



Core Information Model (CoreModel)

TR-512.A.3 Appendix – Model Rationale

Version 1.6
September 2024

ONF Document Type: Technical Recommendation
ONF Document Name: Core Information Model version 1.6

Disclaimer

THIS SPECIFICATION IS PROVIDED "AS IS" WITH NO WARRANTIES WHATSOEVER, INCLUDING ANY WARRANTY OF MERCHANTABILITY, NONINFRINGEMENT, FITNESS FOR ANY PARTICULAR PURPOSE, OR ANY WARRANTY OTHERWISE ARISING OUT OF ANY PROPOSAL, SPECIFICATION OR SAMPLE.

Any marks and brands contained herein are the property of their respective owners.

Open Networking Foundation
1000 El Camino Real, Suite 100, Menlo Park, CA 94025
www.opennetworking.org

©2024 Open Networking Foundation. All rights reserved.

Open Networking Foundation, the ONF symbol, and OpenFlow are registered trademarks of the Open Networking Foundation, in the United States and/or in other countries. All other brands, products, or service names are or may be trademarks or service marks of, and are used to identify, products or services of their respective owners.

Important note

This Technical Recommendations has been approved by the Project TST, but has not been approved by the ONF board. This Technical Recommendation is an update to a previously released TR specification, but it has been approved under the ONF publishing guidelines for 'Informational' publications that allow Project technical steering teams (TSTs) to authorize publication of Informational documents. The designation of '-info' at the end of the document ID also reflects that the project team (not the ONF board) approved this TR.

Table of Contents

Disclaimer 2

Important note 2

Document History..... 3

1 Introduction to the document suite..... 4

 1.1 References 4

 1.2 Definitions..... 4

 1.3 Conventions..... 4

 1.4 Viewing UML diagrams 4

 1.5 Understanding the figures 4

 1.6 Appendix Overview 4

2 Introduction to this Appendix document 5

 2.1 Business need 5

 2.2 Benefit of the ONF CIM..... 6

 2.3 Model evolution..... 7

 2.3.1 Agile approach for model evolution..... 9

 2.3.2 Industry cooperation in model evolution..... 10

List of Figures

Figure 2-1 Normalization in Information Modeling 6

Figure 2-2 Relationship between past current and future network model work..... 8

Figure 2-3 Agile Architecture process 10

Document History

Version	Date	Description of Change
		Appendix material was not published prior to Version 1.3
1.3	September 2017	Version 1.3 [Published via wiki only]
1.3.1	January 2018	Addition of text related to approval status.
1.4	November 2018	No change.
1.5	September 2021	Enhancements to model structure
1.6	January 2024	Updated release and dates.

1 Introduction to the document suite

This document is an appendix of the addendum to the TR-512 ONF Core Information Model and forms part of the description of the ONF-CIM. For general overview material and references to the other parts refer to [TR-512.1](#).

1.1 References

For a full list of references see [TR-512.1](#).

1.2 Definitions

For a full list of definition see [TR-512.1](#).

1.3 Conventions

See [TR-512.1](#) for an explanation of:

- UML conventions
- Lifecycle Stereotypes
- Diagram symbol set

1.4 Viewing UML diagrams

Some of the UML diagrams are very dense. To view them either zoom (sometimes to 400%) or open the associated image file (and zoom appropriately) or open the corresponding UML diagram via Papyrus (for each figure with a UML diagram the UML model diagram name is provided under the figure or within the figure).

1.5 Understanding the figures

Figures showing fragments of the model using standard UML symbols and also figures illustrating application of the model are provided throughout this document. Many of the application-oriented figures also provide UML class diagrams for the corresponding model fragments (see [TR-512.1](#) for diagram symbol sets). All UML diagrams depict a subset of the relationships between the classes, such as inheritance (i.e. specialization), association relationships (such as aggregation and composition), and conditional features or capabilities. Some UML diagrams also show further details of the individual classes, such as their attributes and the data types used by the attributes.

