

OPEN NETWORKING
FOUNDATION

SDN architecture

Issue 1.1 (draft PA8)
2015

Abstract

This document specifies the architecture of software defined networking (SDN). Issue 1.1 extends SDN architecture issue 1 in light of further work in the industry and in the ONF architecture working group. It also clarifies a number of topics in light of experience with issue 1.



ONF Document Type: TR (Technical Reference), non-normative, type 2

ONF Document Name: SDN Architecture 1.1

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
2275 E. Bayshore Road, Suite 103, Palo Alto, CA 94303
www.opennetworking.org

©2014 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.

1	Introduction	6
2	Executive summary.....	7
3	Scope	9
4	Definitions.....	11
4.1	Abstraction.....	11
4.2	Client.....	11
4.3	Client context	11
4.4	Domain.....	11
4.5	Management-control continuum (MCC)	11
4.6	Orchestration	11
4.7	Policy	11
4.8	Recursion.....	12
4.9	Resource.....	12
4.10	Server	12
4.11	Server context.....	12
4.12	Service	12
4.13	Service context	12
4.14	Virtualization	12
5	Concepts.....	12
5.1	Principles of SDN.....	12
5.2	Roles.....	15
5.3	Service and resource oriented models.....	16
5.4	Primitives	19
5.5	Controllers and planes	21
6	SDN controller	22
6.1	SDN controller as feedback node	23
6.2	Orchestration	24
6.3	Virtualization	25
6.4	Resource sharing.....	26
6.5	Delegation.....	27
6.6	Client context	27
6.7	Service context	29
6.8	Server context.....	30
7	Applications.....	31
8	Putting it together: the integrated architecture	32
8.1	Interfaces	34
8.2	Notifications	36

8.3	Peer controllers.....	37
9	Specific perspectives on the architecture	37
9.1	Security.....	37
9.2	Reliability, availability	38
9.3	Identifiers	39
9.4	Realization considerations	40
9.5	Initialization	40
9.6	Complexity	41
9.7	Persistence	42
9.8	Migration and coexistence	43
9.9	Relationship of SDN and NFV.....	43
Appendix A.	Appendices	43
A.1.	Discussion of definitions.....	44
A.1.1.	Abstraction.....	44
A.1.2.	Client.....	44
A.1.3.	Client context	44
A.1.4.	Domain	44
A.1.5.	Management-control continuum (MCC)	45
A.1.6.	Orchestration	45
A.1.7.	Policy	45
A.1.8.	Recursion.....	46
A.1.9.	Resource	47
A.1.10.	Server.....	48
A.1.11.	Server context	48
A.1.12.	Service	48
A.1.13.	Service context.....	48
A.1.14.	Virtualization.....	49
A.2.	Observations about models	49
A.3.	Multiple logins	49
A.4.	Evolution from issue 1 to issue 1.1	53
Appendix B.	Back matter	54
A.5.	Acronyms	54
A.6.	References.....	55
A.7.	Contributors	55

List of Figures

Figure 1 – Basic model	7
Figure 2 – Core of the SDN architecture.....	8
Figure 3 – Administrator role.....	15
Figure 4 – User and provider roles.....	16

Figure 5 – Control as feedback.....	23
Figure 6 – Client context	28
Figure 7 – Server context.....	31
Figure 8 – SDN controller illustrating contexts	33
Figure 9 – Supplementary function example.....	34
Figure 10 – Peers as symmetric requestors and providers.....	37
Figure 11 – Multiple users of a client context	50
Figure 12 – Separated SDN controller Green	51
Figure 13 – Issue 1 architecture.....	53

1 Introduction

This architecture issue 1.1 is a stand-alone document that clarifies and extends issue 1 [1], but does not obsolete it. Clause 10.4 describes and explains the differences.

A great deal of work on SDN has already been done, and continues, in a number of ONF working groups, standards development organizations (SDOs), and open-source communities. Even if it were possible, the architecture would not attempt to invalidate existing work. However, it does aspire to unify the discussion across these groups.

Except for illustrative examples, this architecture intentionally remains abstract and avoids details of target technologies. Additional documents, existing and potential, expand the architecture into focused areas. These include:

- TR-xxx, SDN architecture for transport [ref to onf2014.301 when published] [could be ITU-T G.astdn]
- TR-518, Relationship of SDN and NFV [2]
- <designator>, Intent NBI – Principles and information model [ref to onf2015.327 when published]

Document structure

Ed: revise this when we converge on a final outline

Clause 2 is an executive summary, an abbreviated statement of the essentials. Subsequent clauses build up the terminology and concepts in detail, then use them to support the overall architecture.

Clause 3 describes the scope and purpose of the architecture.

Clause 4 defines a number of key terms.

Clause 5 identifies some of the key concepts of SDN and expands on their interpretation.

Clause 5.5 discusses the SDN controller in some detail.

Clause 7 describes applications, in less detail. Maximum freedom is allowed by intentionally underspecifying their internals and functions.

Clause 8 is an integrated view of SDN interfaces.

Clause 9 views the architecture from a number of particular angles.

Clause 10 includes appendices whose content warrants inclusion in the main architecture text, although not inline. Other material to clarify and expand on inline text may be published in the form of separate white papers or architecture notes.

Finally, clause 11 lists acronyms, references, and contributors.

