Service Models Explained
draft-wu-opsawg-service-model-explained-03

Abstract

   The IETF has produced a considerable number of data models in the
   YANG modelling language. The majority of these are used to model
   devices and they allow access for configuration and to read
   operational status.

   A small number of YANG models are used to model services (for
   example, the Layer Three Virtual Private Network Service Model
   produced by the L3SM working group).

   This document briefly sets out the scope of and purpose of an IETF
   service model, and it shows where a service model might fit into a
   Software Defined Networking architecture or deployment.

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 http://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 March 12, 2017.

Copyright Notice

Table of Contents

1. Introduction

In recent years the number of data models written in the YANG
modelling language [RFC6020] for configuration and monitoring has
blossomed. Many of these are used for device-level configuration
(for example, [RFC7223]) or for control of protocols (for example,
[RFC7407]).

Within the context of Software Defined Networking (SDN) [RFC7426]
YANG data models may be used on Southbound Interfaces (SBIs) between
a controller and network devices, and between network orchestrators
and controllers. There may also be a hierarchy of such components

with super-controllers, domain controllers, and device controllers
all exchanging information and instructions using YANG models.

Recently there has been interest in using YANG to define and document
data models that describe services in a portable way that is
independent of which network operator uses the model. These models
may be used in manual and even paper-driven service request processes
moving to IT-based mechanisms. Ultimately they could be used in
online, software-driven dynamic systems.

This document explains the scope and purpose of service models within
the IETF and describes how a service model can be used by a network
operator. Equally, this document clarifies what a service model is
not, and dispels some common misconceptions.

Other work on classifying YANG data models has been done in
[I-D.ietf-netmod-yang-model-classification]. That document provides
an important reference for this document, and also uses the term
"service model". Section 6.1 provides a comparison between these two
classifications.

2. Terms and Concepts

Readers should familiarize themselves with the description and
classification of YANG models provided in
[I-D.ietf-netmod-yang-model-classification].

The following terms are used in this document:

Network Operator: This term is used interchangeably to refer to the
   company that owns a network that provides Internet connectivity
   and services, or the individual who performs operations and
   management on that network.

Customer: Someone who purchases connectivity and other services from
   a network operator. In the context of this document, a customer
   is usually the company that runs their own network or computing
   platforms and wishes to connect to the Internet or between sites.
   Such a customer may operate an enterprise network or a data
   center. Sometimes this term may also be used to refer to the
   individual in such a company who contracts to buy services from a
   network operator. A customer as described here is a separate
   commercial operation from the network operator, but some companies
   may operate with internal customers so that, for example, an IP/
   MPLS packet network is the customer of an optical transport
   network.

Service: A network operator delivers one or more services to a
   customer. A service in the context of this document (sometimes
   called a Network Service) is some form of connectivity between
   customer sites and the Internet or between customer sites across
   the network operator's network and across the Internet, however a
   distinction should be drawn between the parameters that describe a
   service as included in a customer service model (q.v.) and a
   Service Level Agreement (SLA) as discussed in Section 5 and
   Section 7.2.

   A service may be limited to simple connectivity (such as IP-based
   Internet access), may be a tunnel (such as a virtual circuit), or
   may be a more complex connectivity model (such as a multi-site
   virtual private network). Services may be further enhanced by
   additional functions providing security, load-balancing,
   accounting, and so forth. Additionally, services usually include
   guarantees of quality, throughput, and fault reporting.

   This document makes a distinction between a service as delivered
   to a customer (that is, the service as discussed on the interface
   between a customer and the network operator) and the service as
   realized within the network (as described in
   [I-D.ietf-netmod-yang-model-classification]. This distinction is
   discussed further in Section 6

Data Model: The concepts of information models and data models are
   described in [RFC3444]. That document defines a data model by
   contrasting it with the definition of an information model, so it
   may be helpful to quote some text to give context within this
   document.

      The main purpose of an information model is to model managed
      objects at a conceptual level, independent of any specific
      implementations or protocols used to transport the data. The
      degree of specificity (or detail) of the abstractions defined
      in the information model depends on the modeling needs of its
      designers. In order to make the overall design as clear as
      possible, an information model should hide all protocol and
      implementation details. Another important characteristic of an
      information model is that it defines relationships between
      managed objects.

      Data models, conversely, are defined at a lower level of
      abstraction and include many details. They are intended for
      implementors and include protocol-specific constructs.

   Service Model: A service model is a specific type of data model. It
      describes a service and the parameters of the service in a
      portable way. The service model may divided into to categories:

      Customer Service Model: A customer service model is used to
         describe a service as offer or delivered to a customer by a
         network operator. It can be used by a human (via a user
         interface such as a GUI, web form, or CLI) or by software to
         configure or request a service and may equally be consumed by a
         human (such as via an order fulfillment system) or by a
         software component. Such models are sometimes referred to
         simply as "service models" [I-D.ietf-l3sm-l3vpn-service-model].
         A customer service model is expressed as a core set of
         parameters that are common across network operators: additional
         features that are specific to the offerings of individual
         network operators would be defined in extensions or
         augmentations of the model.

      Service Delivery Model: A service delivery model is used by a
         network operator to define and configure how a service is
         provided by the network. It can be used by a human operator
         (such as via a management station) or by a software tool to
         instruct network components. Such models are sometimes
         referred to as "network service models"
         [I-D.ietf-netmod-yang-model-classification]. A service
         delivery model is expressed as a core set of parameters that
         are common across a network type and technology: additional
         features that are specific to the configuration of individual
         vendor equipment or proprietary protocols would be defined in
         extensions or augmentations of the model.