1.6 Appendix Overview

This document is part of the Appendix to TR-512. An overview of the Appendix is provided in [TR-512.A.1](#).

2 Introduction to this Appendix document

This document explains the reasoning behind the current model and its evolution. The document relates some of the patterns and architecture in [TR-512.A.2](#) to business need and vision. The document:

- Works through business need.
- Discusses the implications
- Discusses the challenge of evolution
- Explains how the current model satisfies the need in the context of the current stage of evolution
- Explains the next steps in the context of evolution and vision

The descriptions in this document are built from descriptions in earlier referenced works.

2.1 Business need

Management of networks and devices today is a complex operational challenge resulting from, and exacerbated by, the plethora of conflicting standards and incompatible implementations. Almost all the existing models use inconsistent terminology and outdated concepts that aren't applicable to SDN/NFV scenarios¹.

The root cause of many of the issues encountered centers around the approaches that have been used to create these standards and interfaces; e.g.:

- A 'handcrafted silo' approach that initially produces good point solutions results but doesn't scale, ultimately creating a morass requiring complex mappings between the silos;
- A vast monolithic 'one central model' that results in slow progress after long lead times with second-rate results that don't meet stakeholder needs.

Whenever new technologies and management/control paradigms arise, such as SDN and virtualization, they initially all seems different, new and special. But eventually we realize that there is rather a lot in common, as we again see the patterns. In fact, neither SDN/NFV² nor virtualization³ are in any way new. There are increasingly complex resource interactions to be managed-controlled such as virtual machines supporting virtual switches. But it seems there are consistent patterns that could be leveraged.

So, there are opportunities, and our goal should be, to massively improve the realization of the value fabric:

- The way in which providers of services interact, intertwine and interdepend to provide every increasing value - we are all providers of service

¹ The current drive for SDN and virtualization give us an opportunity though to change the way of operating.

² Much of what is described in SDN/NFV as control appears already present in networks as Management. SDN/NFV control layer functions are mostly the same as functions in the OSS/NMS.

³ Virtualization is not new but has been ad-hoc. Most of the functions of OSS/BSS environment are necessary even in the most virtualized and automated world.

- There is a need to enable an ecosystem of developers and an ever-changing structure of interacting parties.

It is undeniable that the value fabric has high inherent complexity (deep recursions, fractals and continuums) but there are artificial boundaries (which cause unnecessary cost)⁴. There is a need for coherent treatment, broad perspective and long-term view, and for players to abide by enabling constraints to allow efficient interaction.

Management is control (e.g., EMS/NMS flow through and restoration). Management and control is a continuum (MCC). Any artificial boundaries only complicate the solutions. The solution should focus on a coherent component architecture and structure.

2.2 Benefit of the ONF CIM

Information Modeling is the process of identifying, defining, consistently labeling, and recording the key concepts (functional things, physical things etc.) in a problem space as well as the interrelationships between them. An information model defines and labels the things in a domain—for example, a telecommunications network—in terms of objects, such as links, terminations and controllers; their properties or attributes, such as latency; and their relationships—for example, a link is bounded by terminations. By normalizing the concepts, objects, interfaces, etc., an information model reduces complexity (see Figure 2-1). In addition, an information model can be used, with appropriate tooling, to automate code production, helping accelerate product and service development while reducing coding errors and interoperability issues.