2 Executive summary

An architecture is a necessarily incomplete collection of perspectives over a set of underlying ideas. A consistent architecture reveals no contradictions when viewed from any of these perspectives. A useful architecture facilitates productive development of concepts into realities. An open architecture minimizes the difficulty of extension in previously unforeseen directions. The architecture described in this document aspires to be consistent, useful, and open.

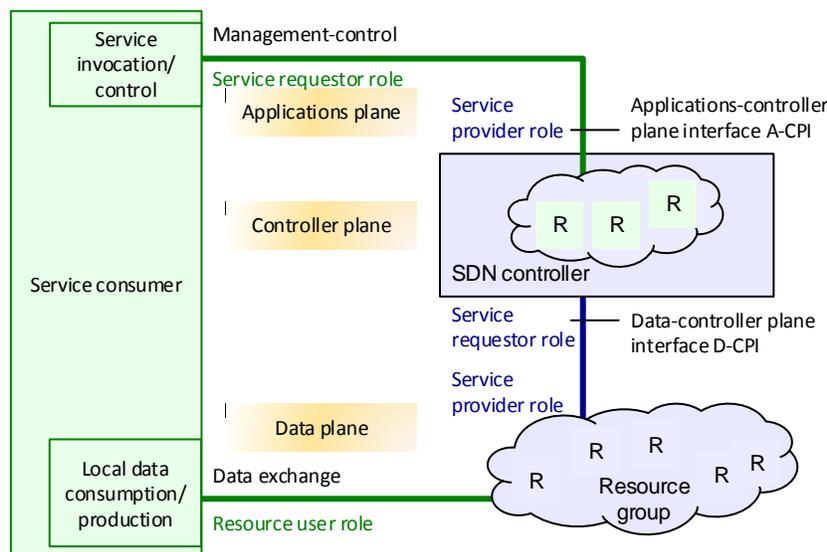


Figure 1 – Basic model

Figure 1 illustrates the basic model of SDN, that of a service consumer (client, user, customer,) Green, who exchanges both data and management-control operations with some SDN server or provider, Blue. Although user data is ultimately forwarded or processed by some set of resources (R) that are owned by Blue, Green controls its service via a management-control association by invoking actions on a set of virtual resources (R) that it perceives to be its own. Among other responsibilities, the SDN controller virtualizes and orchestrates the Green resource and service view onto its own underlying Blue resources and services. The concepts of resources and services are intentionally unbounded.

The SDN architecture extends the basic model and clarifies its implications. Key extensions include sharing resources a) among multiple clients, b) dynamically, c) in an optimum way. Other essentials of a complete architecture include management in the classical sense, both of network resources and of services.

This architecture usually portrays client and server as existing in separate business domains, illustrated with separate colors. The reason for this is to emphasize the need for traffic isolation, information hiding, security and policy enforcement at interface points. Depending on the relationship between client and server, visibility and security criteria may be strict or relaxed in any particular deployment situation, including full transparency when appropriate.

SDN controller relationships

The central entity in an SDN is the SDN controller. Figure 2 illustrates some of its key functions and interfaces.

SDN is modelled as a set of client-server relationships between SDN controllers and other entities that may themselves be SDN controllers. In its role as a server, an SDN controller may offer services to any number of clients, while an SDN controller acting as client may invoke services from any number of servers. As long as they exhibit appropriate interface behavior, the details of entities that are not SDN controllers are beyond the scope of the architecture.