   The distinction between a customer service model and a service
   delivery model needs to be repeatedly clarified. A customer service
   model is not a data model used to directly configure network devices,
   protocols, or functions: it is not something that is sent to network
   devices (i.e., routers or switches) for processing. Equally, a
   customer service model is not a data model that describes how a
   network operator realizes and delivers the service described by the
   model. This issue is discussed further in later sections.

3. Using Service Models

   As already indicated, customer service models are used on the
   interface between customers and network operators. This is shown
   simply in Figure 1

The language in which a customer service model is described is a
choice for whoever specifies the model. The IETF uses the YANG data
modeling language defined in [RFC6020]

The encoding and communication protocol used to exchange a customer
service model between customer and network operator are deployment-
and implementation-specific. The IETF has standardized the NETCONF
protocol [RFC6241] and the RESTCONF protocol
[I-D.ietf-netconf-restconf] for interactions "on the wire" between
software components with data encoded in XML or JSON. However, co-
located software components might use an API, while systems with more
direct human interactions might use web pages or even paper forms.

```
         --------------        Customer        ---------------------
        |              |    Service Model      |                     |
        |   Customer   | <----------------->  |   Network Operator  |
        |              |                       |                     |
         --------------                        ---------------------
```
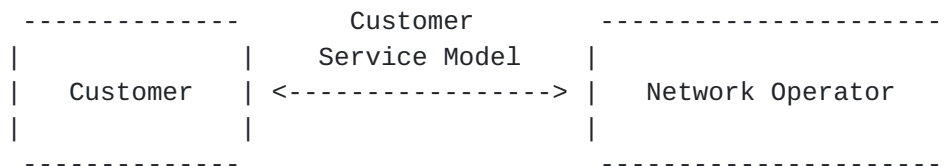
        Figure 1: The Customer Service Models used on the Interface between
                      Customers and Network Operators

How a network operator processes a customer's service request
described with a customer service model will depend on the commercial
and operational tools, processes, and policies used by the operator.
These may vary considerably from one network operator to another.

However, the intent is that the network operator maps the service
request into configuration and operational parameters that control
one or more network to deliver the requested services. That means
that the network operator (or software run by the network operator)
takes the information in the customer service model and determines
how to deliver the service by enabling and configuring network
protocols and devices. They may achieve this by constructing service
delivery models and passing them southbound to network orchestrators
or controllers.

4. Service Models in an SDN Context

In an SDN system, the control and configuration of network resources
and protocols is performed by software systems that determine how
best to utilize the network. Figure 2 shows a common architectural
view of an SDN system where network elements are programmed by a
component called a controller, and where controllers are instructed
by an orchestrator that has a wider view of the whole of, or part of,
a network.

```
                     ------------------
                     |                |
                     |   Orchestrator |
                     |                |
                     .------------------.
                    .          :          .
                   .           :           .
            ------------   ------------   ------------
            |          |   |          |   |          |
            | Controller |   | Controller |   | Controller |
            |          |   |          |   |          |
            ------------   ------------   ------------
              :          .       .          :
              :          .       .          :
              :          .       .          :
            ---------   ---------  ---------   ---------
            | Network |  | Network | | Network |  | Network |
            | Element |  | Element | | Element |  | Element |
            ---------   ---------  ---------   ---------
```
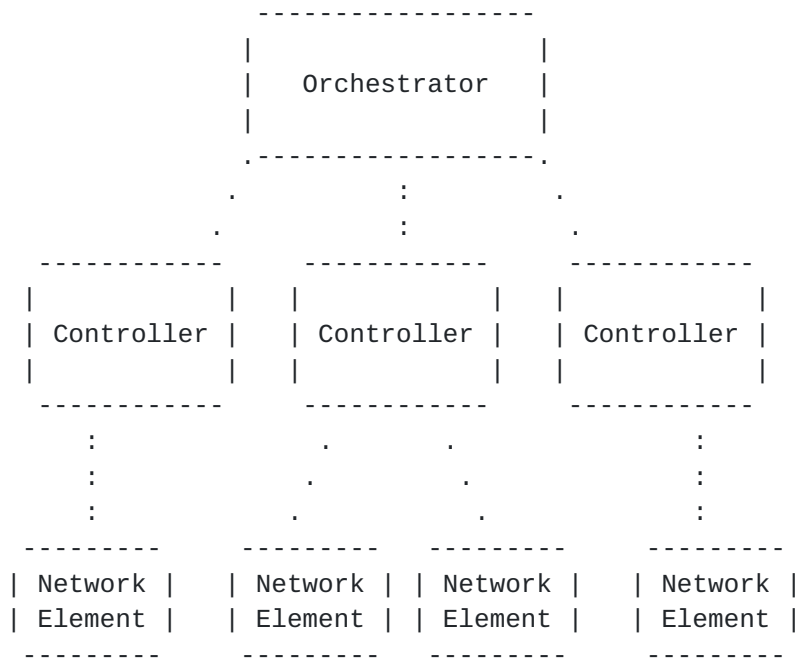
Figure 2: A Common SDN Architecture

But a customer' service request is (or should be) network-agnostic.
That is, there should be an independence between the behavior and
functions that a customer requests and the technology that the
network operator has available to deliver the service. This means
that the service request must be mapped to the orchestrator's view,
and this mapping may include a choice of which networks to use
depending on what technologies are available and which service
features have been requested.

