Current Limitations in Digital Map Modelling

Kilian Matheis, Johannes Späth*, Florian Wiedner*
*Chair of Network Architectures and Services
School of Computation, Information and Technology, Technical University of Munich, Germany
Email: kilian.matheis@tum.de, spaethj@net.in.tum.de, wiedner@net.in.tum.de

***Abstract*—With the increased complexity of managing networks, the need for Digital Twins has been growing rapidly. These are based on digital maps of networks. The IETF has been working on standardising the modelling of networks and their topologies and proposed their base standard in RFC 8345. In this paper we therefore share our experience with modelling an Autonomous System based on RFC 8345. We summarise the current limitations and the possible solutions for them. The two main issues encountered were the lack of bidirectional links and the absence of links between networks. The proposed augmentations to RFC 8345 we and other researchers suggest based on shared experiences take away unnecessary complexity and offer backwards compatibility.**

***Index Terms*—RFC 8345, digital map modelling, digital twin, bidirectional links**

1. Introduction

The rapid growth of networks and rising customer demands require dynamic adaptation, therefore creating new challenges for operators. Managing these advanced networks and services is complex, and introducing innovations is risky and costly without reliable emulation platforms [1]. A rising technology to help manage these networks are Digital Twins. Zhou et al. [1] define a Network Digital Twin (NDT) as 'a digital representation that is used in the context of networking and whose physical counterpart is a data network' [1]. These NDTs can help in many ways. E.g., one of the main applications for an NDT is the testing of so called 'what-if' scenarios. Being able to test a network for possible points of failure or to test new functionalities without having to test in production is not only essential for a good consumer/user-experience but also for the maintaining institutions. Furthermore, there is the possibility to turn a network into an Autonomous Driving Network, which regulates the network autonomously using NDTs as possible test benches or simply as data input.

A Digital Map (DM) represents the topology information for a given network. It provides the core topological entities, their role in the network, core properties and relationships between networks and entities on a multi-level topology [2]. The building of digital maps is essential for digital twins [3].

We modelled a network based on the guidelines and standards provided by the IETF. Previous research has highlighted several limitations in the current version of RFC 8345, particularly regarding bidirectional links and cross-network connectivity [4] [5]. Various augmentations have been proposed to address these challenges, including improved guidelines for multi-layer network modeling [4]. Our work builds on these findings, offering a practical evaluation of the proposed solutions.

The rest of this paper is structured as follows: Chapter 2 presents necessary background information about the IETF standard RFC 8345 and Chapter 3.1 gives a short overview of our network we modelled. In Chapter 3 we cover the limitations we found and what proposed solutions and workaround we could find and also present our own opinion on how these limitations should be handled in the future and Chapter 4 summarises the limitations and the proposed solutions and gives an outlook into future work.

2. Understanding RFC 8345 – A YANG Data Model for Network Topologies

Because networks have grown in size and become more dynamic there was and still is a need for real-time topology data to support automation and programmability. To ensure interoperability and integration across multi-layer multi-vendor networks, the need for a standardised topology description has been growing. In 2018, to meet those demands, the IETF formalised their RFC 8345 [6], in which they proposed a standard describing network topologies in YANG, a data modelling language [7]. YANG organises data into a hierarchical tree structure using constructs such as containers, lists, leafs, and leaf-lists to represent complex configurations and state information. This modular approach allows for the creation of reusable and interoperable models that facilitate efficient network management and automation.

2.1. Structure of RFC 8345

The IETF introduces the data model in two divided parts. This design decision has been made to separate the integration of network topology and network inventory models. It allows the augmentation for inventory information without knowing anything about the underlaying topology of the network. According to the IETF the standard has purposefully been kept very abstract and generic so the model can represent any kind of network [6].

2.1.1. Base Network Model (ietf-network). The `ietf-network` module defines the base network data model. The model includes a container that holds a list of networks. Each network has its own entry and is unique through its primary-key `network-id`. A network can pose multiple `network-types`. This container acts as a target for augmentation, therefore it is implemented as a container, rather than an empty leaf. This approach supports hierarchical representations of network subtypes.

