

Time-Variant Routing  
Internet-Draft  
Intended status: Informational  
Expires: 13 August 2026

L. Zhang  
J. Dong  
Huawei  
M. Boucadair  
Orange  
9 February 2026

Applicability of TVR YANG Data Models  
draft-zdm-tvr-applicability-05

## Abstract

Time-Variant Routing (TVR) is a routing system that can accommodate predicted topology changes caused by internal or external factors. Typical use cases include resource preservation networks, operating efficiency networks and dynamic reachability networks. This document provides examples of how to implement the TVR scheduling capabilities for key use cases. It describes which part of the TVR data model is used and why, and it outlines operational and security considerations when deploying TVR-based technologies.

## Discussion Venues

This note is to be removed before publishing as an RFC.

Discussion of this document takes place on the Time-Variant Routing Working Group mailing list ([tvr@ietf.org](mailto:tvr@ietf.org)), which is archived at <https://mailarchive.ietf.org/arch/browse/tvr/>.

Source for this draft and an issue tracker can be found at <https://github.com/zhangli-abcd/TVR-Applicability-2>.

## 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 13 August 2026.

## 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 . . . . .	3
2. Conventions and Definitions . . . . .	3
3. Applicability of the TVR YANG Model . . . . .	4
3.1. Applicability of TVR Node YANG Module . . . . .	4
3.2. Applicability of TVR Topology YANG Module . . . . .	5
3.2.1. Interactions in Centralized Scenario . . . . .	6
3.2.2. Interactions in Distributed Scenario . . . . .	6
3.3. Encoding of the YANG Model . . . . .	7
3.4. Management Protocols for TVR . . . . .	9
4. Time Synchronization . . . . .	11
4.1. Hardware-based Time Synchronization Mechanisms . . . . .	12
4.2. Software-based Time Synchronization Protocols . . . . .	12
4.2.1. NTP . . . . .	12
4.2.2. SNTP . . . . .	13
5. Schedule Database . . . . .	13
5.1. Data Structure . . . . .	14
5.2. Schedule Operations . . . . .	15
6. Operational Considerations . . . . .	15
6.1. Schedule Dissemination . . . . .	16
6.2. Schedule Execution . . . . .	16
6.2.1. Execution of Node Schedule . . . . .	16
6.2.2. Execution of Topology Module Schedule . . . . .	17
6.3. Schedule Recovery . . . . .	18
6.4. Error Handling . . . . .	18
6.4.1. Consistency Error . . . . .	18
7. Security Considerations . . . . .	18
7.1. Denial-of-Service (DoS) Attack . . . . .	19
7.2. Traffic Analysis and Path Prediction . . . . .	19
7.3. Activity Identification and Privacy . . . . .	20
7.4. Spoofing and Manipulation of Time Information . . . . .	20

7.5. Replay Attacks on Time-Sensitive Data . . . . .	20
7.6. Compromised Time Sources . . . . .	20
7.7. Schedule Tampering and Malicious Schedule Injection . . . . .	21
8. IANA Considerations . . . . .	21
9. References . . . . .	21
9.1. Normative References . . . . .	21
9.2. Informative References . . . . .	22
Acknowledgments . . . . .	23
Appendix A: Code Examples . . . . .	24
Code Examples for "Energy-harvesting Wireless Sensor Network" . . . . .	24
Code Examples for "Cellular Network" . . . . .	27
Code Examples for "Tidal Network" . . . . .	28
Code Examples for "Mobile Satellites" . . . . .	33
Code examples for "Predictable Moving Vessels" . . . . .	37
Contributors . . . . .	38
Authors' Addresses . . . . .	39

## 1. Introduction

The Time-Variant Routing (TVR) Working Group addresses a need in network environments where predictable variations in topology - such as the restoration, activation, or loss of network elements, are part of normal operations. This approach is essential in dynamic networks with mobile nodes, where links may be frequently disrupted and re-established due to mobility. It is also essential in networks with highly predictable traffic patterns, where links may be powered down to conserve or reduce energy.

This document provides examples of implementing TVR scheduling capabilities in identified use cases. It demonstrates the applicability of the TVR data model, methods for disseminating the TVR schedules, and the necessary IETF ancillary technologies for network environments, such as time synchronization and policy, that support TVR capabilities. The examples assume YANG instance data encoding per [RFC7951] for JSON and the TVR schedule YANG modules.

## 2. Conventions and Definitions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

This document uses the following terms:

Managing Device: A centralized entity responsible for generating and

maintaining TVR schedules. The managing device distributes schedules to network controllers and/or managed devices using the TVR YANG modules. In some deployments, the managing device may also serve as the network controller.

**Network Controller:** An entity that receives topology schedules from the managing device and performs route computation based on time-variant network conditions. The controller then distributes routing results to managed devices.

- \* **Managed Device:** A network device (e.g., router, switch) that receives schedules and/or routing instructions, and executes them according to the specified time windows. Managed devices may receive schedules directly from the managing device or routing results from the network controller.

### 3. Applicability of the TVR YANG Model

The TVR data model [I-D.ietf-tvr-schedule-yang] defines the TVR node YANG module and TVR topology YANG module. This clause discusses the applicability of these two modules separately.

#### 3.1. Applicability of TVR Node YANG Module

As specified in Section 5.2 of [I-D.ietf-tvr-schedule-yang], module `ietf-tvr-node.yang` is a device model and designed to manage a single node with scheduled attributes. It is not necessary in all TVR use cases.

The applicability of TVR node YANG module depends on whether changes in the attributes of network devices are caused by the environment or centrally controlled.

- \* When the changes are caused by the environment changes (such as movement, sunlight changes, and weather changes) or by decisions made by the devices themselves, the network device does not need to get the managed information through the YANG module. For example, when the distance between two nodes' is too far to establish a connection, then the link is down. Another case is that when the weather is not good and leading to the link degradation, then the nodes decide to disconnect the link.
- \* When the changes are caused by the centralized control (such as a controller, or an orchestrator), the devices themselves do not know when to adjust the attributes. In this case, the scheduled attributes changes should be delivered to the network devices through TVR node YANG module.

### 3.2. Applicability of TVR Topology YANG Module

As specified in Section 5.3 of [I-D.ietf-tvr-schedule-yang], module `ietf-tvr-topology.yang` describes a network topology with a time-variant availability schedule. This YANG module is also not applicable for all TVR use cases.

According to the description of Section 3.1 of [I-D.ietf-tvr-requirements], the scheduling generation locality and execution locality may be centralized or distributed.