This mapping can be achieved by splitting the orchestration function
between a "Service Orchestrator" and a "Network Orchestrator" as
shown in Figure 3. In a system that is fully implemented in
software, this could lead to agile service delivery or service
automation.

```
                                        Customer
                    -----------------   Service  ----------
                    |               |   Model   |        |
                    |    Service    |<-------->| Customer |
                    |  Orchestrator |          |        |
                    |               |           ----------
                    -----------------
                        .         .
                       .           .                 -----------
                      .             .      ......|Application|
                     .               .     :     |  BSS/OSS  |
                    .                 .     :      -----------
                   .  Service Delivery  .  :
                   .      Model         .  :
              -----------------   -----------------
              |             |   |               |
              |   Network   |   |    Network    |
              |  Orchestrator |   |  Orchestrator |
              |             |   |               |
             .-----------------   -----------------.
              .         :                 :         .
               .        : Network Configuration :         .
                .       :        Model          :         .
         ------------  ------------   ------------   ------------
         |          | |          |   |          |   |          |
         | Controller | | Controller |   | Controller |   | Controller |
         |          | |          |   |          |   |          |
         ------------  ------------   ------------   ------------
            :            .       .            :              :
            :            .        .    Device   :              :
            :            .         . Configuration :              :
            :            .          .   Model      :              :
         ---------    ---------  ---------    ---------      ---------
         | Network |  | Network | | Network |  | Network |    | Network |
         | Element |  | Element | | Element |  | Element |    | Element |
         ---------    ---------  ---------    ---------      ---------
```

        Figure 3: An SDN Architecture with a Service Orchestrator

   Figure 3 also shows where different data models might be applied
   within the architecture.

   The split between control components that exposes a "service
   interface" is present in many figures showing extended SDN
   architectures:

o Figure 1 of [RFC7426] shows a separation of the "Application
   Plane", the "Network Services Abstraction Layer (NSAL)", and the
   "Control Plane". It marks the "Service Interface" as situated
   between the NSAL and the Control Plane.

o [RFC7491] describes an interface between an "Application Service
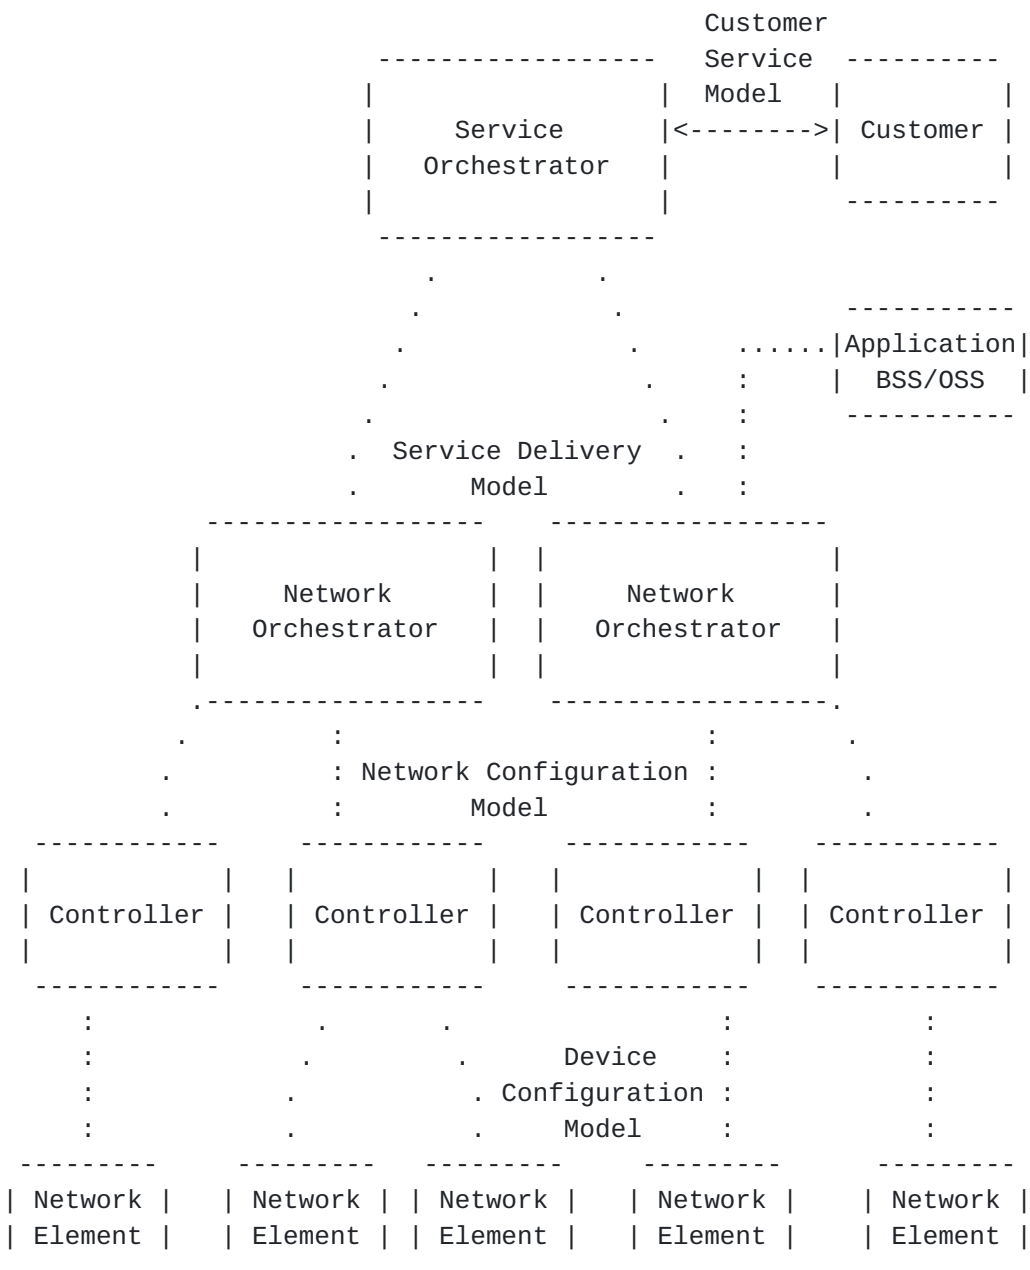   Coordinator" and an "Application-Based Network Operations
   Controller".

o Figure 1 of [I-D.ietf-netmod-yang-model-classification] shows an
   interface from an Operations Support System (OSS) or a Business
   Support System (BSS) that is expressed in "service models".