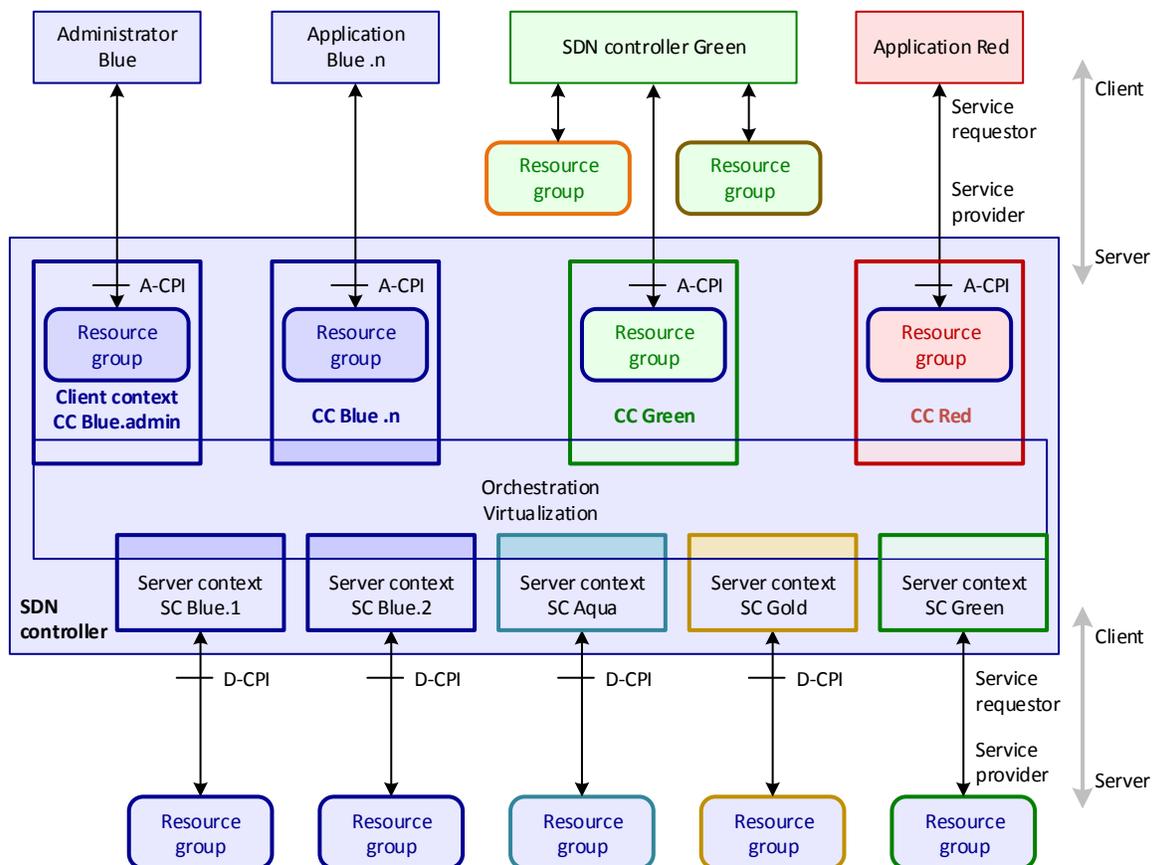


Figure 2 – Core of the SDN architecture

The architecture recognizes dual perspectives on the nature of client-server interfaces. The services perspective is particularly appropriate from a top-down or customer-provider viewpoint. The resources perspective is particularly appropriate from the bottom-up viewpoint of a resource owner, especially an internal administrator. The perspectives are complementary, but emphasize different things. The construction of views and mappings on a common underlying information model helps tie these perspectives together.

The SDN controller satisfies client requests by virtualizing and orchestrating its underlying resources. As the network environment changes and client demands change, the SDN controller is responsible for continuously updating network and service state toward a policy-based optimum configuration.

An SDN controller exposes services and resources to clients via applications-controller plane interfaces (A-CPIs), and consumes underlying services and resources via data-controller plane interfaces (D-CPIs) (note). Each of these interfaces is a reference point for information hiding, traffic and namespace isolation, and policy enforcement.

Management and control are viewed as a continuum, in which an administrator role differs from that of ordinary applications by having greater scope and privilege. The administrator has authority to configure the SDN controller itself, along with client and server contexts.

Recursion may be deduced from figure 1 by recognizing the repetitive service requestor and service provider roles. SDN controllers may associate with other SDN controllers and non-SDN management-control entities in hierarchical or peer arrangements, within or across administrative domains. This permits SDN to span the range from end-user service negotiation to world-wide service provisioning to finely granular control of individual network elements.

3 Scope

Software-defined networking (SDN) applies the flexibility of software to the entirety of the networking space. It includes unlimited numbers of business relationships, geography spanning the world, and everything from end-user service negotiation and delivery to the planning, installation, provisioning and maintenance of network infrastructure. As well as forwarding traffic, an SDN may process traffic, either as part of added-value services or for service and network maintenance purposes.

The overall architecture is intended to span the entirety of the space. Nevertheless, it is structured for easy subsetting, for example into carrier, cloud, campus or enterprise environments that purchase much of their network service from third parties, across opaque or semi-opaque business domain boundaries.

SDN is based on three principles, which are further explored in clause 5.1.

- Decoupling of traffic forwarding and processing from control
- Centralized control
- Programmability of network services

The major components of an SDN are resources and controllers. SDN controllers mediate between clients and resources to deliver services. Most resources are related to traffic in some way, but support resources are also recognized, for example security credentials and notification subscriptions.

The primitive roles in an SDN are those of administrator, service requestor and provider, and resource user. Additional roles are not precluded.

This architecture document expands the meaning and implications of the principles, the required and optional functions of the components, and the interfaces, rights and responsibilities of the roles.

Architecture goals

The overall goal of the architecture is to assist providers in better serving their customers – in any number of dimensions – while reducing their own cost to deliver those services. SDN addresses all aspects of this goal, a few examples of which are:

- An environment that reduces the time and cost of developing new services
- Flexible definition and availability of resources, including virtual network functions (VNFs)
- On demand assembly of resources into services
- More efficient resource loading
- Facilitating global semantic agreement with a common information model
- Merging traditional BSS/OSS functions with control

Because of its wide scope, not all aspects of SDN can be addressed in depth immediately. The architecture is intended to create a common understanding that facilitates parallel development as cost-benefit opportunities arise across the broadly defined problem space. While alignment on principles is important, the single most important underpinning of consistency is arguably a common information model.

The architecture describes principles, components and roles in abstract ways. Any number of implementations can therefore claim to comply with the architecture. Rather than as a vehicle for compliance statements, the architecture may better be used as a reference to which an implementation can be compared. Many open-source projects implement aspects of SDN; the architecture helps understand the gaps. Gaps are not necessarily deficiencies; partial solutions may be completely adequate for their target markets.

One of the long-standing problems of the industry has been silos, separate areas with separate expertise, separate staffing, separate business, service, and investment considerations. While the classic voice-data silo partition is finally gone, or nearly so, other silos persist. A few examples:

- Transport vs packet
- Wireline vs wireless
- Local vs long-haul (access vs core)

And new silos:

- Data center, cloud vs dedicated and dispersed physical elements
- NFV vs SDN

There will continue to be important differences along these, and other, dimensions. However, SDN ought never provide a technical justification for the perpetuation or creation of silos. The goal of the architecture is to open possibilities, and especially to expose common ground, such that silos can be collapsed whenever and however it makes sense.