To model network hierarchies, a network has a list of `supporting-networks` with references to underlay networks. For a Layer-3 network this could be a reference to a Layer-2 network.

To model parts of the network the RFC 8345 introduces so-called nodes. A node is intended as an abstract construct, meaning in one network it can model a physical device and in a different network it could be a processor or a router. Each node is uniquely identified by its node-id and strictly bound to its network. Each network has its own list of nodes. Just like networks can have supporting-networks, a node can have `supporting-nodes` in underlay networks [6].

2.1.2. Base Network Topology Model (ietf-network-topology). The `ietf-network-topology` module defines the base network topology data model. It augments `ietf-network` from above. To describe the topology of networks in an abstract way, it adds two crucial elements, links and termination points.

Nodes get extended by a list of termination-points (TPs). Just like nodes, a termination point is an abstract construct unique by its `tp-id` that represents one end of a link. This could be an interface. Termination points can have `supporting-termination-points` in underlaying nodes of underlaying networks.

Links are captured in a list in networks. Each link is uniquely identified by its key `link-id`. A link consists of a source and destination, both represented by the before mentioned termination-points. A link can also be supported by underlaying links in underlay networks. Links are point-to-point and only unidirectional [6].

2.2. Augmentations of RFC 8345

The YANG model, as mentioned above, has been kept very generic. This allows for YANG augmentations to the two YANG modules. The main focus in this paper is on augmentations on `ietf-network-topology`. Havel et al. [4] analysed a plethora of augmentations and came to the conclusion that work needs to be done on creating guidelines for augmentations as the current state is not consistent and therefore very hard to combine them into multi-layer networks based on one single standard.

3. Limitations of the Current RFC 8345 Version

To represent our test-network we wanted to describe our system using the proposed RFC 8345 standard. During this process, we experienced current limitations of the base model, which have also already been encountered by other researchers [5] [4]. This motivated us to go into further research and find existing workarounds or proposed solutions to RFC 8345 and evaluate them.

3.1. Our Autonomous System Network

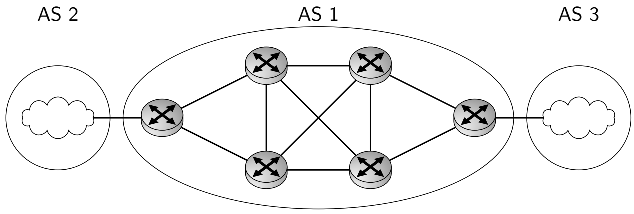

Figure 1: Network topology diagram

In our study of Autonomous Driving Networks, we modelled an Autonomous System (AS) comprising three interconnected systems. In our system we wanted to transfer files from AS2 to AS3 through AS1. Our AS1 consists of six routers. as1-host1 is connected to AS2 and as1-host6 is connected to AS3. The six routers are set up so that for data to flow from host-1 to host-6 it has to go through at least two more routers. All connections are wired, therefore bidirectional. Between the systems, the Border Gateway Protocol (BGP) is used, and our middle-system, AS1, uses the Intermediate System to Intermediate System Protocol (IS-IS).

3.2. Missing Bidirectional Links

The first limitation we experienced is that there is currently no way to model a link as bidirectional as "Links are point-to-point and unidirectional" [6]. This is an intentional design decision made by the IETF. The decision has been made to be in line with the philosophy to keep the standard on a generic level. The IETF further reasons that unidirectionality improves the applicability of graph algorithms. Although Havel et al. [5] regard this as inaccurate in practise, as in most cases further transformations have to be made before applying a graph algorithm of any kind.

3.2.1. Proposed Solutions. First it has to be said that the IETF, aware of the debate on this challenge, not only states their above mentioned reasons, but also offers their way of representing bidirectional-links. They want them to be modelled as two separate and independent uni-directional links [6]. This is a quite simple way to circumnavigate this challenge. From our perspective, this still seems rather unnecessary as it doubles the amount of links and therefore the networks complexity. This decision and the resulting independence of the two parts of a link that is supposed to be a one logical unit create the possibility of just one unidirectional link existing, therefore increasing the rate of errors.