- \* When the schedules are generated and executed in distributed manner, which means that each node generates and executes its specific schedules. In this scenario, the topology YANG module is not necessary, the devices can collect topology schedules by other means. This scenario is outside of the scope of this document.
- \* When the schedules are generated and executed in centralized manner and within the same device, the topology YANG module is also not applicable. Therefore, this scenario is also outside of the scope of this document.
- \* When the schedules are generated and executed in centralized manner but on different devices. For example, the schedule is generated by the managing device, and executed by the network controller. In this scenario the scheduled topology changes need to be sent to the execution device through the topology YANG module. This scenario is called "Centralized Scenario".
- \* When the schedules are generated in a centralized manner and executed in a distributed manner, the YANG module needs to be used to deliver the scheduled topology changes to the managed device. This scenario is called "Distributed Scenario".

To summarize the key differences between these scenarios:

- \* Centralized Scenario: Schedules are generated by the managing device, stored in the network controller, and executed by network devices. Route computation runs in the controller, and routing results are pushed to the network devices.
- \* Distributed Scenario: Schedules are generated by the managing device, stored in and executed by the network devices. Route computation runs in the network devices themselves.

### 3.2.1. Interactions in Centralized Scenario

In the centralized scenario, the network managing device generates and maintains schedules, the routing application is deployed in the network controller, and the network devices execute the schedules and routing results. The following figure shows the components of the centralized scenario.




Figure 1: Components of Centralized Scenario

A centralized scenario involves the interaction between the managing device and network controller, the managing device and network devices, and the controller and network devices.

The managing device may need to deliver node-specific schedules to network devices by TVR Schedule Node YANG module Section 5.2 of [I-D.ietf-tvr-schedule-yang]. This is optional and only necessary when the node attribute changes are instructed by the controller. Meanwhile, the network devices need to report their own status data to the managing device.

The managing device needs to deliver the schedules of network topology to the network controller by the TVR Network Topology YANG module Section 5.3 of [I-D.ietf-tvr-schedule-yang], so that the routing application in the controller can consider the impact of topology changes on routes when calculating routes.

The network controller should deliver the route calculation results to the network devices. The format of the routing results depend on the protocols deployed (BGP, PCEP, etc.). Routing results MAY be delivered ahead and activated using an explicit time reference (e.g., absolute UTC), or via pre-staging plus an activation trigger.

### 3.2.2. Interactions in Distributed Scenario

In the distributed scenario, the managing device generates and maintains schedules, the routing application is deployed in the network devices which also executes schedules and route calculation.




Figure 2: Components of Distributed Scenario

The distributed model only involves the interaction between managing devices and network devices(managed devices).

The managing device need to deliver network topology schedules to all the network devices by TVR Network Topology YANG module for route calculation. The managing device may also need to deliver node-specific schedules to network devices by TVR Schedule Node YANG module, this is optional and only necessary when the node attributes changes are instructed by the controller. The network devices need to report their own status data to the managing device.

### 3.3. Encoding of the YANG Model

The TVR data model [I-D.ietf-tvr-schedule-yang] can manage network resources and topologies with scheduled attributes. There are modules defined in the TVR data model, these are:

- \* The "ietf-tvr-schedule" module contains the schedule YANG definitions. This module uses groupings from [I-D.ietf-netmod-schedule-yang] data model;
- \* The "ietf-tvr-topology" module defines a network topology with a time-variant availability schedule;
- \* The "ietf-tvr-node" module is to be used to manage the scheduled attributes of a single node.

To create a schedule, the following TVR data model objects and subsequent branches are used:

- \* 'node-schedule'
- \* 'interface-schedule'
- \* 'attribute-schedule'

When using these YANG modules with NETCONF or RESTCONF, implementations SHOULD target the intended configuration datastore for schedule provisioning and MAY read from the operational datastore to retrieve execution status and applied schedules. Clients can use NMDA (Network Management Datastore Architecture, [RFC8342]) operations to distinguish between intended configuration and actual operational state. For example, the managing device writes schedules to the intended or running datastore, and network devices report execution status via the operational datastore.

A TVR scenario example is provided below, where a wireless link is shut down for 12 hours, from 19:00 to 7am the next day. The schedule is identified using a unique identifier that is conveyed in 'schedule-id', and the recurring schedule can be applied for multiple days using Coordinated Universal Time (UTC). More detailed examples of the JSON example is provided in this documents Appendix.



```

{
  "ietf-tvr-node:node-schedule":[
    {
      "node-id":1234567890,
      "node-power-schedule":{
        "power-default":true,
      },
      "interface-schedule":[
        {
          "name":"Wlan0",
          "default-available":false,
          "attribute-schedule":{
            "schedules":[
              {
                "schedule-id":111111,
                "recurrence-first":{
                  "start-time-utc":"2025-12-01T19:00:00Z",
                  "duration":43200
                },
                "utc-until":"2026-12-01T00:00:00Z",
                "frequency":"ietf-schedule:daily",
                "interval":1,
                "attr-value":{
                  "available":true
                }
              }
            ]
          }
        }
      ]
    }
  ]
}

```

The methods for disseminating and propagating the generated schedules are discussed in the following subsections.

#### 3.4. Management Protocols for TVR

The TVR data model is designed to be accessed via YANG-based management protocols such as NETCONF [RFC6241], RESTCONF [RFC8040] and CORECONF[I-D.ietf-core-comi-19]. This section discusses the applicability of these protocols for configuring time-variant network resources using the TVR YANG data models.

NETCONF provides a robust mechanism for managing complex network configurations, particularly when transactional integrity and operational consistency are required. Its ability to execute atomic

transactions ensures that schedules involving multiple resources are applied fully, preventing partial updates that could lead to configuration inconsistencies. This feature is important for time-sensitive scheduling in TVR environments. Additionally, NETCONF supports the validation of configurations prior to commitment, allowing operators to verify the correctness of schedules before they are applied. It also includes rollback capabilities, such as restoring a previous configuration during scheduling errors.

In contrast, RESTCONF offers a simpler, stateless method for interacting with network devices, making it suitable for use cases requiring lightweight, rapid configuration. RESTCONF utilizes a RESTful interface over HTTP, providing a streamlined approach to network configuration and management. Therefore, RESTCONF may be advantageous in scenarios where quick adjustments to schedules are needed or where integration with web-based or cloud-native systems is a priority.

CORECONF provides a lightweight, stateless method for managing small network devices where saving bytes to transport a message is very important. CORECONF uses CoAP[RFC7252] methods to access structured data defined in YANG which is a complementary to RESTCONF. Unlike RESTCONF, many CORECONF design decisions are motivated by minimizing the message size. Therefore, CORECONF is advantageous in networks with constrained devices and very limited transmission bandwidth, especially in IoT devices that already deployed CoAP.

Depending on the type of node in the TVR network, NETCONF would be the preferred protocol for large-scale, critical scheduling operations requiring validation and rollback mechanisms. For smaller-scale or isolated scheduling tasks, RESTCONF provides an efficient and straightforward option without the need for the transactional features offered by NETCONF. CORECONF is preferred in resource constrained IoT networks where saving message bytes is a priority. The choice of protocol to use with the TVR YANG model should be driven by the specific requirements of the network environment and the complexity of the scheduling tasks involved.