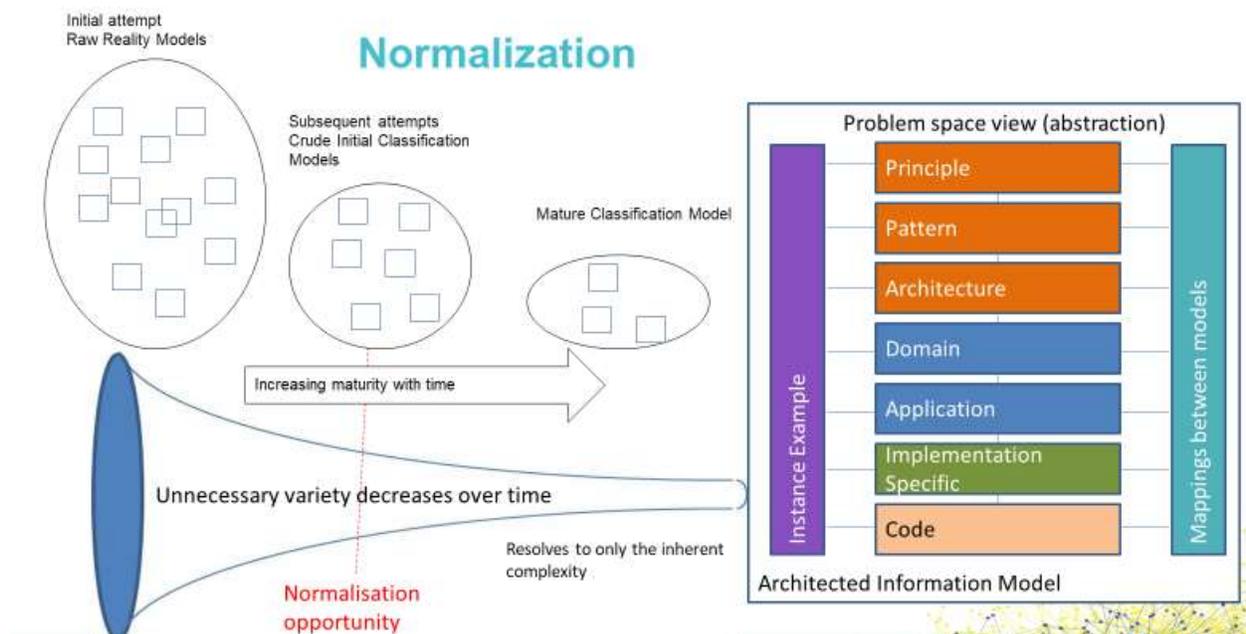


Figure 2-1 Normalization in Information Modeling

To avoid the pitfalls mentioned in section 2.1 above, the ONF Information Modelling team:

⁴ Existing divisions should be ignored so that to allow reconsideration of structure of the solution.

- Has defined a consistent set of fundamental concepts and the relationship among them by leveraging the knowledge gained from years of management standards evolution and pragmatic implementation/software development experience. These concepts are capable of representing both legacy management and SDN/NFV concepts/scenarios, while allowing for consistent management in hybrid environments.
- Employs a realistic federated model with a layered model architecture and managed dependencies. This is comprised of a stable core model (which is itself modular for scalability) and technology-specific model extensions that can be added in a timely manner without destabilizing the core.

The ONF Information model is not a purist, theoretical artifact, but a pragmatic one that forms part of a tooling chain, enabling context and technology specific interfaces in different languages to be generated from a key set of definitions. Automating the interface writing avoids inconsistent silos as a change can be simultaneously made in many interfaces in a consistent manner. This approach has led to the development of TAPI, which formally incorporates transport technology specific properties from ITU-T.

The value of a common model approach has been recognized both within ONF and by external SDOs such as ITU-T SG15, ETSI NFV, MEF and TMF, as well as by some major service providers. These external SDOs recognized that the ONF model draws on many years of experience to consolidate the key aspects of existing standards; is a stable base to work from and is more advanced than their current work in these areas. To further its development, the MEF, ONF and TMF have signed a formal collaboration agreement to continue develop this common information model. In addition, ITU-T has re-published the Core Model as Recommendation G.7711.

2.3 Model evolution

The ONF CIM, which built on 30 years of implementation experience and standards evolution (see Figure 2-2 below), provides a rich⁵, compact⁶ canonical⁷ model covering the problems spaces relevant to management/control for open networking.