This can all lead to some confusion around the definition of a
"service interface" and a "service model". Some previous literature
considers the interface that is northbound of the Network
Orchestrator to be a "service interface" used by an application (as
shown in Figure 3), but the service described at this interface is
network-centric and is aware of many features such as topology,
technology, and operator policy. Thus, we make a distinction between
this type of service interface and the more abstract service
interface where the service is described by a service model and the
interaction is between customer and operator. Further discussion of
this point is provided in Section 5.

5. Possible Causes of Confusion

In discussing service models, there are several possible causes of
confusion:

o The services we are discussing are services provided by network
   operators to customers. This is a completely different thing to
   "Foo as a Service" (for example, Infrastructure as a Service
   (IaaS)) where a service provider offers a service at some location
   that is reached across a network. The confusion arises not only
   because of the use of the word "service", but also because network
   operators may offer value-added services as well as network
   connection services to their customers.

o Network operation is completely out of scope in the discussion of
   services between a network operator and a customer. That means
   that the customer service model does not reveal to the customer
   anything about how the network operator delivers the service. The
   model does not expose details of technology or network resources
   used to provide the service. For example, in the simple case of
   point-to-point virtual link connectivity provided by a network
   tunnel (such as an MPLS pseudowire) the network operator does not
   expose the path through the network that the tunnel follows. Of

course, this does not preclude the network operator from taking
guidance from the customer (such as to avoid routing traffic
through a particular country) or from disclosing specific details
(such as might be revealed by a route trace), but these are not
standard features of the service as described in the customer
service model.

o The network operator may use further data models (service delivery
   models) that help to describe how the service is realized in the
   network. These models might be used on the interface between the
   Service Orchestrator and the Network Orchestrator as shown in
   Figure 3 and might include many of the pieces of information from
   the customer service model alongside protocol parameters and
   device configuration information.
   [I-D.ietf-netmod-yang-model-classification] also terms these data
   models as "service models" and a comparison is provided in
   Section 6.1. It is important that the Service Orchestrator should
   be able to map from a customer service model to these service
   delivery models, but they are not the same things.

o Commercial terms are generally not a good subject for
   standardization. It is possible that some network operators will
   enhance standard customer service models to include commercial
   information, but the way this is done is likely to vary widely
   between network operators.

o Service Level Agreements (SLAs) have a high degree of overlap with
   the definition of services present in customer service models.
   Requests for specific bandwidth, for example, might be present in
   a customer service model, and agreement to deliver a service is a
   commitment to the description of the service in the customer
   service model. However, SLAs typically include a number of fine-
   grained details about how services are allowed to vary, by how
   much, and how often. SLAs are also linked to commercial terms
   with penalties and so forth, and so are also not good topics for
   standardization.

6. Comparison With Other Work

   Other work has classified YANG models, produced parallel
   architectures, and developed a range of YANG models. This section
   briefly examines that other work and shows how it fits with the
   description of service models introduced in this document.

6.1. Comparison With Network Service Models

As previously noted, [I-D.ietf-netmod-yang-model-classification]
provides a classification of YANG data models. It introduces the
term "Network Service YANG Module" to identify the type of model used
to "describe the configuration, state data and operations of an
abstract representation of a service implemented on one or multiple
network elements." These are service delivery models as described in
this document, that is, they are the models used on the interface
between the Service Orchestrator or OSS/BSS and the Network
Orchestrator as shown in Figure 3.