4 Definitions

A good definition is concise, correct and pedantic, avoids direct or indirect self-reference, and aligns well with at least one of the understandings in common usage. If multiple understandings exist in common usage, a definition must take a position, rather than sanctioning confusion.

Further, a good definition should be narrow enough and crisp enough to exclude invalid candidates. Fuzzy outer bounds are a major flaw in many definitions in the industry.

The implications of a definition are often important, but not obvious. A definition may therefore require informative material that is not essential to the definition itself, but that assists in its understanding. Informal descriptions of the category and qualifiers may be helpful. Examples are often useful, both of candidates that do, and also of candidates that do not, satisfy the criteria. For brevity in the main text, these discussions appear in clause 10.1.

4.1 Abstraction

Definition: The representation of an entity or group of entities according to some criterion, while ignoring aspects that do not match the criterion.

4.2 Client

Definition: An entity that receives services from a server.

4.3 Client context

Definition: The conceptual component of a server that represents all information about a given client and is responsible for participation in active server-client management-control operations.

4.4 Domain

Definition: A grouping of entities according to some criterion.

4.5 Management-control continuum (MCC)

Definition: The principle that the functions of management and of control are largely, if not entirely, the same.

Note – In keeping with SDN convention, this document mostly uses the term *control*, and refers to one of the key entities as an SDN *controller*.

4.6 Orchestration

Definition: The ongoing selection and use of resources by a server to satisfy client demands according to optimization criteria.

4.7 Policy

Definition: A rule that guides and constrains subsequent actions of, and interactions between, other parties.

4.8 Recursion

Definition: The repeated application of a process in which the input to each iteration except the first is derived from the output of the previous iteration.

4.9 Resource

Definition: Anything that can be used to deliver a service.

4.10 Server

Definition: An entity that provides services to a client.

4.11 Server context

Definition: The conceptual component of a client that represents all information about a given server and is responsible for participation in active server-client management-control operations.

4.12 Service

Definition: The delivery of value for some time interval by a provider to a consumer.

4.13 Service context

Definition: The conceptual component of a client context that represents all information about a given service.

4.14 Virtualization

Definition: An abstraction of underlying resources on a server, whose criterion is allocation of those resources to a particular client, application, or service.

5 Concepts

5.1 Principles of SDN

SDN is based on three principles, as described below. All require interpretation.

5.1.1 Decoupling of traffic forwarding and processing from control

The purpose of this principle is to permit independent deployment of control and traffic forwarding and processing entities. This principle is not obviously a game changer per se. Transport equipment, for example, has long decoupled control from forwarding and traffic processing. However, decoupling is a necessary precondition of centralized control and of recursion, specifically hierarchical recursion. Decoupling also allows for separate optimization of platform technology and life cycles.

The architecture reflects the decoupling principle in the form of an entity called the SDN controller, which has management-control responsibility for some set of resources. The resources are

considered to be in a data plane, so-called because most of them are directly or indirectly associated with processing client traffic. Recursive decoupling is embodied in the idea that a client has management-control access to a set of resources and services that is exposed by an SDN controller.

Further reading: open interfaces, client context, control as feedback, recursion, delegation, resources and resource groups

5.1.2 Centralized control

Decoupling of traffic processing from control is a precondition of centralized control. The centralized control principle asserts that resources can be used more efficiently when viewed from a wider perspective. A centralized SDN controller can orchestrate resources that span a number of subordinate entities, and thereby offer better abstractions to its clients than if it could only abstract subsets of individual subordinate entities. The best example of this is the exposure of a single monolithic forwarding domain that is built atop an arbitrarily large and complex underlying network.

However, a number of factors argue against a single monolithic control point for the entire global telecoms network.

- Scale is the obvious reason, along with related factors such as sheer propagation delay. The optimal choice among a large number of options may not justify the increased complexity, as compared with a slightly suboptimal choice among a smaller set of options.
- Management-control information exchange is highly constrained across trust boundaries. By separating SDN controllers along (at least) trust domain boundaries, the architecture exposes the necessary security and policy enforcement points.
- It is necessary that SDN co-exist with non-SDN technology in all possible ways: as controlling or controlled entities or as peers, within and across administrative, technology and other domains, responsible for greater or lesser functionality.
- Resources can themselves be active, in a number of ways and for a number of reasons. Local control loops ought not be centralized without justification.
- A more subtle reason is that management-control communication is distributed around the network on a network of its own, which itself requires management-control. Domain partitioning is likely to be a better solution than an all-inclusive network attempting to control itself.

The principle of centralized control is best understood as a recommendation to consider cost-benefit trade-offs of centralization.

Further reading: complexity, control as feedback

5.1.3 Programmability of network services

This principle permits a client to exchange information with an SDN controller, either by discovery or negotiation prior to the establishment of a service, or during the lifetime of a service according to changes in client needs or the state of the client's virtual resources. Agility is the pay-

back, both in negotiating a satisfactory service and in turning it up. This is tied into the ability of SDN to dynamically tap a wide domain of existing resources, or to create new resources on demand, especially virtual network functions (VNFs).

The programmability principle is based on the premise that service-consuming applications and resources benefit from collaborating in detail in negotiating and delivering services, even on a continuing basis. This may be true of some applications. At the other end of a continuum, the intent interface model expects the client to express its desired outcome – possibly in considerable detail, possibly after considerable negotiation – but then leave the realization and real-time optimization to the SDN controller.

