurement Protocol
(OWAMP)", [RFC 4656](#), DOI 10.17487/RFC4656, September 2006,
<<https://www.rfc-editor.org/info/rfc4656>>.

[RFC5357] Hedayat, K., Krzanowski, R., Morton, A., Yum, K., and J.
Babiarz, "A Two-Way Active Measurement Protocol (TWAMP)",
[RFC 5357](#), DOI 10.17487/RFC5357, October 2008,
<<https://www.rfc-editor.org/info/rfc5357>>.

[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/info/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/info/rfc6241>>.
- [RFC6812] Chiba, M., Clemm, A., Medley, S., Salowey, J., Thombare, S., and E. Yedavalli, "Cisco Service-Level Assurance Protocol", [RFC 6812](#), DOI 10.17487/RFC6812, January 2013, <<https://www.rfc-editor.org/info/rfc6812>>.
- [RFC7011] Claise, B., Ed., Trammell, B., Ed., and P. Aitken, "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information", STD 77, [RFC 7011](#), DOI 10.17487/RFC7011, September 2013, <<https://www.rfc-editor.org/info/rfc7011>>.
- [RFC7276] Mizrahi, T., Sprecher, N., Bellagamba, E., and Y. Weingarten, "An Overview of Operations, Administration, and Maintenance (OAM) Tools", [RFC 7276](#), DOI 10.17487/RFC7276, June 2014, <<https://www.rfc-editor.org/info/rfc7276>>.
- [RFC7540] Belshe, M., Peon, R., and M. Thomson, Ed., "Hypertext Transfer Protocol Version 2 (HTTP/2)", [RFC 7540](#), DOI 10.17487/RFC7540, May 2015, <<https://www.rfc-editor.org/info/rfc7540>>.
- [RFC7575] Behringer, M., Pritikin, M., Bjarnason, S., Clemm, A., Carpenter, B., Jiang, S., and L. Ciavaglia, "Autonomic Networking: Definitions and Design Goals", [RFC 7575](#), DOI 10.17487/RFC7575, June 2015, <<https://www.rfc-editor.org/info/rfc7575>>.
- [RFC7799] Morton, A., "Active and Passive Metrics and Methods (with Hybrid Types In-Between)", [RFC 7799](#), DOI 10.17487/RFC7799, May 2016, <<https://www.rfc-editor.org/info/rfc7799>>.
- [RFC7854] Scudder, J., Ed., Fernando, R., and S. Stuart, "BGP Monitoring Protocol (BMP)", [RFC 7854](#), DOI 10.17487/RFC7854, June 2016, <<https://www.rfc-editor.org/info/rfc7854>>.
- [RFC8321] Fioccola, G., Ed., Capello, A., Cociglio, M., Castaldelli, L., Chen, M., Zheng, L., Mirsky, G., and T. Mizrahi, "Alternate-Marking Method for Passive and Hybrid Performance Monitoring", [RFC 8321](#), DOI 10.17487/RFC8321, January 2018, <<https://www.rfc-editor.org/info/rfc8321>>.

[RFC8639] Voit, E., Clemm, A., Gonzalez Prieto, A., Nilsen-Nygaard, E., and A. Tripathy, "Subscription to YANG Notifications", [RFC 8639](#), DOI 10.17487/RFC8639, September 2019, <<https://www.rfc-editor.org/info/rfc8639>>.

[RFC8641] Clemm, A. and E. Voit, "Subscription to YANG Notifications for Datastore Updates", [RFC 8641](#), DOI 10.17487/RFC8641, September 2019, <<https://www.rfc-editor.org/info/rfc8641>>.

[Appendix A](#). A Survey on Existing Network Telemetry Techniques

In this non-normative appendix, we provide an overview of some existing techniques and standard proposals for each network telemetry module.

[A.1](#). Management Plane Telemetry

[A.1.1](#). Push Extensions for NETCONF

NETCONF [[RFC6241](#)] is one popular network management protocol, which is also recommended by IETF. Although it can be used for data collection, NETCONF is good at configurations. YANG Push [[RFC8641](#)] [[RFC8639](#)] extends NETCONF and enables subscriber applications to request a continuous, customized stream of updates from a YANG datastore. Providing such visibility into changes made upon YANG configuration and operational objects enables new capabilities based on the remote mirroring of configuration and operational state. Moreover, distributed data collection mechanism [[I-D.ietf-netconf-distributed-notif](#)] via UDP based publication channel [[I-D.ietf-netconf-udp-notif](#)] provides enhanced efficiency for the NETCONF based telemetry.

[A.1.2](#). gRPC Network Management Interface

gRPC Network Management Interface (gNMI) [[I-D.openconfig-rtgwg-gnmi-spec](#)] is a network management protocol based on the gRPC [[I-D.kumar-rtgwg-grpc-protocol](#)] RPC (Remote Procedure Call) framework. With a single gRPC service definition, both configuration and telemetry can be covered. gRPC is an HTTP/2 [[RFC7540](#)] based open source micro service communication framework. It provides a number of capabilities which are well-suited for network telemetry, including:

- o Full-duplex streaming transport model combined with a binary encoding mechanism provided further improved telemetry efficiency.
- o gRPC provides higher-level features consistency across platforms that common HTTP/2 libraries typically do not. This

characteristic is especially valuable for the fact that telemetry data collectors normally reside on a large variety of platforms.