Figure 1 of [I-D.ietf-netmod-yang-model-classification] can be
modified to make this more clear and to add an additional example of
a Network Service YANG model as shown in Figure 4.

```
            +---------------+
            |               |
            |   Customers   |
            |               |
            +---------------+


      - - - - - - - - - - - - - - -
       Service YANG Modules

        +--------------------------+      +--------------------------+
        |                          |      | Operations and Business  |
        |   Service Orchestrator   |      |     Support Systems      |
        |                          |      |         (OSS/BSS)        |
        +--------------------------+      +--------------------------+


      - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - -
       Network Service YANG Modules

      +------------+    +-------------+    +-------------+    +-------------+
      |            |    |             |    |             |    |             |
      |  - L2VPN   |    |  - L2VPN    |    |    EVPN     |    |    L3VPN     |
      |  - VPWS    |    |  - VPLS     |    |             |    |             |
      |            |    |             |    |             |    |             |
      +------------+    +-------------+    +-------------+    +-------------+


      - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - -
       Network Element YANG Modules

      +------------+    +------------+    +-------------+    +------------+
      |            |    |            |    |             |    |            |
      |    MPLS    |    |    BGP     |    | IPv4 / IPv6 |    |  Ethernet  |
      |            |    |            |    |             |    |            |
      +------------+    +------------+    +-------------+    +------------+

         L2VPN: Layer 2 Virtual Private Network
         L3VPN: Layer 3 Virtual Private Network
         VPWS: Virtual Private Wire Service
         VPLS: Virtual Private LAN Service
```
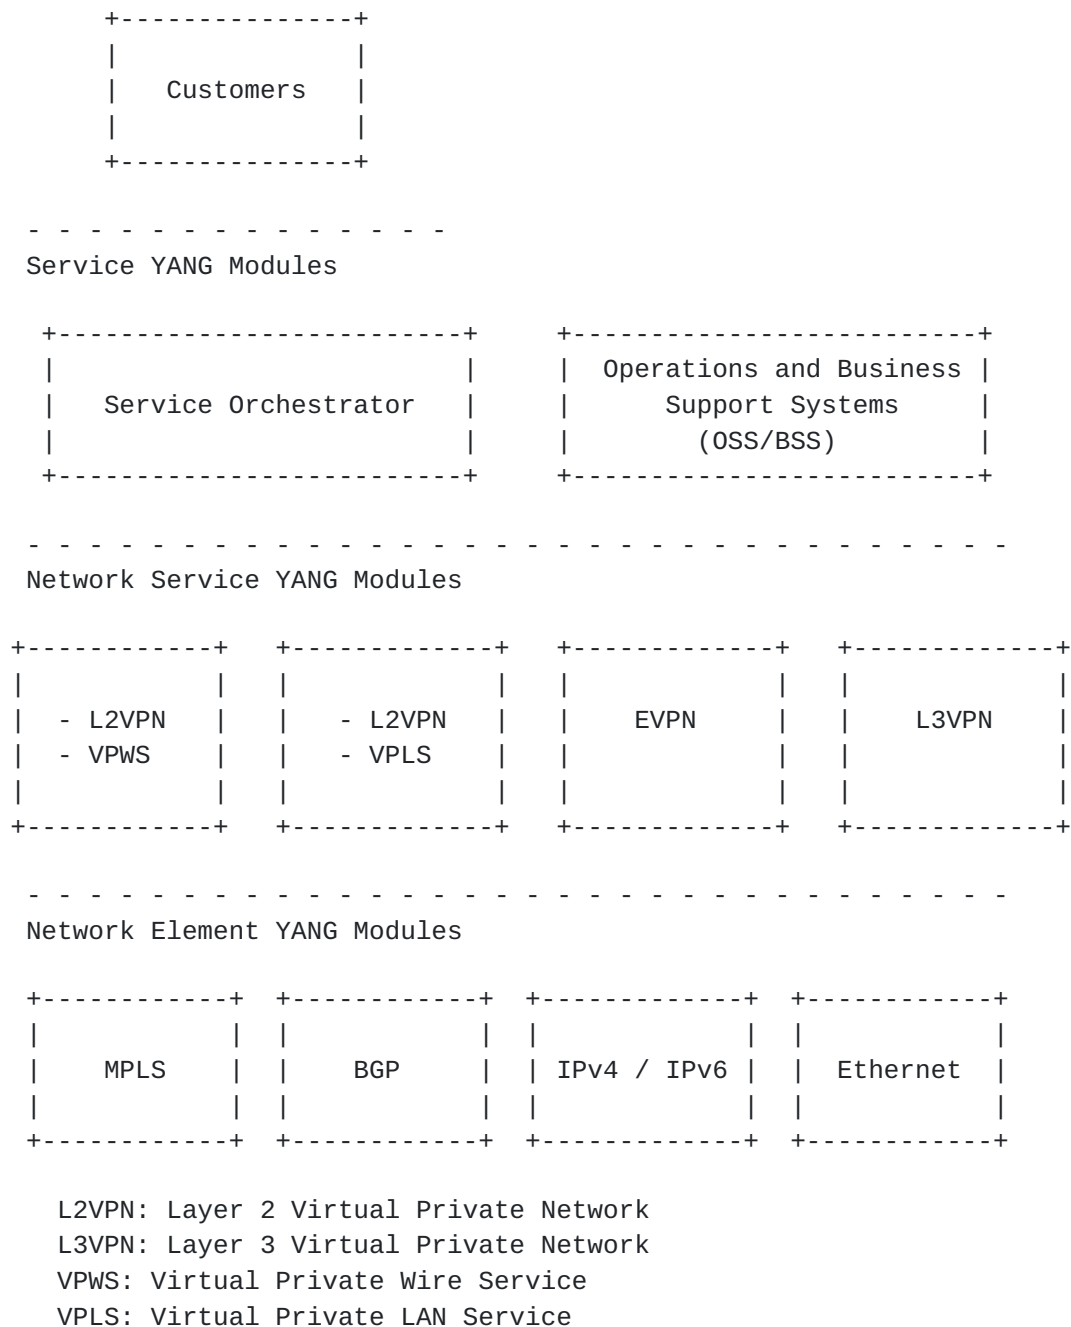
           Figure 4: YANG Module Layers Showing Service Models

6.2. Service Delivery Model Work

   A number of IETF working groups are developing YANG models related to
   services. These models focus on how the network operator configures
   the network through protocols and devices to deliver a service.

A sample set of these models is listed here:

o [I-D.dhjain-bess-bgp-l3vpn-yang] defines a YANG model that can be
   used to configure and manage BGP Layer 3 VPNs.

o [I-D.ietf-bess-l2vpn-yang] documents a YANG model that it is
   expected will be used by the management tools run by the network
   operators in order to manage and monitor the network resources
   that they use to deliver L2VPN services.

o [I-D.ietf-bess-evpn-yang] defines YANG models for delivering an
   Ethernet VPN service.

All of these models are service delivery models in the context of
this document.