The merit of the intent model is that it disentangles resource operation from the client's purpose, allowing for simpler clients that may be more independent of the infrastructure. The client trades off simplicity for participation in fine tuning or continuing optimization, while if the ongoing service cannot be delivered as committed, the feedback is more likely to be an abrupt failure indication.

Further reading: client context, multiple logins, service and resource models

5.1.4 Open interfaces

This fourth principle concerns SDN implementation and deployment, not the fundamentals of the architecture. It presupposes the well-defined partition of functions and interfaces, and specifies that the interfaces be public and open to community definition. The purpose of this principle is to encourage competition. The desired competition is most likely to emerge in the support of completely mainstream features on completely mainstream white box hardware.

In other circumstances, open interfaces are of less value, and it may be preferable for both sides of an interface to be provided by the same vendor.

- When proprietary hardware offers features that outweigh the disadvantage of proprietary interfaces. Such features may be supported with proprietary extensions of standard interfaces.
- When the market for a given hardware class is so small that few competitors find it worthwhile.
- When the data-plane functionality is implemented in software, e.g., as a virtual network function VNF.

Interface standardization in the service-consuming applications space is difficult, not only because the scope of new and innovative services is intentionally unbounded, but because the best solutions may only be apparent in retrospect. The experimental approach of the open-source community may be able to evolve best practices more effectively than would be possible through formal standardization.

It may be considered that the principle is better thought of as open integration, rather than open interfaces. Whatever the nature of the functionality and interfaces, integral fit into an industry information model is key.

5.2 Roles

The roles described in this clause are useful to the SDN architecture, but it is not implied that no other roles exist.

5.2.1 Administrator role

An administrator is characterized by having greater visibility and privilege than an ordinary client. Normally, an administrator would be a trusted employee of the same organization that operates the SDN controller. The administrator's responsibility is to create an environment that can offer services, to modify the environment from time to time, to monitor the environment for proper operation, and to act on exceptions beyond the ability of the environment to resolve internally.

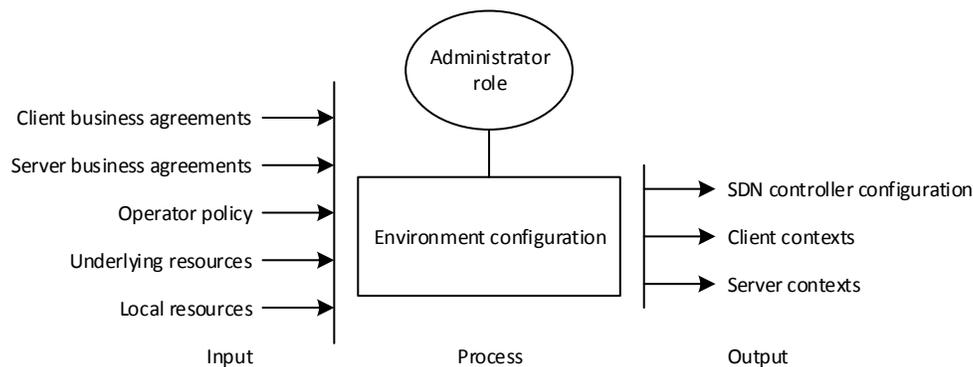


Figure 3 – Administrator role

Figure 3 shows that the administrator configures SDN controllers and the necessary contexts.

The administrator creates a working SDN controller from some substrate (creation of an SDN controller, possibly as a VNF, downloading its code, etc, is beyond the scope of the current discussion). The SDN controller is created by default with an administrator client context with unrestricted visibility and authority to perform all other operations. The administrator configures the controller with server contexts to expose underlying resources, and updates them from time to time as needed. The underlying resources are themselves configured by their own administrators at their own (underlying) levels.

The administrator then creates client contexts for each of its clients, which includes allocation of underlying resources to clients in the process called virtualization, as well as supplementary configuration. The administrator configures each client context with policy that defines the actions and bounds permitted to the client. An administrator may modify a client context during its lifetime, and may destroy a client context if the client relationship terminates.

An SDN controller is responsible for continuously optimizing its use of resources according to a global optimization policy. The administrator installs and modifies the optimization policy as needed. An administrator may also be expected to create subscriptions and logs on its own behalf.

Further reading: client context, server context, controller as feedback node, realization considerations

5.2.2 User and provider roles

Clients or service consumers satisfy their needs by requesting services from SDN controllers as providers, and they achieve their data transfer and data processing objectives as users of those resources (figure 1).

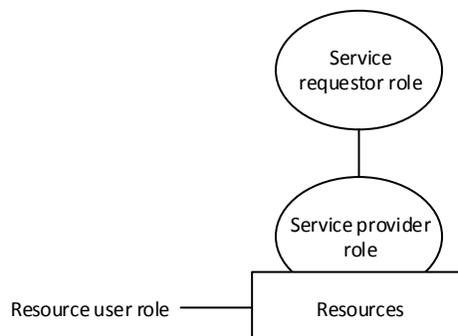


Figure 4 – User and provider roles

As illustrated in figure 4, the resources explicitly or implicitly available to a client play a service provider role, offering services to a service requestor for configuration. The service requestor represents the management-control aspect of the client in setting up the desired service; it may also subscribe to notifications from the provider. A given service requestor may invoke and manage-control any number of simultaneous services.