This challenge can be solved simply through a small augmentation by adding a `direction-of-link` to the link container in `ietf-network-topology`. If one only wants to address this directionality challenge, this augmentation would be sufficient. This approach is fully backward compatible.

3.2.2. Evaluation. We would advise against resolving this limitation by oneself through augmentation and rather accept the additional complexity by modelling bidirectional links with two unidirectional ones. As the solution to this limitation is fully backwards compatible and easy to implement, we recommend the IETF to implement it in an updated version of the RFC 8345. From our perspective, the benefits would outweigh the drawbacks and an augmentation regarding this would not go against the principle of generic design. Our view is in line with other papers [5] [4].

3.3. Missing Links Between Networks

The second limitation we encountered when we modelled our AS-System is the constriction on link paths. Links can only have nodes and termination points in the same network [6]. Therefore links between two (or more) separate networks are not allowed. This poses a problem for our intended description as our network to model has 3 different AS domains (AS1, AS2, AS3).

3.3.1. Proposed Solutions. If one wants to model links between two different networks using only RFC 8345, there is just one way to approach this problem. One needs to create a new, higher-level, network of domains [6]. In our case this would be a network with nodes AS1, AS2, AS3 and links connecting AS1 with AS2 and AS2 with AS3. Then describing those networks in detail in their own network. To model the links between them you would need to create duplicate nodes for the devices in the neighbouring networks and then create a link to this node from the corresponding note in the current network. This creates/emphasises another limitation that will be explained in chapter 3.5.

A more complex solution in the current state of RFC 8345 is to use the augmentation for Traffic Engineering (TE) [8]. TE-topologies are specialised augmentations used for traffic engineering, which often necessitate additional complexity due to their specific application in optimising traffic flow in complex networks. This adds the possibility to link nodes between two different networks by adding a network-ref to the source and destination. This change in reference makes it unfortunately impossible for programs that only understand RFC 8345 to retrieve the topology of the network [4]. This goes against the original intent of the standard. Additionally it adds possible unwanted complexity to the describing as the network has to be of type `te-topology`.

Modelling a IS-IS Topology makes this challenge even more apparent IS-IS areas and IS-IS Domains can currently not be modelled without using a workaround. De Dios et al. [9] use node-attributes to define the IS-IS area. This method makes it impossible for a RFC 8345 to understand the network and its topology without having prior IS-IS knowledge.

3.3.2. Evaluation. We recommend extending the RFC 8345 to allow links between nodes in separate networks by adding a network-ref in both source and destination of a given link. This solution would also be backward compatible [4]. This extension would minimise the complexity to model AS/IS-IS/OSPF networks, allow for simple RFC 8345 algorithms to understand those networks. We do not think that this addition goes against the spirit of having an abstract base model, because multiple connected networks are of high occurrence and not technology specific. This view is shared by Havel et al. [4]. Technology-specific augmentations, e.g. IS-IS and Open-Shortest-Path-First (OSPF), are also on hold regarding this challenge until it will be addressed in the base module.

3.4. Missing Option for Subnetworks/Partitioning

This limitation is closely related to the limitation on links between networks from Section 3.3. Currently, RFC 8345 does not provide a method for describing one network as part of another. This limitation is especially relevant in protocols such as OSPF and IS-IS, where networks are often divided into areas that need to be modeled as part of a larger structure. Havel et al. [4] propose to add a simple part-of relationship between networks to RFC 8345 . This proposal is aligned with our view, because it would greatly simplify the modelling of our own AS domain. It is also backward compatible.

3.5. Duplicate Nodes, TPs in Multiple Networks

In the current version of RFC 8345, nodes and TPs are only allowed to be in one network, meaning that if the same physical or logical node or TP exists in two networks, it must have different keys and therefore be two independent nodes from a modelling point-of-view. This would be the case for our AS-Domain as well, if we choose to use the workaround with two duplicate nodes in their respective network from Section 3.3. Havel et al. [4] state that this is the case with OSPF Networks and shows the possibility to have this challenge in IS-IS. This introduces the challenge of consistency when writing to a model, because one would have to assume that all duplicate nodes are in the same state. This is not guaranteed by the model as they are totally independent from the models point of view [4].