6.3. Customer Service Model Work

Several initiatives within the IETF are developing customer service
models. The most advanced presents the Layer Three Virtual Private
Network (L3VPN) service as described by a network operator to a
customer. This L3VPN service model (L3SM) is documented in
[I-D.ietf-l3sm-l3vpn-service-model] where its usage is described as
in Figure 5 which is reproduced from that document. As can be seen,
the L3SM is a customer service model as described in this document.

```
        L3VPN-SVC |
          MODEL   |
                  |
           +-----------------+          +-----+
           |   Orchestration | < --- > | OSS |
           +-----------------+          +-----+
                 |         |
         +----------------+   |
         | Config manager |   |
         +----------------+   |
                 |         |
                 | Netconf/CLI ...
                 |         |
       +-----------------------------------------------+
                     Network
```
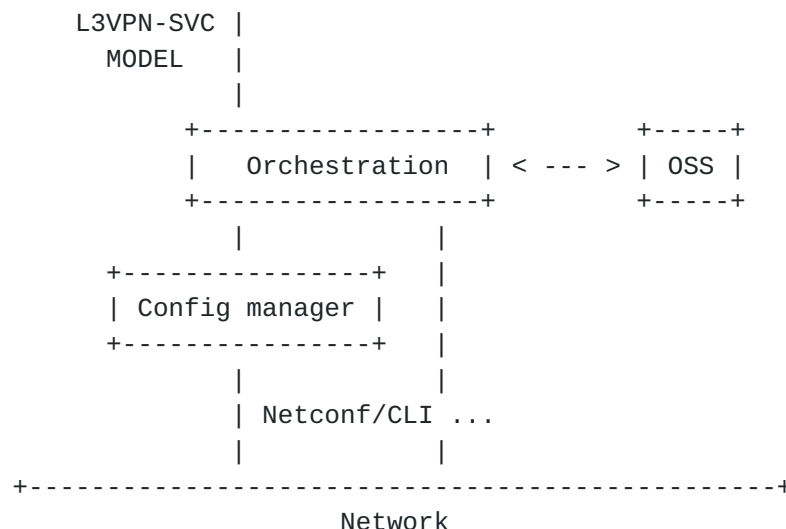
Figure 5: The L3SM Service Architecture

A Layer Two VPN service model (L2SM) is defined in
[I-D.wen-l2sm-l2vpn-service-model]. That model's usage is described

as in Figure 6 which is a reproduction of Figure 5 from that
document. As can be seen, the L2SM is a customer service model as
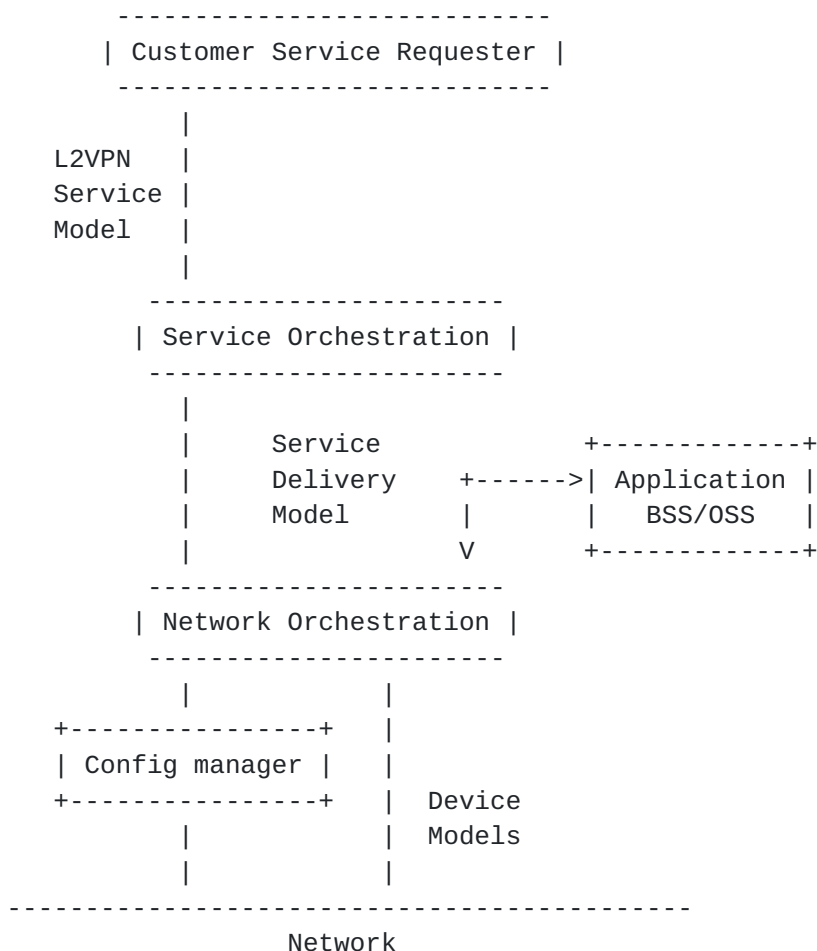described in this document.

```
            -----------------------------
           | Customer Service Requester |
            -----------------------------
                  |
       L2VPN      |
       Service    |
       Model      |
                  |
            -----------------------
           | Service Orchestration |
            -----------------------
               |
               |      Service              +-------------+
               |      Delivery    +------>| Application |
               |      Model       |       |   BSS/OSS   |
               |                  V        +-------------+
            -----------------------
           | Network Orchestration |
            -----------------------
               |           |
      +----------------+   |
      | Config manager |   |
      +----------------+   |  Device
               |           |  Models
               |           |
       ---------------------------------------------
                     Network
```

Figure 6: The L2SM Service Architecture

6.4. The MEF Architecture

The MEF has developed an architecture for network management and
operation. It is documented as the Lifecycle Service Orchestration
(LSO) Reference Architecture and illustrated in Figure 2 of [MEF-55].
Part of this is depicted in Figure 7.