- o The built-in load-balancing and failover mechanism.

A.2. Control Plane Telemetry

A.2.1. BGP Monitoring Protocol

BGP Monitoring Protocol (BMP) [[RFC7854](#)] is used to monitor BGP sessions and intended to provide a convenient interface for obtaining route views.

The BGP routing information is collected from the monitored device(s) to the BMP monitoring station by setting up the BMP TCP session. The BGP peers are monitored by the BMP Peer Up and Peer Down Notifications. The BGP routes (including Adjacency_RIB_In [[RFC7854](#)], Adjacency_RIB_out [[I-D.ietf-grow-bmp-adj-rib-out](#)], and Local_Rib [[I-D.ietf-grow-bmp-local-rib](#)] are encapsulated in the BMP Route Monitoring Message and the BMP Route Mirroring Message, in the form of both initial table dump and real-time route update. In addition, BGP statistics are reported through the BMP Stats Report Message, which could be either timer triggered or event-driven. More BMP extensions can be explored to enrich the applications of BGP monitoring.

A.3. Data Plane Telemetry

A.3.1. The Alternate Marking (AM) technology

The Alternate Marking method is efficient to perform packet loss, delay, and jitter measurements both in an IP and Overlay Networks, as presented in [[RFC8321](#)] and [[I-D.fioccola-ippm-multipoint-alt-mark](#)].

This technique can be applied to point-to-point and multipoint-to-multipoint flows. Alternate Marking creates batches of packets by alternating the value of 1 bit (or a label) of the packet header. These batches of packets are unambiguously recognized over the network and the comparison of packet counters for each batch allows the packet loss calculation. The same idea can be applied to delay measurement by selecting ad hoc packets with a marking bit dedicated for delay measurements.

Alternate Marking method needs two counters each marking period for each flow under monitor. For instance, by considering n measurement points and m monitored flows, the order of magnitude of the packet counters for each time interval is $n*m*2$ (1 per color).

Since networks offer rich sets of network performance measurement data (e.g packet counters), traditional approaches run into limitations. One reason is the fact that the bottleneck is the generation and export of the data and the amount of data that can be reasonably collected from the network. In addition, management tasks related to determining and configuring which data to generate lead to significant deployment challenges.

Multipoint Alternate Marking approach, described in [\[I-D.fioccola-ippm-multipoint-alt-mark\]](#), aims to resolve this issue and makes the performance monitoring more flexible in case a detailed analysis is not needed.

An application orchestrates network performance measurements tasks across the network to allow an optimized monitoring and it can calibrate how deep can be obtained monitoring data from the network by configuring measurement points roughly or meticulously.

Using Alternate Marking, it is possible to monitor a Multipoint Network without examining in depth by using the Network Clustering (subnetworks that are portions of the entire network that preserve the same property of the entire network, called clusters). So in case there is packet loss or the delay is too high the filtering criteria could be specified more in order to perform a detailed analysis by using a different combination of clusters up to a per-flow measurement as described in Alternate-Marking (AM) [\[RFC8321\]](#).

In summary, an application can configure end-to-end network monitoring. If the network does not experiment issues, this approximate monitoring is good enough and is very cheap in terms of network resources. However, in case of problems, the application becomes aware of the issues from this approximate monitoring and, in order to localize the portion of the network that has issues, configures the measurement points more exhaustively. So a new detailed monitoring is performed. After the detection and resolution of the problem the initial approximate monitoring can be used again.

[A.3.2.](#) Dynamic Network Probe

Hardware-based Dynamic Network Probe (DNP) [\[I-D.song-opsawg-dnp4iq\]](#) provides a programmable means to customize the data that an application collects from the data plane. A direct benefit of DNP is the reduction of the exported data. A full DNP solution covers several components including data source, data subscription, and data generation. The data subscription needs to define the complex data which can be composed and derived from the raw data sources. The data generation takes advantage of the moderate in-network computing to produce the desired data.

While DNP can introduce unforeseeable flexibility to the data plane telemetry, it also faces some challenges. It requires a flexible data plane that can be dynamically reprogrammed at run-time. The programming API is yet to be defined.

A.3.3. IP Flow Information Export (IPFIX) protocol

Traffic on a network can be seen as a set of flows passing through network elements. IP Flow Information Export (IPFIX) [[RFC7011](#)] provides a means of transmitting traffic flow information for administrative or other purposes. A typical IPFIX enabled system includes a pool of Metering Processes collects data packets at one or more Observation Points, optionally filters them and aggregates information about these packets. An Exporter then gathers each of the Observation Points together into an Observation Domain and sends this information via the IPFIX protocol to a Collector.