3.5.1. Proposed Solutions. In our case we would have to model the as1-host1, as2-host1 routers in the as1 AND the as2 network and the as2-host6 and as3-host1 routers in the as2 AND as3 network. The IETF is aware of the challenge and proposes the solution to have an extra network of physical nodes that represent the existing devices and are linked to the logical nodes from networks through the supporting-node logic [6]. In our case we would have to create a 5th physical-network with all our routers and then add those routers to their respective nodes in as1, as2 and as3.

Havel et al. [4] also propose to allow the definition of nodes outside the scope of networks. This solution would likely be backwards compatible but greatly alter the topology tree and therefore needs to be researched further.

3.5.2. Evaluation. In our opinion this limitation is, just like the limitation in Chapter 3.3, not technology specific and it would only be sensible to allow nodes to be defined outside of networks, to make them uniquely identifiable across multiple networks, as long as the solution is backwards compatible. Allowing nodes to exist independently of specific networks would streamline the process of maintaining consistency across domains, as it would eliminate the need to synchronise duplicate nodes and their states/interfaces.

3.6. Missing Multi-Point Links

While modelling our Autonomous System, we did not initially encounter the absence of multipoint links as a limitation. However, during our research into the other shortcomings of RFC 8345, this challenge became apparent. Multipoint links, which enable the connection of multiple devices through a single logical link, are a critical feature in many modern network environments such as Layer 2 topologies. For instance, in Virtual-Local-Area-Network (VLAN) configurations, multipoint links allow multiple endpoints to communicate across a shared infrastructure. Unfortunately, RFC 8345 does not natively support this type of connection. Instead, it only allows for point-to-point links, which limits the ability to accurately model networks that require multipoint communication. This limitation originates from the core design of RFC 8345, where links are described as unidirectional and strictly point-to-point.

3.6.1. Proposed Solutions. The current workaround proposed by RFC 8345 is to model multipoint connections using pseudonodes [6]. Pseudonodes act as intermediary points to represent multipoint connectivity indirectly. By creating a new node for each multipoint connection, network designers can model the individual connections as point-to-point links between the devices and the pseudonode. However, this approach significantly increases network complexity, making the model more difficult to manage and prone to inconsistencies. The RFC acknowledges this limitation but offers no further improvements, stating that this method preserves the standard's generic, technology-agnostic design.

An alternative solution, discussed by Davis et al. [5], suggests augmenting RFC 8345 to allow links to directly have multiple termination points, enabling more effective support for multipoint connections. This approach would preserve backward compatibility, ensuring that existing models remain functional while simplifying the representation of networks using multipoint links. Like the bidirectional link augmentation mentioned earlier in Section 3.2, this enhancement would streamline the modeling process by reducing redundant link definitions, making network maps cleaner and more manageable. It introduces an additional, optional link-end definition, which fits within the current structure, ensuring that current point-to-point links continue to operate without modification. The simplicity of this solution would allow for broad adoption without disrupting existing network models.

A more advanced approach is also suggested to improve the integrity of the existing model [5]. This method involves extracting the current source and destination structures and then re-augmenting them back into the link, all within the same module. Since there is no namespace change, the augment remains within the module. While this solution is not fully backward compatible with YANG, it still produces the same instance structures (e.g., in JSON) and supports any existing augmentations for source and destination. One major advantage of this approach is that it allows for point inclusion to be controlled based on feature support, effectively separating the structures supporting existing capabilities from those designed to handle new functionalities [5].

3.6.2. Evaluation. From our perspective, we strongly advise the adoption of the simpler, backward-compatible solution into RFC 8345. This approach would resolve the limitation of multipoint links without increasing the complexity of the model unnecessarily, allowing for easier and more accurate network representation. Given the high occurrence of multipoint connections in modern networks, this change would make RFC 8345 far more applicable in real-world scenarios while preserving the standard's abstract and generic nature.