The security aspects of these management protocols, including their strengths and weaknesses, are discussed further in Section 7 of this document.

#### 4. Time Synchronization

Time Synchronization is fundamental for ensuring that TVR mechanisms, which depend on highly accurate timing, function as intended across an entire network. Misalignment in time could lead to serious routing issues, including inefficiency in path forwarding, instability in routing tables, and traffic outages.

Time Synchronization mechanisms will be used to ensure:

- \* Coordination of Planned Network Events;
- \* Verification of TVR Data Model Time Stamps
- \* Accurate Scheduling of Paths;
- \* Fault Tolerance.

Different time-variant scenarios may require different granularities of time synchronization. For example, the period of traffic and topology changes in tidal networks is usually a day or week. Therefore, a second-level time synchronization is enough. However, for the dynamic reachability scenarios, a fine-granularity time synchronization may be necessary, as the nodes may moving very fast in some cases (the moving speed of a low earth orbit satellite is more than 7900 m/s)

Operators SHOULD derive a maximum acceptable time-error bound based on the schedule granularity, execution jitter tolerance, and activation window requirements. For instance, if a schedule has a 1-second activation window and the system can tolerate up to 100ms of execution jitter, the time synchronization error MUST be kept well below 900ms. The chosen time synchronization protocol and configuration MUST be capable of meeting this derived bound under all expected network conditions.

Existing clock synchronization protocols can be classified into hardware-based protocols and software-based protocols.

Hardware-based protocols often rely on dedicated hardware to ensure clock synchronization, such as Satellite Based Timing Service (SBTS) and Precision Time Protocol (PTP). The SBTS includes but not limited to Global Position System (GPS), BeiDou Navigation Satellite System(BDS), Global Navigation Satellite System(GLONASS), and Galileo satellite navigation system. Software-based protocols, on the other hand, synchronize clocks through software packages running on systems, such as Network Time Protocol (NTP) [RFC5905] and Simple Network Time Protocol (SNTP) [RFC4330].

The security aspects of time synchronization mechanisms are discussed further in Section 7 of this document.

#### 4.1. Hardware-based Time Synchronization Mechanisms

Hardware-based protocols typically have higher precision and stability, but also have higher cost due to the dedicated hardware. SBTS and PTP are the typical hardware-based time synchronization mechanisms.

SBTS provides a precise time synchronization service based on the signals transmitted by the satellites. SBTS can realize the micro-second level time synchronization among the devices installed with SBTS reviver hardware.

PTP is a network protocol that complies with the IEEE 1588 standard and is used to implement high-precision time synchronization between network nodes. PTP implements time synchronization by transmitting synchronization messages between master and slave devices. Based on the hardware timestamp, the precision of time synchronization is much higher than that of NTP, and can reach the sub-microsecond level or even tens of nanoseconds. When deploying PTP in TVR networks, the managing devices should be the master and the network devices and controller should be the slaves which get time from the master.

Both SBTS and PTP can realize micro-second level time synchronization. Depending on the features of TVR network, the SBTS would be the preferred mechanisms for large-scale, high dynamic and open-air networks, especially networks with unreliable links as it does not require network links to exchange time information. For the small-scale networks with stable links but have high-precision time synchronization requirements, the PTP is much preferred.

#### 4.2. Software-based Time Synchronization Protocols

Software-based protocols are simple and applicable to common hardware devices, but have lower precision (For example, the NTP can realize the synchronization at tens of milliseconds level).

##### 4.2.1. NTP

NTP uses a hierarchical structure of time sources. Each level of this hierarchy is termed a stratum. Generally, an NTP server synchronized to an authoritative clock runs at stratum 1. Such NTP server functions as the primary time server to provide clock synchronization for other devices on the network. Stratum 2 servers obtain time from stratum 1 servers, stratum 3 servers obtain time from stratum 2 servers, and so on.

In TVR use cases, the managing device functions as a level-1 NTP server and synchronized to an authoritative clock source. The network controller and network devices behave as clients to obtain accurate time from the managing device. Figure 3 shows an NTP deployment scenario for obtaining clock from a GPS clock source.




Figure 3: Deployment Case of NTP in Tidal Networks

NTP is preferred in large-scale networks with reliable links and long-term changes, which does not require a high-precision time synchronization.

#### 4.2.2. SNTP

SNTP is a subset of the NTP used to synchronize computer clocks in the Internet. It simplifies the complex NTP synchronization function and has lower clock precision, but the synchronization precision still can be kept within seconds. SNTP is also preferred in large-scale networks with reliable links, long-term changes, and loose synchronization precision. In addition, it is more suitable for networks with limited resources than NTP.

### 5. Schedule Database

The schedule database is used to store and maintain schedules, the database may be deployed on a managing device and managed devices based on requirements.

The source of the schedule database may be created from multiple sources, for example, configuration from an administrator or YANG model from the management interface. The schedule entries of different databases may be different, but the content of the same schedule entry in the schedule databases of different devices in the same domain must be consistent. There are at least two ways to make the content of the same schedule entry in different schedule databases consistent:

- \* All the schedule entries are generated at a specific device;
- \* Schedule entries are generated at different devices, but there is a synchronization mechanism to synchronize the schedule databases among devices.

Option 1 is simplest and easy to implement. In a time-variant domain, the managing device may receive scheduling requests and generate all schedule entries. Then the schedule entries are delivered to the necessary network devices in the domain through the TVR YANG model.

Option 2 relies on advertisement mechanisms (such as routing techniques) to advertise the scheduling data generated by itself to other devices. This could be achieved using extensions to existing routing schemes and techniques.

Detailed schedule distribution mechanisms beyond YANG-based configuration are outside of this document.

### 5.1. Data Structure

[I-D.ietf-tvr-schedule-yang] defines a TVR Node and TVR Topology YANG modules. The Node YANG module includes node power schedules and interface schedules. The Topology YANG module includes node schedules and link schedules.

Based on the preceding four schedule types, the schedule database should contain four types of schedule entries in different formats:

- \* Node power schedule entry;
- \* Interface schedule entry;
- \* Node schedule entry;
- \* Link schedule entry;

The detailed format and fields of different types of schedule entries could reference to the definitions of the corresponding YANG modules.

## 5.2. Schedule Operations

This section provides general requirements for using the TVR schedules.

The schedule database should support the add, update, and delete operations.

When adding or updating a schedule entry, the execution node needs to check whether resource conflicts exist between the current schedule and existing schedules. If a conflict exists, the operation should be failed.

Schedules are updated and deleted based on schedule IDs. Therefore, schedule IDs must be unique in a time-variant domain. This can be handled, e.g., by a dedicated allocation agent within the time-variant domain.

## 6. Operational Considerations

Several operational considerations exist when using TVR techniques and data models in a network. This section provides some high-level observations and more detailed sub-sections for specific consideration related to schedule dissemination, execution, and recovery in case of failure to apply a schedule or partial change.