While the service requestor represents the setup of service by a client, the resource user role represents the client's use of the according resources to satisfy its service needs. Usually this is accomplished by exchanges in the data plane. The services offered by the provider are constrained only by the nature and functions of the available resources and the agreement between server and client.

Further reading: client context, server context

5.3 Service and resource oriented models

As described in clause 2, a service-oriented model is especially appropriate from the top-down viewpoint of a customer, while a resource-oriented model is appropriate from the bottom-up viewpoint of an operator, and especially from the viewpoint of an operator's administrator.

5.3.1 Service-oriented model

From the service-oriented perspective, the basic operation across an SDN interface is service invocation and management. According to this model, a client requests creation, read, update or delete (CRUD) operations on all or a component of a service context object. The server is an SDN controller, which is expected to validate the request against policy and available resources, and either satisfy the request or provide an appropriate exception response.

In most cases (note), the service context governs the behavior of data-plane resources that permit the client to exchange traffic with other network entities. Examples of such service contexts range from contracts for residential telephony to corporate VPNs. At a very detailed level, a service context might represent a connection through a network.

Note – Examples of value received by a client that do not involve data-plane exchanges include information discovery and negotiation of various kinds, and directory lookups. In some such interactions, it may not be necessary to instantiate a service context object.

Service-related resources are released when the service context is deleted. A service context may persist indefinitely, i.e., until explicitly deleted, or it may terminate due to events such as signaling or the expiration of an explicit schedule. The management-control association between client and server may or may not exist continuously during the lifetime of a service context.

A service request (CRUD operation) is necessarily expressed in terms of entities, actions, and names/addresses known to the client. Internal functions of the client are not subject to architectural standardization, but it would be expected that the client retain a view of the service, as expressed in its own terms, possibly updated with state information derived by query or notification from the server.

The server necessarily retains the service context, which may include the service invocation as expressed by the client. The client's request is the desired outcome that guides the server in orchestrating and virtualizing underlying resources to satisfy the client's demand. This is a continuing process, not only as clients amend their existing services, but as network state and resource contention change over time.

The architecture does not constrain the entities and actions known to the client. In general, these client entities and actions must be mapped into entities and actions known to the SDN controller acting as server (note). In general, neither client nor server is in a position to supply complete mapping information, which must therefore be a shared responsibility.

Note – Transparent mappings are not precluded. They may be appropriate, especially for server administrators.

A number of ways may be appropriate to populate the mapping function. A few of them include:

- Offline business or technical level negotiation, e.g., of access points of presence (PoPs), with subsequent provisioning of explicit mapping equivalents into a mapping database.
- Auto-discovery, e.g., of the speed of an Ethernet PoP. Protocols such as LLDP [8] may operate across the client-server data-plane boundary.
- Publication by the server of a service catalog, or the equivalent, the schema of which may be used to express the desired service. Some such schemas may be known a priori from standards.
- Negotiation of agreement by protocol, for example through proposing and accepting inter-domain labels.

A given client may have any number of service contexts in place on a given server at any given time. The service context identifier allows for unambiguous CRUD operations on new or existing services.

A server may offer services to a number of clients. All of the server's information for a particular client is contained in a conceptual component called a client context (CC). A CC may contain any number of service contexts.

A client may orchestrate services that span one or more servers. Such a client may be modelled as an SDN controller in its own right. Orchestration of composite services requires that the client be responsible for choosing and provisioning intermediate interface points, as well as the service-specific behavior of each node. Successive partitioning is one expected way to realize a service, where a single SDN controller may perform the entire process for the subnetwork within its own domain, or a hierarchical stack of controllers may successively process successively more detailed service invocations, or anything in between. Peer controllers may invoke neighbor recursion to achieve the same ultimate goal. In all cases, the edge points must be coordinated and therefore recursively visible, while the internal details of the service are optionally visible, but are often left for the immediate controller.

The selection, connection, and provisioning of network topology and nodes shades into the resource-based perspective on SDN.

Further reading: Peering, client context, service context, additional interfaces

5.3.2 Resource-oriented model

Resources represent the things that are available for use by SDN. An underlying resource is an instance (actual or potential) of some managed object class in an information model, and is regarded by its client as a black box (note). Virtual resources exposed to clients are views of the underlying resources, with static mapping detail in the virtualization function, or dynamic mapping shared by the orchestration and virtualization functions.

Note – The internals of the black box would be visible in the underlying (server) frame of reference.

The inventory of resources is not necessarily static. Given a suitable substrate and policy, non-physical resources can often be created and deleted by SDN operations, especially as they invoke NFV capabilities. Physical resources can be created and deleted by operations such as network construction and renovation, and brought into the scope of an SDN domain via discovery or provisioning.

In the simplest resource-based model, the owner of the SDN controller (server administrator role) allocates resources to clients. In this case, the mapping process for the service perspective may be the static result of a collaboration (note) between client and server administrators. The collaboration may occur offline, e.g., as part of a business-technical negotiation, but online or real-time negotiation is not precluded.

Note – When an organization constructs a client-server relationship within its own domain, for example for scaling reasons, a single entity may play both client administrator and server administrator roles.

Virtualized client resources are views of complete, partitioned, composite, sliced or shared parts of underlying resources, with namespace and possibly address space translation. These virtual resources are things that a client can use as black-box entities in its service invocations. An example of detailed pre-allocated resources would be the leasing of an entire subnetwork by one company to another.