The MEF terms the functional interfaces between components as
Management Interface Reference Points (MIRPs). These are the logical
points of interaction between functional management entities and are
further defined by Interface Profiles and implemented by APIs. As

can be seen from Figure 7, each of these MIRPs has been given a name.
In the context of this document:

o ALEGRO is the MIRP on which customer service models are used

o PRESTO is the MRP on which service delivery models are used.
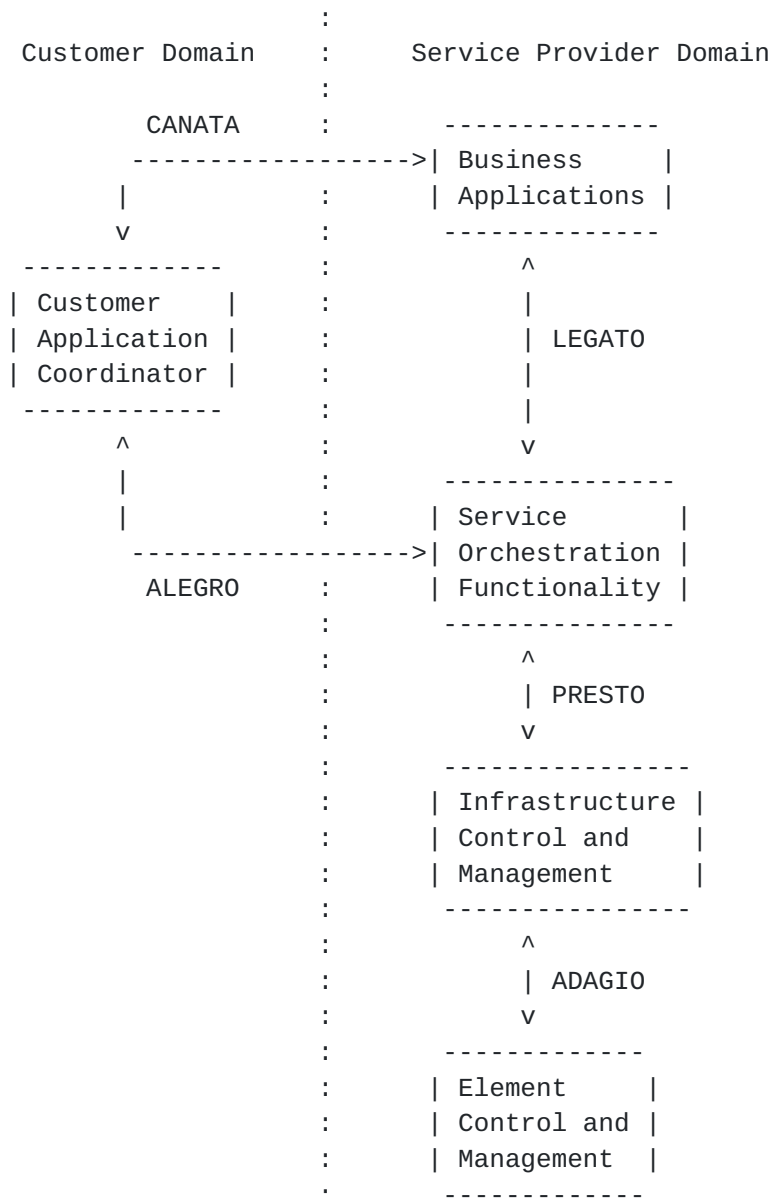
```
                         :
     Customer Domain     :       Service Provider Domain
                         :
             CANATA      :           --------------
             ----------------->| Business    |
         |               :        | Applications |
         v               :        --------------
   -------------         :               ^
   | Customer   |        :               |
   | Application|        :               | LEGATO
   | Coordinator|        :               |
   -------------         :               |
         ^               :               v
         |               :        ---------------
         |               :        | Service      |
      ----------------->| Orchestration |
         ALEGRO          :        | Functionality |
                         :        ---------------
                         :               ^
                         :               | PRESTO
                         :               v
                         :        ----------------
                         :        | Infrastructure |
                         :        | Control and    |
                         :        | Management     |
                         :        ----------------
                         :               ^
                         :               | ADAGIO
                         :               v
                         :        -------------
                         :        | Element     |
                         :        | Control and |
                         :        | Management  |
                         :        -------------
```

Figure 7: Part of The MEF's LSO Reference Architecture

7. Further Concepts

   This section introduces a few further, more advanced concepts

7.1. Technology Agnostic

   Service models should generally be technology agnostic. That is to
   say, the customer should not care how the service is provided so long
   as the service is delivered.

   However, some technologies reach the customer site and make a
   definition to the type of service delivered. Such features do need
   to be described in the service model.

   Two examples are:

   o The data passed between customer equipment and network operator
      equipment will be encapsulated in a specific way, and that data
      plane type forms part of the service.

   o Protocols that are run between customer equipment and network
      operator equipment (for example, Operations, Administration, and
      Maintenance protocols, or protocols for exchanging routing
      information) need to be selected and configured as part of the
      service description.

7.2. Relationship to Policy

   Policy appears as a crucial function in many places during network
   orchestration. A Service Orchestrator will, for example, apply the
   network operator's policies to determine how to provide a service for
   a particular customer (possibly considering commercial terms).
   However, the policies within a service model are limited to those
   over which a customer has direct influence, but which are acted on by
   the network operator.

   The policies that express desired behavior of services on occurrence
   of specific events are close to SLA definitions: they should only be
   included in the base service model where they are common to all
   network operators' offerings. Policies that describe who at a
   customer may request or modify services (that is, authorization) are
   close to commercial terms: they, too, should only be included in the
   base service model where they are common to all network operators'
   offerings.

   Nevertheless, policy is so important that all service models should
   be designed to be easily extensible to allow policy components to be
   added and associated with services as needed.