⁵ Including all problem space semantics. Some aspects of the problem space are covered by other standards activities (the properties of particular transport protocols for example) and, as described, a method for federating activities has been developed so that the work of each body can be applied efficiently and coherently to an implementation.

⁶ A particular level of encapsulation of semantics is used that has been developed over many years of practical experience. This reduces volume in implementations compared to an atomic form. The model takes full advantage of the patterns of networking such that one small set of entities cover ALL technologies without the need for network technology based subclassing specialization (i.e. the model is network technology agnostic).

⁷ Each concept is represented once.

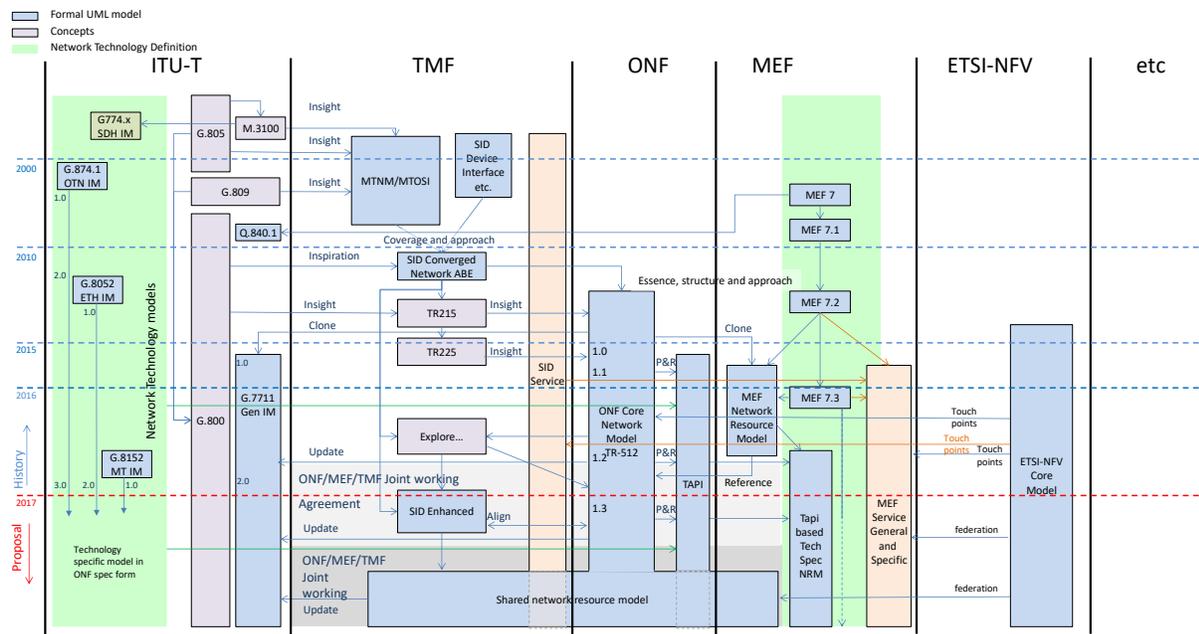


Figure 2-2 Relationship between past current and future network model work

Whilst the ONF Core Model could be used directly to construct interfaces, the modeling team recognizes that an interface should be oriented towards the user's needs from the user's viewpoint (terminology/structure etc.). That orientation includes removal of aspects not relevant to the viewpoint and also refactoring/relabeling the representation to suit the viewpoint. A process called "Pruning and Refactoring" is used to construct information model representations of relevant interface views from the Core Model⁸. This supports the current variety of terminology whilst providing a rigorous mapping back to the canonical form⁹.

Having a Core Model, a process for using the Core Model to construct interface views (Pruning and Refactoring) and to convert those views to implementation schema (e.g., UML-to-YANG Mapping Guidelines (see [UML-YANG GUIDE])) along with tools that support the process is vital. Using tooling (e.g., UML-to-YANG tool (see [UML-YANG TOOL])) that consistently generates the implementations is key to ensuring unambiguous and interoperable products and open source software. This also helps reduce errors and enables code to be more easily produced. The [OSSDN-Eagle] (IISOMI) activity currently focuses on process/methodology and tooling.