It may also be the case that only a few resources are mapped one-to-one to client resources. A common example of this case is the interface to a residential or business subscriber, in which on-

ly the subscriber's (client's) points of presence need to be mapped. The remainder of the server's resources are hidden from the client, including access to the E.164 telephony number space or the Internet. Underlying resources are orchestrated by the server as required, depending on network state, service demand, or traffic load. Even with client-specified class of service and traffic processing constraints, this minimal mapping allows maximum freedom to the server to orchestrate resources according to its own optimization policy.

Further reading: resource groups, client context, service context, resource data base

5.4 Primitives

5.4.1 Information model

One of the SDN value propositions asserts that the communications environment of the future will include more vendors with a wide range of product offerings, on short and unsynchronized release schedules. Especially if a service invocation goes through several APIs from several vendors before reaching a data-plane device, semantic mismatches in the information conveyed will clearly result in chaos. A common information model (note) is key. As new fragments are developed, it is important that they be integrated into the common model.

Note – RFC 3444 [9] describes the difference between an information model and a data model. Compared to information model mismatch, data model incompatibility is far less difficult to diagnose and resolve.

Purpose-specific data models and customized views can be derived from the common information model as necessary. These are comparatively easy to adapt pragmatically, as long as semantics are preserved. This allows purpose-specific APIs to be produced, while enforcing consistent semantics.

5.4.2 Resources and resource groups

Any SDN service is built upon some set of resources, whose functions and interfaces are configured to the particular need. Resources may be physical or virtual, active or passive, and in many cases, may be created, scaled, or destroyed, by or at the behest of the client or the server. The ETSI network functions virtualization (NFV) initiative [3] is an important source of SDN resources.

A resource is modeled as an instance of a managed object class in an information model. Resources can be repeatedly subdivided or combined into larger resources in any way that makes sense, either by definition of the underlying information model, by management-control configuration, or by real-time orchestration and virtualization. This includes so-called slices of resources, in which it may be desired to only abstract part of an underlying resource, for example some part of the capacity of a link.

A client sees a resource as defined by a view. A view may restrict the information available to the client, the actions that can be invoked, and the notifications that it publishes. If it is shared, the same underlying resource may be virtualized to present different views to different clients. A view may also represent information such as subdivided allocation, potential or actual existence, and potential scalability.

The resource or data plane is sometimes considered to be simply an undifferentiated cloud. However, resources naturally tend to fall into groups. Resource groups are not architecturally fundamental; they are convenient abstractions.

Resource groups may exhibit some or all of the following properties:

- Fate sharing; the likelihood that all resources in the group lose or regain connectivity to the SDN controller together, that all fail or recover together, that all are subject to backup, restoration and migration together.
- Through common security attributes, the ability for a single outer management-control association (e.g., IPsec) to serve one or more inner management-control associations.
- The ability to use relative distinguished names and addresses of resources within the context of a group, rather than more global names and addresses.
- Tightly integrated functionality within the resources of the group. Delegated or autonomous functions exemplify this property: dynamic sharing of bandwidth, local fault recovery, fault correlation, the ability to inject test or other stimuli at one point and detect the consequences locally at some other point.
- A common notifications publication and subscription scope for all resources in the group.
- A knowable upper bound on the ability of resources to be augmented or scaled without having to migrate into other groups.
- A holding space for resources that are available for use, but are not actually in use at any given time.
- A topologically significant point in a network graph, useful in optimizing path or other computations.

Further reading: client context, service context, resource-oriented model, resource data base, persistence, notifications

5.4.3 Resource data base

The resource data base (RDB) is the conceptual container for information that needs to be retained in an SDN controller beyond some small locality in process space and time (note). Among other things, the RDB is conceived as the local repository of all underlying resources. The RDB does not appear on the architectural drawings because it exists everywhere in an SDN controller, as needed to support the various active functions and interfaces.

Note – This is to exclude cached information that remains within a single function and calling parameters between functions.

An important criterion for retention of a data fragment in the RDB is the need for its availability after controller reinitialization or replacement. Large fractions of the information in client and server contexts are appropriate for the RDB, including views of underlying resources that are presented to clients, services and their resource mappings, topology and resource inventories, policies, logs, profiles, and subscriptions. The identifier mapping between underlying resources

and client resources must be retained here, as must the optimization criteria used by the orchestration function.

A second criterion for RDB storage recognizes that some information is created by one entity and may subsequently be used on a continuing basis by other entities. A prime example is the mirror of an underlying resource, conceptually contained in a server context.

Note – It must be understood that CRUD operations on a mirrored copy of a resource would have no effect on the corresponding underlying resource.

Local persistent objects may also be regarded as stored in the RDB. Examples include subscriptions and profiles. Logs and PM history would **probably not** be regarded as RDB material.

Ed: do we take a definite position?

Nothing is said about the implementation of the RDB, whether it is a formal database of some type, whether it is distributed or centralized, whether all data is represented in a mutually consistent format, etc. It is recognized that different backup and restoration regimens will be appropriate for different data fragments.

An administrator role would have unrestricted access to the entire RDB. For other users, fine-grained access policy must be enforced on data fragments in the RDB, as they relate to provider, client or server contexts. Some policy may be implicit, e.g., no client can view information in the provider's or another client's context, but it may also be necessary to explicitly define additional access policy. Being persistent, such a policy would itself be contained in the conceptual RDB.