- \* Coordinated Network Events: TVR often coordinates routing changes anticipating events like predictable low-traffic periods or link downtimes (e.g., scheduled maintenance or traffic demand).
- \* Accurate Scheduling of Paths: TVR schedules capable routers and network nodes will dynamically adjust forwarding paths based on planned changes in link availability or network conditions.
- \* Time-Stamped Data Models: TVR will require the use time-stamped data models (e.g., schedules for link changes or availability windows) to make interface management decisions. This ensures that all TVR nodes interpret the timing of events consistently and implement time-based policies correctly.

Therefore, network time accuracy and time-stamped data models are critical to ensure that coordinated network events and scheduled path decisions across the network are based on a consistent time reference. Without accurate time sync, nodes could apply different schedules, causing routing inconsistencies, path flapping, or packet loss.

Since the all the schedules are generated by the managing devices, it is necessary to make sure that the managing device and managed devices within a schedule domain (Section 2.1.1 of [I-D.ietf-tvr-requirements]) are synchronized with the same time source. It ensures they have the same notion of when that is.

### 6.1. Schedule Dissemination

When distributing schedules, the following problems should be considered:

- \* The managed devices that receives a schedule should have the ability to execute the schedule. If the device does not support schedule execution, it SHOULD reject the update with an appropriate error, or ignore it if the management interface does not support error reporting.
- \* The distributing of a schedule SHOULD be delivered at least x (operator-defined) time before the earliest start time of the schedule, this ensures that the managed device has enough time to execute this schedule.

### 6.2. Schedule Execution

Schedule execution means that a component (e.g., device) undertakes an action (e.g., allocates and deallocates resources) at specified time points. The schedule execution of Node Module and Topology Module should be considered separately.

#### 6.2.1. Execution of Node Schedule

Node schedule execution indicates a node to change its node/interfaces availability/power up and down, and other attributes directly or by commands.

when executing a node schedule, the schedule executor should undertake an action at specified time points as indicated in the schedule.



### 6.2.2. Execution of Topology Module Schedule

Topology schedule execution means a node take some measures before or upon the scheduled topology changes.

The schedule executor should understand the consequences of the schedule execution. The addition and deletion of the topology need to be considered separately.

A link coming up or a node joining a topology should not have any functional change until the change is proven to be fully operational. The routing paths may be pre-computed but should not be installed before all the topology changes are confirmed to be operational. Pre-computation MAY be used, operators should weigh compute cost versus reduced activation latency. The network may choose to not do any pre-installation or pre-computation in reaction to topological additions at a small cost of some operational efficiency.

Topological deletions are an entirely different matter. If a link or node is to be removed from the topology, then the network should act before the anticipated change to route traffic around the expected topological change. Specifically, at some point before the planned topology change, the routing paths should be pre-computed and installed before the topology change takes place. The required time to perform such planned action will vary depending on the exact network and configuration. When using an IGP or other distributed routing protocols, the affected links may be set to a high metric to direct traffic to alternate paths. This type of change does require some time to propagate through the network, so the metric change should be initiated far enough in advance that the network converges before the actual topological change.

In addition to the addition and deletion of topology, a schedule may indicate the attributes change of some links, such as bandwidth and delay. If an attribute changes better (such as latency decrease and bandwidth increase), then the executor should take actions later or until the topology change is proven to be fully operational. If an attribute changes worse (such as latency increase and bandwidth decrease), then the node should react to it before the change take place.

### 6.3. Schedule Recovery

Schedule recovery means when a node lost its specific schedule data, it should be able to recover these schedule data. Typical scenarios include data loss due to device restart, disconnection from managing devices and failure to receive new schedules for a long time. In these scenarios, once the connection between the managed device and the managing device is established again, the managing device needs to re-distribute all schedules that start time is later than the current moment to the managed device after time synchronization is finished.

### 6.4. Error Handling

#### 6.4.1. Consistency Error

Consistency error means that some time parameters conflict with other time parameters in the same schedule or in other schedules.

- \* If the time parameters of a schedule conflict with each other, for example, the period-start later than period-end, the duration is longer than the product of frequency and interval, or the duration is longer than utc-until, then the schedule should be discarded and an error should be returned to the schedule originator (e.g., the managing device or management client).
- \* If there is a conflict between schedules with different schedule IDs, for example, schedule1 indicates that interface B is closed at time A, but schedule2 indicates that interface B is open at time A, then only the conflicting update should be rejected (retaining the last-known-good schedule). An error should be returned to the schedule originator, and the conflict MUST be logged for audit purposes. If two schedules have the same schedule ID, then it is considered as an update of the former schedule. Updates with the same schedule ID SHALL completely replace the previous schedule (full replacement, not merge). Implementations SHOULD support versioning or etag-based mechanisms to detect concurrent updates. Partial updates to a schedule are NOT permitted; clients MUST send the complete schedule definition.

## 7. Security Considerations

The integration of time-variant mechanisms in network operations presents distinct security challenges that require thorough analysis to safeguard the network's integrity, availability, and confidentiality. Networks that rely on time-sensitive data for routing and forwarding decisions are particularly susceptible to attacks that exploit timing dependencies.

The "Security Considerations" section of [I-D.ietf-tvr-requirements] outlines various threat vectors and categories specific to time-variant environments.

#### 7.1. Denial-of-Service (DoS) Attack

In a time-variant network, malicious actors could attack the network by disrupting or manipulating the time synchronization process. For example, an attacker could intentionally delay or corrupt time signals exchanged within the network, leading to routing errors and widespread denial-of-service (DoS) attacks.

This kind of attack could be mitigated by the redundant time synchronization mechanisms, for example, multiple NTP sources or multiple time synchronization protocols could be deployed in a TVR network. The network devices could guarantee the correctness of the time by checking whether the time signals from different sources or protocols.

In addition, peer authentication is also an important way to protect the time signals being tampered by attackers. Some security extensions for time synchronization protocols (such as NTS (Network Time Security)) are recommended to be applied.

#### 7.2. Traffic Analysis and Path Prediction

In a time-variant network, if time information is not adequately protected, attackers could conduct traffic analysis to infer routing decisions, network load, or usage patterns. The schedule ability could enable attackers to launch highly targeted attacks, such as selectively overloading certain links or intercepting sensitive communications.

This kind of attack could be mitigated by the encryption of schedules and the authentication of managing devices. For the networks using NETCONF to deliver schedules, NETCONF over TLS[RFC7589] is recommended to achieve the bidirectional authentication and encryption of YANG model data. RESTCONF supports TLS originally, so it can be deployed without additional configurations or modifications.

In addition, in time variant networks with centralized scenarioSection 3.2.1, the encryption of routing path information is also necessary to avoid the fake routing information. Considering the most typical protocols used to deliver the routing path information between controller and network devices are BGP and PCEP, and both are based on TCP. Therefore, the TLS is RECOMMENDED to provide confidentiality and integrity protection.