As the ONF Core Model is network technology agnostic (see [TR-512.2](#)) it is important that there is a mechanism to enable consistent addition of specific technology details. The CIM provides such a mechanism that is also designed to support network technology evolution. As ONF does not own the definition of any particular network technology, the ONF modeling team looks to other bodies (e.g. ITU-T) to provide the interface definitions (e.g., the OTN IM in [ITU-T

⁸ Tooling is under agile development in the ONF-sponsored Open Source [OSSDN-EAGLE] project to ease the Pruning and Refactoring process.

⁹ The current state of unnecessary variety of terminology that has emerged as a result of the various origins of what is now one piece of work needs to be accommodated. The methodology promotes the reduction of this variety over time (as machines do not benefit from the variety but instead run most efficiently using the patterns).

G.874.1]). The mechanism also allows vendor proprietary addition. For example, a pre-standard vendor innovation may be accommodated on an equal footing with the technology standard properties.

Addition of network technology details employs a specification mechanism defined in the ONF Core Model (see [TR-512.7](#)). Technology specific specification classes are used to augment the core model classes (e.g., media for information transfer). Where a new technology introduces a capability that has structure that could be applied more broadly, this structure will be incorporated in the Core Model. A vendor may take a technology specification, prune out parts that are not supported by their implementation, add vendor proprietary parts and then publish the results in a per-case specification. This specification can then be used to drive the augmentation of the Core Model and any derivative of the model.

2.3.1 Agile approach for model evolution

An agile architecture approach has been used by the modeling team to develop the ONF Information model:

- Take some specific and general use cases in the context of control of networking
- Explore and generate/refine the model¹⁰/guidelines/tooling building on the existing model/guidelines/tooling
- Use the Prune and Refactor process to propagate the model changes to each existing interface views [ONF TAPI] and to generate new interface views
- Update the data schema related to each view [ONF TAPI] using the appropriate tooling (e.g. UML to YANG automated generator)
- Generate the appropriate Interface schema
- Generate reference implementations
- Construct proprietary and open source solutions [e.g. OIF/ONF and MEF PoC/Trial activities]
- Gain insights from implementation and feedback insights/issues to further evolve the model/guidelines/tooling