4. Conclusion

In summary, our evaluation of RFC 8345 has brought to light several limitations that complicate network topology modeling, particularly in Autonomous Systems. Key challenges include the lack of support for bidirectional links, restrictions on inter-network connectivity, and the absence of a mechanism for representing subnetworks or partitions. While manageable through workarounds, these issues unnecessarily increase complexity and risk of inconsistency.

To improve the applicability of RFC 8345, we believe the IETF should consider adjustments that remove these limitations while retaining backward compatibility. Specifically, enabling bidirectional links and allowing cross-network references would provide a more accurate and streamlined model for real-world networks. These enhancements would preserve the flexibility and abstraction of the standard, while making it more practical for advanced network architectures. Addressing these concerns would lead to more efficient implementations and broader adoption within the networking field.

By addressing these limitations, RFC 8345 would better support the development of Digital Twins and other advanced network architectures, ensuring its relevance in the evolving landscape of network management.

Further work has to be done on how augmentations for specific layers and protocols can or need to be changed when these improvements to RFC 8345 are implemented.

References

[1] C. Zhou, H. Yang, X. Duan, D. Lopez, A. Pastor, Q. Wu, M. Boucadair, and C. Jacquenet, "Network Digital Twin: Concepts and Reference Architecture," Internet Engineering Task Force, Internet-Draft draft-irtf-nmrg-network-digital-twin-arch-06, Jul. 2024, work in Progress. [Online]. Available: https://datatracker.ietf.org/doc/draft-irtf-nmrg-network-digital-twin-arch/06/

[2] Benoît Claise, "Digital Map Modelling: IETF 117 Latest Developments," 2023, available online at https://www.claise.be/digital-map-modelling-ietf-117-latest-developments/; last accessed on 2024/09/18.

[3] O. Havel, B. Claise, O. G. de Dios, and T. Graf, "Digital Map: Concept, Requirements, and Use Cases," Internet Engineering Task Force, Internet-Draft draft-havel-nmop-digital-map-concept-00, Jul. 2024, work in Progress. [Online]. Available: https://datatracker.ietf.org/doc/draft-havel-nmop-digital-map-concept/00/

[4] O. Havel, B. Claise, O. G. de Dios, A. Elhassany, and T. Graf, "Modeling the Digital Map based on RFC 8345: Sharing Experience and Perspectives," Internet Engineering Task Force, Internet-Draft draft-havel-nmop-digital-map-01, Jul. 2024, work in Progress. [Online]. Available: https://datatracker.ietf.org/doc/draft-havel-nmop-digital-map/01/

[5] N. Davis, O. Havel, and B. Claise, "Some Refinements to Network Topologies (RFC8345)," Internet Engineering Task Force, Internet-Draft draft-davis-nmop-some-refinements-to-rfc8345-00, Jul. 2024, work in Progress. [Online]. Available: https://datatracker.ietf.org/doc/draft-davis-nmop-some-refinements-to-rfc8345/00/

[6] A. Clemm, J. Medved, R. Varga, N. Bahadur, H. Ananthakrishnan, and X. Liu, "A YANG Data Model for Network Topologies," RFC 8345, Mar. 2018. [Online]. Available: https://www.rfc-editor.org/info/rfc8345

[7] M. Björklund, "The YANG 1.1 Data Modeling Language," RFC 7950, Aug. 2016. [Online]. Available: https://www.rfc-editor.org/info/rfc7950

[8] X. Liu, I. Bryskin, V. P. Beeram, T. Saad, H. C. Shah, and O. G. de Dios, "YANG Data Model for Traffic Engineering (TE) Topologies," RFC 8795, Aug. 2020. [Online]. Available: https://www.rfc-editor.org/info/rfc8795

[9] O. G. de Dios, S. Barguil, V. Lopez, D. Ceccarelli, and B. Claise, "A YANG Data Model for Intermediate System to intermediate System (IS-IS) Topology," Internet Engineering Task Force, Internet-Draft draft-ogondio-nmop-isis-topology-00, Mar. 2024, work in Progress. [Online]. Available: https://datatracker.ietf.org/doc/draft-ogondio-nmop-isis-topology/00/