### 7.3. Activity Identification and Privacy

In certain scenarios, precise time information exchanged within the network could be correlated with specific user or device behavior, inadvertently revealing private information.

This risk could also be mitigated by the solutions mentioned in Section 7.2.

### 7.4. Spoofing and Manipulation of Time Information

In a time-variant network, if an attacker were to inject false or manipulated time data into the network, it could cause routers and devices to make incorrect decisions, potentially leading to traffic misrouting, network partitions, or inefficient use of resources.

This risk could also be mitigated by the solutions mentioned in Section 7.1.

### 7.5. Replay Attacks on Time-Sensitive Data

Time-variant network data and schedules updates may be susceptible to replay attacks, which could cause network devices to act on outdated information, leading to inconsistent routing decisions, misaligned schedules, or security gaps. In particular, attackers could exploit replay attacks to force devices into outdated configurations or interfere with the synchronization of schedules across the network.

This kind of attack could be mitigated by encrypting time signals, schedules and routing path data, and adding a unique number to the encrypted section of a packet. This has been implemented in existing protocols, for example, the NTS supports unique identifier extension field (EF) containing a random number, the TLS supports Message Authentication Code generated from sequence number.

### 7.6. Compromised Time Sources

The reliance on external time sources for synchronization purposes presents a potential attack surface for time-variant networks. If a trusted time source, such as a GPS signal or an NTP server, is compromised, the attacker could feed erroneous time information to the entire network, disrupting its operation.

This kind of attack could be mitigated by the solutions mentioned in Section 7.1.

### 7.7. Schedule Tampering and Malicious Schedule Injection

Unauthorized modification or injection of malicious TVR schedules poses significant operational and security risks. An attacker who successfully tampers with schedules could cause traffic blackholing (by scheduling link or node shutdowns at critical times), trigger costly network-wide reroutes, degrade service-level agreement (SLA) performance, or enable targeted interception of sensitive traffic flows. Such attacks undermine the predictability and reliability that TVR aims to provide.

Mitigating schedule tampering requires a defense-in-depth approach:

- \* **Authentication and Authorization:** All schedule updates **MUST** be authenticated to verify the identity of the originator. Role based access control (RBAC) or attribute-based access control (ABAC) **SHOULD** be enforced to ensure that only authorized entities can modify schedules.
- \* **Integrity Protection:** Schedules **MUST** be protected against tampering in transit and at rest using cryptographic integrity mechanisms (e.g., digital signatures, HMAC). NETCONF over TLS [RFC7589], RESTCONF over TLS, and similar secure transport protocols provide such protection.
- \* **Audit Logging:** All schedule creation, modification, and deletion operations **SHOULD** be logged with timestamps, originator identity, and a description of the change. These logs are essential for forensic analysis and detecting anomalous behavior.
- \* **Rate Limiting and Anomaly Detection:** Implementations **SHOULD** enforce rate limits on schedule update operations and deploy anomaly detection mechanisms to identify suspicious patterns (e.g., rapid schedule churn, schedules from unexpected sources).

## 8. IANA Considerations

This document has no IANA actions.

## 9. References

### 9.1. Normative References

[I-D.ietf-tvr-schedule-yang]

Qu, Y., Lindem, A., Kinzie, E., Fedyk, D., and M. Blanchet, "YANG Data Model for Scheduled Attributes", Work in Progress, Internet-Draft, draft-ietf-tvr-schedule-yang-07, 19 January 2026, <<https://datatracker.ietf.org/doc/html/draft-ietf-tvr-schedule-yang-07>>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/rfc/rfc2119>>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/rfc/rfc8174>>.

[RFC9657] Birrane, III, E., Kuhn, N., Qu, Y., Taylor, R., and L. Zhang, "Time-Variant Routing (TVR) Use Cases", RFC 9657, DOI 10.17487/RFC9657, October 2024, <<https://www.rfc-editor.org/rfc/rfc9657>>.

## 9.2. Informative References

[I-D.ietf-core-comi-19]

Veillette, M., Van der Stok, P., Pelov, A., Bierman, A., and C. Bormann, "CoAP Management Interface (CORECONF)", Work in Progress, Internet-Draft, draft-ietf-core-comi-19, 3 November 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-core-comi-19>>.

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

Ma, Q., Wu, Q., Boucadair, M., and D. King, "A Common YANG Data Model for Scheduling", Work in Progress, Internet-Draft, draft-ietf-netmod-schedule-yang-10, 7 August 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-netmod-schedule-yang-10>>.

[I-D.ietf-tvr-requirements]

King, D., Contreras, L. M., Sipos, B., and L. Zhang, "TVR (Time-Variant Routing) Requirements", Work in Progress, Internet-Draft, draft-ietf-tvr-requirements-07, 10 October 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-tvr-requirements-07>>.

[I-D.lj-rtg-sat-routing-consideration]

Liu, P. and T. Jiang, "Routing in Satellite Networks: Challenges & Considerations", Work in Progress, Internet-

Draft, draft-lj-rtg-sat-routing-consideration-00, 2 March 2025, <<https://datatracker.ietf.org/doc/html/draft-lj-rtg-sat-routing-consideration-00>>.

- [RFC4330] Mills, D., "Simple Network Time Protocol (SNTP) Version 4 for IPv4, IPv6 and OSI", RFC 4330, DOI 10.17487/RFC4330, January 2006, <<https://www.rfc-editor.org/rfc/rfc4330>>.
- [RFC5905] Mills, D., Martin, J., Ed., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", RFC 5905, DOI 10.17487/RFC5905, June 2010, <<https://www.rfc-editor.org/rfc/rfc5905>>.
- [RFC6241] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed., and A. Bierman, Ed., "Network Configuration Protocol (NETCONF)", RFC 6241, DOI 10.17487/RFC6241, June 2011, <<https://www.rfc-editor.org/rfc/rfc6241>>.
- [RFC7252] Shelby, Z., Hartke, K., and C. Bormann, "The Constrained Application Protocol (CoAP)", RFC 7252, DOI 10.17487/RFC7252, June 2014, <<https://www.rfc-editor.org/rfc/rfc7252>>.
- [RFC7589] Badra, M., Luchuk, A., and J. Schoenwaelder, "Using the NETCONF Protocol over Transport Layer Security (TLS) with Mutual X.509 Authentication", RFC 7589, DOI 10.17487/RFC7589, June 2015, <<https://www.rfc-editor.org/rfc/rfc7589>>.
- [RFC7951] Lhotka, L., "JSON Encoding of Data Modeled with YANG", RFC 7951, DOI 10.17487/RFC7951, August 2016, <<https://www.rfc-editor.org/rfc/rfc7951>>.
- [RFC8040] Bierman, A., Bjorklund, M., and K. Watsen, "RESTCONF Protocol", RFC 8040, DOI 10.17487/RFC8040, January 2017, <<https://www.rfc-editor.org/rfc/rfc8040>>.
- [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/rfc/rfc8342>>.