¹⁰ Includes Core Model and Profiles

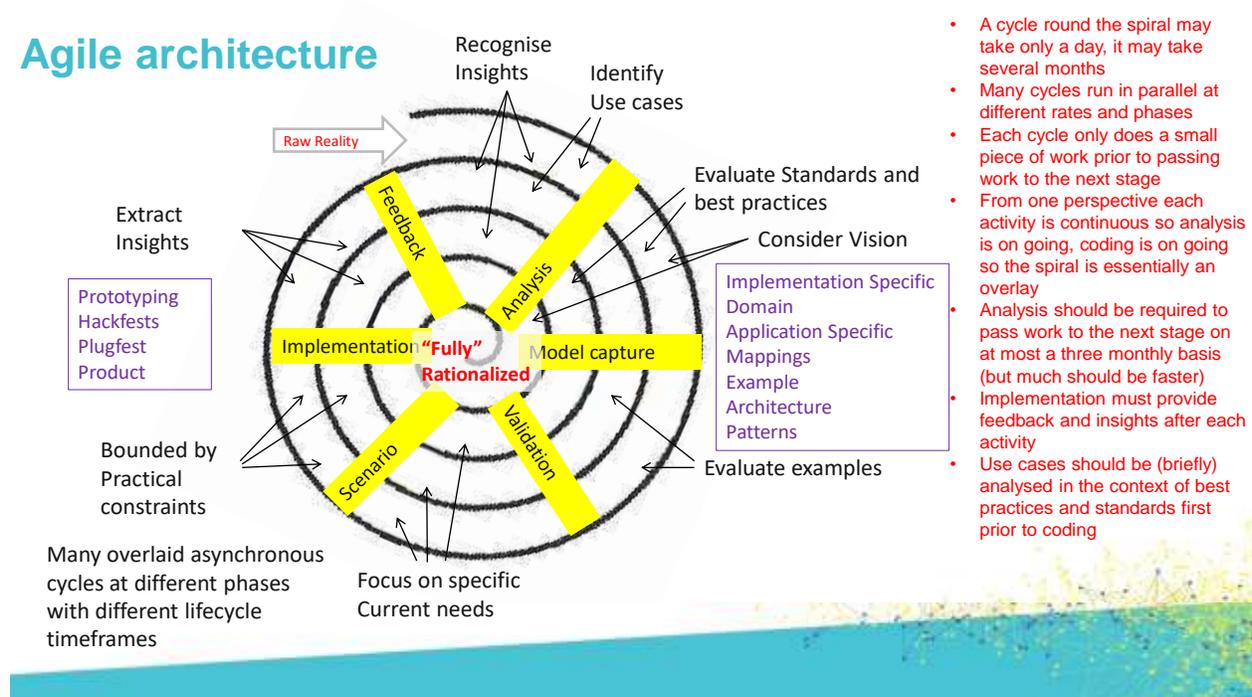


Figure 2-3 Agile Architecture process

The above cycle can be run asynchronously at different rates depending upon the specific aspect/use case. Experimental changes can take place in the model/guidelines/tooling within a day (in the ongoing work).

The modeling team has promoted open source software tooling and hence there is a strong association between the modeling work and software in that dimension. This software is being developed using an agile approach and hence there have been incremental releases on a regular basis supporting use cases defined for interface views. This software is a step in a progression towards a model driven implementation environment. The intention is to reduce the need to hand craft interface (and in the longer term potentially algorithmic) implementations.

To help enable evolution and to deal with the stages of innovation value lifecycle, each aspect and element of the Information Model is developed through a lifecycle where the lifecycle stage/state (Experimental, Preliminary, Mature, etc.) of the aspect/element is marked using Lifecycle Stereotypes define in the Open Model profile [ONF TR-514]. Any particular part of the model may include elements with any mix of lifecycle stages. The full model, including experimental and preliminary work, is published as a snapshot.

2.3.2 Industry cooperation in model evolution

The ONF modeling team has strongly promoted the use of the ONF CIM within ONF and also with other industry groups, including MEF, TMF, ITU-T, ETSI-NFV, [OASIS TOSCA], OIF, IETF, and BBF etc. ONF has signed a collaboration agreement [ONF-TMF-MEF] with MEF and TMF to work jointly on information modeling and the intention is that the ONF CIM form a key part of this joint modeling work. Work in TMF provides guidance and content for the new

solution, such as SDF and DSRA provide architecture, eTOM provides processes, and SID provides information structuring. The MEF and TMF have already chosen to take advantage of the ONF CIM in place of their own work and ITU-T have republished the ONF Core Model [ONF TR-512] in [ITU-T G.7711]. Work in ETSI-NFV has led to the recognition of the CIM as a suitable representation of networking in an NFV context and this has led to the construction of a touch point model relating ETSI-NFV classes to ONF networking classes.

The ONF modeling team is also tasked with providing recommendations for modeling tools and guidelines. Tooling and guidelines are dealt with in the ONF-sponsored open source SDN project EAGLE/IISOMI [OSSDN-EAGLE]. This work is shared and co-developed with individuals who are active in ETSI-NFV, ONF, TMF, MEF and ITU-T. [OSSDN-EAGLE] is also the home of tooling that takes an appropriately formed information model in Papyrus UML and generates interface definitions (currently in YANG and JSON schema but the intention is to generate any appropriate schema/API definition).

End of Document