7.3. **Operator-Specific Features**

   When work in Layer Three Virtual Private Network Service Model (L3SM)
   was started, there was some doubt as to whether network operators
   would be able to agree on a common description of the services that
   they offer to their customers because, in a competitive environment,
   each markets the services in a different way with different
   additional features. Thus, when a basic description of the core
   service is agreed and documented in a service model, it is important
   that that model should be easily extended or augmented by each
   network operator so that the standardized model can be used in a
   common way and only the operator- specific features vary from one
   environment to another.

7.4. **Supporting Multiple Services**

   Network operators will, in general, offer many different services to
   their customers. Each would normally be the subject of a separate
   service model.

   It is an implementation and deployment choice whether all service
   models are processed by a single Service Orchestrator that can
   coordinate between the different services, or whether each service
   model is handled by a specialized Service Orchestrator able to
   provide tuned behavior for a specific service.

8. **Security Considerations**

   TBD

9. **Manageability Considerations**

   TBD

10. **IANA Considerations**

   This document makes no requests for IANA action

11. **Acknowledgements**

   Thanks to Daniel King, Xian Zhang, and Michael Scharf for useful
   review and comments.

12. **References**

12.1. Normative References

   [I-D.ietf-l3sm-l3vpn-service-model]
              Litkowski, S., Shakir, R., Tomotaki, L., Ogaki, K., and K.
              D'Souza, "YANG Data Model for L3VPN service delivery",
              draft-ietf-l3sm-l3vpn-service-model-12 (work in progress),
              July 2016.

   [I-D.ietf-netmod-yang-model-classification]
              Bogdanovic, D., Claise, B., and C. Moberg, "YANG Module
              Classification", draft-ietf-netmod-yang-model-
              classification-02 (work in progress), June 2016.

   [RFC3444] Pras, A. and J. Schoenwaelder, "On the Difference between
              Information Models and Data Models", RFC 3444,
              DOI 10.17487/RFC3444, January 2003,
              <http://www.rfc-editor.org/info/rfc3444>.

   [RFC7426] Haleplidis, E., Ed., Pentikousis, K., Ed., Denazis, S.,
              Hadi Salim, J., Meyer, D., and O. Koufopavlou, "Software-
              Defined Networking (SDN): Layers and Architecture
              Terminology", RFC 7426, DOI 10.17487/RFC7426, January
              2015, <http://www.rfc-editor.org/info/rfc7426>.

12.2. Informative References

   [I-D.dhjain-bess-bgp-l3vpn-yang]
              Jain, D., Patel, K., Brissette, P., Li, Z., Zhuang, S.,
              Liu, X., Haas, J., Esale, S., and B. Wen, "Yang Data Model
              for BGP/MPLS L3 VPNs", draft-dhjain-bess-bgp-l3vpn-yang-02
              (work in progress), August 2016.

   [I-D.ietf-bess-evpn-yang]
              Brissette, P., Sajassi, A., Shah, H., Li, Z.,
              Tiruveedhula, K., Hussain, I., and J. Rabadan, "Yang Data
              Model for EVPN", draft-ietf-bess-evpn-yang-01 (work in
              progress), July 2016.

   [I-D.ietf-bess-l2vpn-yang]
              Shah, H., Brissette, P., Chen, I., Hussain, I., and B.
              Wen, "YANG Data Model for MPLS-based L2VPN", draft-ietf-
              bess-l2vpn-yang-00 (work in progress), June 2016.

   [I-D.ietf-netconf-restconf]
              Bierman, A., Bjorklund, M., and K. Watsen, "RESTCONF
              Protocol", draft-ietf-netconf-restconf-16 (work in
              progress), August 2016.

   [I-D.wen-l2sm-l2vpn-service-model]
              Wen, B., Fioccola, G., Xie, C., and L. Jalil, "A YANG Data
              Model for L2VPN Service Delivery", draft-wen-l2sm-l2vpn-
              service-model-00 (work in progress), September 2016.

   [MEF-55]   MEF Forum, "Service Operations Specification MEF 55 :
              Lifecycle Service Orchestration (LSO) Reference
              Architecture and Framework", March 2016.

   [RFC6020] Bjorklund, M., Ed., "YANG - A Data Modeling Language for
              the Network Configuration Protocol (NETCONF)", RFC 6020,
              DOI 10.17487/RFC6020, October 2010,
              <http://www.rfc-editor.org/info/rfc6020>.

   [RFC6241] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed.,
              and A. Bierman, Ed., "Network Configuration Protocol
              (NETCONF)", RFC 6241, DOI 10.17487/RFC6241, June 2011,
              <http://www.rfc-editor.org/info/rfc6241>.

   [RFC7223] Bjorklund, M., "A YANG Data Model for Interface
              Management", RFC 7223, DOI 10.17487/RFC7223, May 2014,
              <http://www.rfc-editor.org/info/rfc7223>.

   [RFC7407] Bjorklund, M. and J. Schoenwaelder, "A YANG Data Model for
              SNMP Configuration", RFC 7407, DOI 10.17487/RFC7407,
              December 2014, <http://www.rfc-editor.org/info/rfc7407>.

   [RFC7491] King, D. and A. Farrel, "A PCE-Based Architecture for
              Application-Based Network Operations", RFC 7491,
              DOI 10.17487/RFC7491, March 2015,
              <http://www.rfc-editor.org/info/rfc7491>.

Authors' Addresses

   Qin Wu
   Huawei Technologies

   Email: bill.wu@huawei.com

   Will Liu
   Huawei Technologies

   Email: liushucheng@huawei.com

Adrian Farrel
Juniper Networks

Email: adrian@olddog.co.uk