## Acknowledgments

TBD

## Appendix A: Code Examples

All examples in this appendix are intended as [RFC7951] JSON instance data, following the YANG instance data encoding rules for JSON. The duration field values are expressed in seconds unless otherwise noted.

## Code Examples for "Energy-harvesting Wireless Sensor Network"

As described in Section 2.3 of [RFC9657], in an energy-harvesting wireless sensor network, nodes rely exclusively on environmental sources for power, such as solar panels. On-board power levels may fluctuate based on various factors. This example assumes that a node will only power its radio when available power is over some threshold. In this scenario, the TVR Node YANG Module is not applicable, since the node attributes changes are caused by the environment, not by the instructions from managing device. The TVR Topology YANG Module may be necessary to convey the schedule of topology changes to all the nodes.

Considering a topology with three nodes, the connectivity of this three-node networks changes over time and repeats daily. The detailed change information is shown in Figure 4. The link between node1 and node2 is powered on at 8:00 and powered off at 11:00. The link between node2 and node3 is powered on at 11:00 and powered off at 16:00.




Figure 4: An example of topology connectivity of Energy-harvesting Wireless Sensor Network

The corresponding JSON example is shown in Figure 5.



```
{
  "ietf-tvr-topology:topology-schedule": {
    "nodes": [
      {
        "node-id": "192.168.0.1",
        "available": {
          "default-node-available": true,
          "schedules": []
        }
      },
      {
        "node-id": "192.168.0.2",
        "available": {
          "default-node-available": true,
          "schedules": []
        }
      },
      {
        "node-id": "192.168.0.3",
        "available": {
          "default-node-available": true,
          "schedules": []
        }
      }
    ],
    "links": [
      {
        "source-node": "192.168.0.1",
        "source-link-id": "link1",
        "available": {
          "schedules": [
            {
              "schedule-id": 1,
              "recurrence-first": {
                "start-time-utc": "2026-01-01T08:00:00Z",
                "duration": 10800
              },
              "utc-until": "2027-01-01T00:00:00Z",
              "frequency": "ietf-schedule:daily",
              "attr-value": {
                "link-available": true
              }
            }
          ],
          "default-link-available": false
        }
      }
    ]
  }
}
```

```
"source-node": "192.168.0.2",
"source-link-id": "link1",
"available": {
  "schedules": [
    {
      "schedule-id": 2,
      "recurrence-first": {
        "start-time-utc": "2026-01-01T08:00:00Z",
        "duration": 10800
      },
      "utc-until": "2027-01-01T00:00:00Z",
      "frequency": "ietf-schedule:daily",
      "attr-value": {
        "link-available": true
      }
    }
  ],
  "default-link-available": false
},
{
  "source-node": "192.168.0.2",
  "source-link-id": "link2",
  "available": {
    "schedules": [
      {
        "schedule-id": 3,
        "recurrence-first": {
          "start-time-utc": "2026-01-01T11:00:00Z",
          "duration": 18000
        },
        "utc-until": "2027-01-01T00:00:00Z",
        "frequency": "ietf-schedule:daily",
        "attr-value": {
          "link-available": true
        }
      }
    ],
    "default-link-available": false
  },
  {
    "source-node": "192.168.0.3",
    "source-link-id": "link1",
    "available": {
      "schedules": [
        {
          "schedule-id": 4,
```

```

        "recurrence-first": {
            "start-time-utc": "2026-01-01T11:00:00Z",
            "duration": 18000
        },
        "utc-until": "2027-01-01T00:00:00Z",
        "frequency": "ietf-schedule:daily",
        "attr-value": {
            "link-available": true
        }
    },
    "default-link-available": false
}
]
}
}

```

Figure 5: JSON example of topology schedule for Energy-harvesting  
Wireless Sensor Network

#### Code Examples for "Cellular Network"

As described in Section 3.3 of [RFC9657], the "Cellular Network" is a network where the nodes operating over cellular connections that charge both peak and off-peak data rates. In this case, individual nodes may be allocated a fixed set of "peak" minutes such that exceeding that amount of time results in expensive overage charges. As a result, links are just available at specific "peak" minutes. In this scenario, both the TVR node YANG module and TVR topology YANG module are applicable to manage the state of node interfaces and deliver the predicted topology changes to each node.

Considering a topology with three nodes, the connectivity variation of it is shown in Figure 4. Taking the nodel as an example, the corresponding node YANG module JSON example for nodel is shown in Figure 6

```

{
  "ietf-tvr-node:node-schedule": {
    "node-id": "192.168.0.1",
    "node-power-schedule": {
      "power-default": true,
      "schedules": []
    },
    "interface-schedule": {
      "interface": [
        {
          "name": "interfacel",
          "default-available": false,
          "schedules": [
            {
              "schedule-id": 100,
              "recurrence-first": {
                "start-time-utc": "2026-01-01T08:00:00Z",
                "duration": 10800
              },
              "utc-until": "2027-01-01T00:00:00Z",
              "frequency": "ietf-schedule:daily",
              "attr-value": {
                "available": true
              }
            }
          ]
        }
      ]
    }
  }
}

```

Figure 6: TVR node YANG module JSON example of node1

The corresponding topology yang module JSON example is the same as Figure 5

#### Code Examples for "Tidal Network"

As described in Section 3.4 of [RFC9657], the "Tidal Network" is a network where traffic volume undergoes significant fluctuations at different times. In the context of a tidal network scenario, energy-saving methods may include the deactivation of some or all components of network nodes as planned. In this scenario, both the TVR node YANG module and TVR topology YANG module are applicable to manage the state of node (interfaces) and deliver the predicted topology changes to each node. Figure 7 shows a tidal network topology with 4 nodes.




Figure 7: An example topology of Tidal Network

Taking the node N1 as an example, assuming the node-ids of N1, N2, N3, N4 are 192.168.0.1, 192.168.0.2, 192.168.0.3, and 192.168.0.4. The corresponding node YANG module JSON example for it is shown in Figure 8

```

{
  "ietf-tvr-node:node-schedule": {
    "node-id": "192.168.0.1",
    "node-power-schedule": {
      "power-default": true,
      "schedules": []
    },
    "interface-schedule": {
      "interface": [
        {
          "name": "interface2",
          "default-available": false,
          "schedules": [
            {
              "schedule-id": 100,
              "recurrence-first": {
                "start-time-utc": "2026-01-01T08:00:00Z",
                "duration": 54000
              },
              "utc-until": "2027-01-01T00:00:00Z",
              "frequency": "ietf-schedule:daily",
              "attr-value": {
                "available": true
              }
            }
          ]
        }
      ]
    }
  }
}

```

Figure 8: TVR node YANG module JSON example of node N1

The corresponding topology YANG module JSON example is shown in Figure 9

```

{
  "ietf-tvr-topology:topology-schedule": {
    "nodes": [
      {
        "node-id": "192.168.0.1",
        "available": {
          "default-node-available": true,
          "schedules": []
        }
      },
      {