A.3.4. In-Situ OAM

Traditional passive and active monitoring and measurement techniques are either inaccurate or resource-consuming. It is preferable to directly acquire data associated with a flow's packets when the packets pass through a network. In-situ OAM (iOAM) [[I-D.ietf-ippm-ioam-data](#)], a data generation technique, embeds a new instruction header to user packets and the instruction directs the network nodes to add the requested data to the packets. Thus, at the path end, the packet's experience gained on the entire forwarding path can be collected. Such firsthand data is invaluable to many network OAM applications.

However, iOAM also faces some challenges. The issues on performance impact, security, scalability and overhead limits, encapsulation difficulties in some protocols, and cross-domain deployment need to be addressed.

A.3.5. Postcard Based Telemetry

PBT [[I-D.song-ippm-postcard-based-telemetry](#)] is an alternative to IOAM. PBT directly exports data at each node through an independent packet. PBT solves several issues of IOAM. It can also help to identify packet drop location in case a packet is dropped on its forwarding path.

A.4. External Data and Event Telemetry

A.4.1.1. Sources of External Events

To ensure that the information provided by external event detectors and used by the network management solutions is meaningful for the management purposes, the network telemetry framework must ensure that such detectors (sources) are easily connected to the management solutions (sinks). This requires the specification of a simple taxonomy of detectors and match it to the connectors and/or interfaces required to connect them.

Once detectors are classified in such taxonomy, their definitions are enlarged with the qualities and other aspects used to handle them and represented in the ontology and information model (e.g. YANG). Therefore, differentiating several types of detectors as potential sources of external events is essential for the integrity of the management framework. We thus differentiate the following source types of external events:

- o Smart objects and sensors. With the consolidation of the Internet of Things~(IoT) any network system will have many smart objects attached to its physical surroundings and logical operation environments. Most of these objects will be essentially based on sensors of many kinds (e.g. temperature, humidity, presence) and the information they provide can be very useful for the management of the network, even when they are not specifically deployed for such purpose. Elements of this source type will usually provide a specific protocol for interaction, especially one of those protocols related to IoT, such as the Constrained Application Protocol (CoAP). It will be used by the telemetry framework to interact with the relevant objects.
- o Online news reporters. Several online news services have the ability to provide enormous quantity of information about different events occurring in the world. Some of those events can impact on the network system managed by a specific framework and, therefore, it will be interested on getting such information. For instance, diverse security reports, such as the Common Vulnerabilities and Exposures (CVE), can be issued by the corresponding authority and used by the management solution to update the managed system if needed. Instead of a specific protocol and data format, the sources of this kind of information usually follow a relaxed but structured format. This format will be part of both the ontology and information model of the telemetry framework.
- o Global event analyzers. The advance of Big Data analyzers provides a huge amount of information and, more interestingly, the identification of events detected by analyzing many data streams

from different origins. In contrast with the other types of sources, which are focused in specific events, the detectors of this source type will detect very generic events. For example, a sports event takes place and some unexpected movement makes it highly interesting and many people connects to sites that are covering such event. The systems supporting the services that cover the event can be affected by such situation so their management solutions should be aware of it. In contrast with the other source types, a new information model, format, and reporting protocol is required to integrate the detectors of this type with the management solution.

Additional types of detector types can be added to the system but they will be generally the result of composing the properties offered by these main classes. In any case, future revisions of the network telemetry framework will include the required types that cover new circumstances and that cannot be obtained by composition.

A.4.2. Connectors and Interfaces

For allowing external event detectors to be properly integrated with other management solutions, both elements must expose interfaces and protocols that are subject to their particular objective. Since external event detectors will be focused on providing their information to their main consumers, which generally will not be limited to the network management solutions, the framework must include the definition of the required connectors for ensuring the interconnection between detectors (sources) and their consumers within the management systems (sinks) are effective.

In some situations, the interconnection between the external event detectors and the management system is via the management plane. For those situations there will be a special connector that provides the typical interfaces found in most other elements connected to the management plane. For instance, the interfaces will accomplish with a specific information model (YANG) and specific telemetry protocol, such as NETCONF, SNMP, or gRPC.

Authors' Addresses

Haoyu Song
Futurewei
2330 Central Expressway
Santa Clara
USA

Email: hsong@futurewei.com

Fengwei Qin
China Mobile
No. 32 Xuanwumenxi Ave., Xicheng District
Beijing, 100032
P.R. China

Email: qinfengwei@chinamobile.com

Pedro Martinez-Julia
NICT
4-2-1, Nukui-Kitamachi
Koganei, Tokyo 184-8795
Japan

Email: pedro@nict.go.jp

Laurent Ciavaglia
Nokia
Villarceaux 91460
France

Email: laurent.ciavaglia@nokia.com

Aijun Wang
China Telecom
Beiqijia Town, Changping District
Beijing, 102209
P.R. China

Email: wangaj.bri@chinatelecom.cn