```

```

    "node-id": "192.168.0.2",
    "available": {
      "default-node-available": true,
      "schedules": []
    }
  },
  {
    "node-id": "192.168.0.3",
    "available": {
      "default-node-available": true,
      "schedules": []
    }
  },
  {
    "node-id": "192.168.0.4",
    "available": {
      "default-node-available": true,
      "schedules": []
    }
  }
],
"links": [
  {
    "source-node": "192.168.0.1",
    "source-link-id": "link2",
    "available": {
      "schedules": [
        {
          "schedule-id": 100,
          "recurrence-first": {
            "start-time-utc": "2026-01-01T08:00:00Z",
            "duration": 54000
          },
          "utc-until": "2027-01-01T00:00:00Z",
          "frequency": "ietf-schedule:daily",
          "attr-value": {
            "link-available": true
          }
        }
      ]
    },
    "default-link-available": false
  }
],
{
  "source-node": "192.168.0.2",
  "source-link-id": "link2",
  "available": {
    "schedules": [

```

```
{
  "schedule-id": 200,
  "recurrence-first": {
    "start-time-utc": "2026-01-01T08:00:00Z",
    "duration": 54000
  },
  "utc-until": "2027-01-01T00:00:00Z",
  "frequency": "ietf-schedule:daily",
  "attr-value": {
    "link-available": true
  }
},
"default-link-available": false
},
{
  "source-node": "192.168.0.3",
  "source-link-id": "link2",
  "available": {
    "schedules": [
      {
        "schedule-id": 300,
        "recurrence-first": {
          "start-time-utc": "2026-01-01T08:00:00Z",
          "duration": 54000
        },
        "utc-until": "2027-01-01T00:00:00Z",
        "frequency": "ietf-schedule:daily",
        "attr-value": {
          "link-available": true
        }
      }
    ],
    "default-link-available": false
  }
},
{
  "source-node": "192.168.0.4",
  "source-link-id": "link2",
  "available": {
    "schedules": [
      {
        "schedule-id": 400,
        "recurrence-first": {
          "start-time-utc": "2026-01-01T08:00:00Z",
          "duration": 54000
        },

```



```

        "utc-until": "2027-01-01T00:00:00Z",
        "frequency": "ietf-schedule:daily",
        "attr-value": {
            "link-available": true
        }
    },
    "default-link-available": false
}
]
}
}

```

Figure 9: JSON example of topology schedule for Tidal Network

#### Code Examples for "Mobile Satellites"

As described in Section 4.3 of [RFC9657], the "Mobile Satellites" generally refers to the Low Earth Orbit(LEO) network, which includes hundreds to thousands of spacecrafts that can communicate both with their orbital neighbors as well as down to any ground station that they happen to be passing over. The connection between the spacecrafts and the ground station depend on the flight trajectories of the spacecrafts, so the link changes between them is predictable.

Section 4.3 of [RFC9657] introduces a scenario with 3 spacecrafts and 1 ground station. The changes of topology are shown in Figure 10. The ground station connects to spacecraft N3 at time t1, connects to N2 at time t2, and connects to N1 at time t3. The duration of the connection depends on the satellite altitude and the elevation angle. According to Section 2.1 of [I-D.lj-rtg-sat-routing-consideration], for the spacecrafts at the 500km altitude, and the connection between the spacecraft and ground station can keep for 7 minutes.

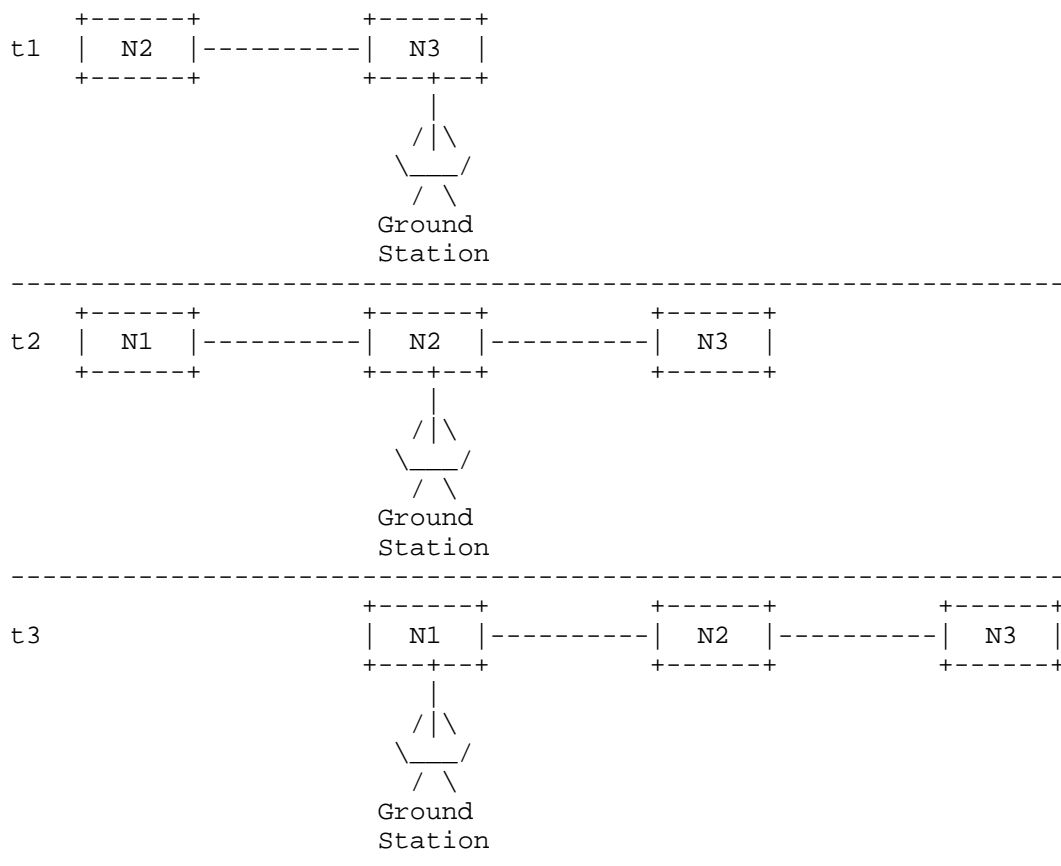


Figure 10: An example topology for Mobile Satellites

In this scenario, the TVR topology YANG module is applicable to deliver the predicted topology changes to each node. However, the TVR node YANG module is not applicable, this depends on whether the link changes are controlled by satellites themselves or by the managing device. Here, we provide the JSON examples for TVR topology YANG module and node YANG module as a reference.

Taking the spacecraft N1 as an example, assuming the time t1 is 10:00:00 1 July 2025 and the node-ids of N1, N2, N3, and ground station is 192.168.0.1, 192.168.0.2, 192.168.0.3, and 192.168.0.4., then the corresponding node YANG module JSON example for it is shown in Figure 11.

```

{
  "ietf-tvr-node:node-schedule": {
    "node-id": "192.168.0.1",
    "node-power-schedule": {
      "power-default": true,
      "schedules": []
    },
    "interface-schedule": {
      "interfaces": [
        {
          "name": "satellite2ground-interface",
          "default-available": false,
          "schedules": [
            {
              "schedule-id": 100,
              "period-start": "2025-07-01T10:00:00Z",
              "duration": 420,
              "attr-value": {
                "available": true
              }
            }
          ]
        }
      ]
    }
  }
}

```

Figure 11: TVR node YANG module JSON example of spacecraft N1

Assuming that time t1 is 10:00:00 1 July 2025, time t2 is 10:10:00 1 July 2025, and time t3 is 10:20:00 1 July 2025, then the corresponding topology YANG module JSON example is shown in Figure 12.

```

{
  "ietf-tvr-topology:topology-schedule": {
    "nodes": [
      {
        "node-id": "192.168.0.1",
        "available": {
          "default-node-available": true,
          "schedules": []
        }
      },
      {
        "node-id": "192.168.0.2",
        "available": {

```

```
        "default-node-available": true,
        "schedules": []
      },
    },
    {
      "node-id": "192.168.0.3",
      "available": {
        "default-node-available": true,
        "schedules": []
      }
    },
    {
      "node-id": "192.168.0.4",
      "available": {
        "default-node-available": true,
        "schedules": []
      }
    }
  ],
  "links": [
    {
      "source-node": "192.168.0.3",
      "source-link-id": "gs-link",
      "available": {
        "schedules": [
          {
            "schedule-id": 100,
            "period-start": "2025-07-01T10:00:00Z",
            "duration": 420,
            "attr-value": {
              "available": true
            }
          }
        ]
      },
      "default-link-available": false
    },
    {
      "source-node": "192.168.0.2",
      "source-link-id": "gs-link",
      "available": {
        "schedules": [
          {
            "schedule-id": 200,
            "period-start": "2025-07-01T10:10:00Z",
            "duration": 420,
            "attr-value": {
              "available": true
            }
          }
        ]
      }
    }
  ]
}
```

```

        }
      },
      "default-link-available": false
    },
    {
      "source-node": "192.168.0.1",
      "source-link-id": "gs-link",
      "available": {
        "schedules": [
          {
            "schedule-id": 300,
            "period-start": "2025-07-01T10:20:00Z",
            "duration": 420,
            "attr-value": {
              "available": true
            }
          }
        ],
        "default-link-available": false
      }
    }
  ]
}

```

Figure 12: JSON example of topology schedule for Mobile Satellites

## Code examples for "Predictable Moving Vessels"

As described in Section 4.4 of [RFC9657], the "Predictable Moving Vessels" involves the movement of vessels with predictable trajectories, such as ferries or planes. These endpoints often rely on a combination of satellite and terrestrial systems for Internet connectivity. Consider a scenario where a vessel uses satellites to access the Internet, including a ship and three satellites. As the satellite and the vessel move, the vessel establishes connections with the satellites N1, N2, and N3 at t1, t2, and t3 respectively. It is assumed that each connection lasts for 1 minutes. The changes of topology are shown in Figure 13.

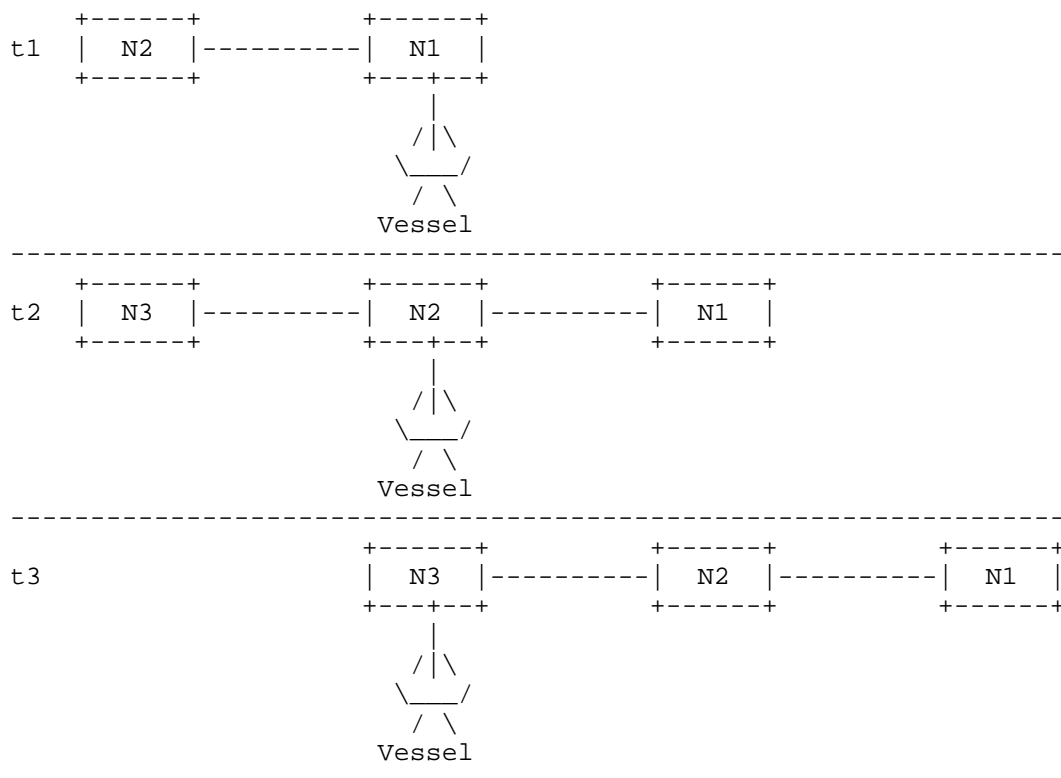


Figure 13: An example topology for Predictable Moving Vessels

This scenario is similar to the "Mobile Satellites" example, so the TVR node YANG module JSON example and topology YANG JSON example of this scenario can refer to the JSON example of the "Mobile Satellites" example.

#### Contributors

Daniel King  
Lancaster University  
United Kingdom  
Email: d.king@lancaster.ac.uk

Charalampos (Haris) Rotsos  
Lancaster University  
Email: c.rotsos@lancaster.ac.uk

Peng Liu  
China Mobile  
Email: liupengyjy@chinamobile.com

Tony Li  
Juniper Networks  
Email: tony.li@tony.li

#### Authors' Addresses

Li Zhang  
Huawei  
Email: zhangli344@huawei.com

Jie Dong  
Huawei  
Email: jie.dong@huawei.com

Mohamed Boucadair  
Orange  
Email: mohamed.boucadair@orange.com