Further reading: Persistence, Client context, Server context, Controller as feedback node, Identifiers

5.5 Controllers and planes

The conventional view of SDN is structured into planes. A data plane processes user traffic, a control or controller plane hosts SDN controller instances, and instances of service-consuming applications reside in an applications plane. While this is a good introductory model, it is less useful when conceptualizing SDN in depth.

A given SDN controller governs some set of resource groups (in a data plane), whose state it controls in a feedback loop. It virtualizes these underlying resources and exposes a customized virtual resource environment (as a data plane) to each of its clients.

In much the same way that a controller is thus sandwiched between data planes, a network resource (note) is supported by an underlying virtualization (controller plane), and is used by a client or SDN controller (controller or applications plane). Even a hardware network element exposes virtual resources to its SDN controller by way of a (possibly minimal) client context, which must have been configured by an administrator.

Note – The term *network resources* distinguishes entities that have some relation to traffic processing from support resources such as client-server security credentials, notifications subscriptions, profiles, etc. These latter resources are typically not virtual.

As to applications, the SDN architecture can say nothing about what may or may not exist as their superior entities, whether and how they serve end users directly or are themselves intermediaries of some kind.

The strongly recursive nature of the architecture reduces the conceptual value of the original SDN data, controller, and applications planes. Sometimes, management is also regarded as a plane, which does not fit well into pictorial representations of the architecture (note). As management and control are recognized as different aspects of the same thing, roles are a more useful way to model the relationships.

Note – Parallel planes are meaningful and useful in terms of geometric analogy. Each plane is independent and sees only the planes above and below itself. The geometric implications of intersecting planes (i.e., lines) convey no useful analogies.

A final objection to planes as a fundamental concept is that everything of interest exists in instances. It can be argued that amorphous plane models have obscured important aspects of the problem space, such as controller-controller peer interactions and resource grouping.

For these reasons, this SDN architecture document uses the concept of planes occasionally, descriptively, informally, colloquially, but not rigorously.

Further reading: recursion, client context, roles, observations about models, SDN controller as feedback node

6 SDN controller

The SDN controller is at the heart of the architecture. It is the intelligent node that controls resources to deliver services. Its core function is the real-time multi-dimensional convergence of a changing resource environment and a changing service demand environment toward an optimum, where the optimization criteria may also change in time.

A given resource is controlled by only one SDN controller, this being an important resource grouping criterion (note). An SDN controller exercises management-control over some set of resource groups, conventionally portrayed in a resource or data plane to the south of the controller, and accessed through a data-controller plane interface (D-CPI). The aggregate of resource groups under a given SDN controller is referred to as an SDN domain; the term may also include the SDN controller itself.

Note – It is recognized that some aspects of a resource may be controlled by other means. Coordination of these activities is necessary, but the details are beyond the scope of the architecture. An example would be an operations support system (OSS) locking a resource via the SDN controller before performing a software upgrade.

An SDN controller offers services to its clients by exposing a virtual resource group to each client. The interface to clients is conventionally drawn to the north of the controller, and called an applications-control plane interface (A-CPI). The interface is also often called a northbound interface (NBI), sometimes an intent interface.

An SDN controller may act as a service-consuming application by consuming resources and invoking services offered by a subordinate or peer SDN controller. An SDN controller may have peer interfaces to non-SDNs, and may invoke services from non-SDNs or offer services to them.

An SDN controller may be implemented via centralized or distributed computing technology, with or without redundancy. These are important considerations, but are implementation details of the SDN controller as a functional block, and lie beyond the scope of the SDN architecture.

Note – Architectures are recursive. It would be perfectly legitimate to develop an implementation architecture for SDN controllers, subject to alignment with the functional view of this top-level SDN architecture.

Further reading: control as feedback, business boundaries, virtualization, orchestration, client environment, non-CPI interfaces, complexity, reliability and availability, interfaces

6.1 SDN controller as feedback node

As an architectural functional block, an SDN controller may appear in any number of environments. It is fair to ask what makes it an SDN controller, what aspect, should it be absent, would preclude a candidate body of software from being an SDN controller.

The core function of an SDN controller is to continually adapt the state of its resources to serve its clients according to an optimization policy. The concept of state is interpreted broadly. The very existence of a particular resource is itself a state, as are the values of all of its attributes and its relationship to other resources, both within and beyond the controller's domain. Figure 5 illustrates this, the SDN controller as the active node in a feedback loop.

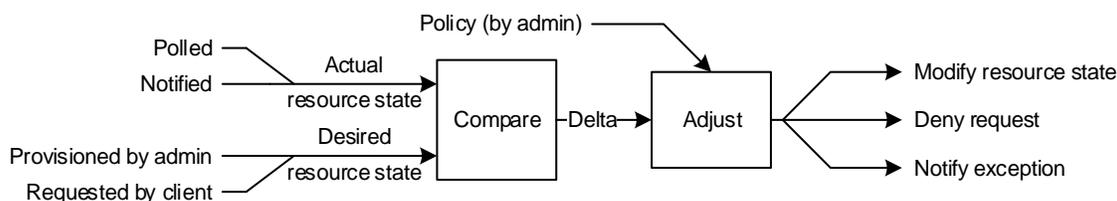


Figure 5 – Control as feedback

Control is the process of establishing and maintaining a desired state. Open-loop control is of little value in the SDN context; feedback is the essence of control.

Figure 5 illustrates that the administrator of the resources configures some desired state, which may change from time to time. The various clients invoke service requests, which translate into desired state. These also change from time to time, both as clients come and go, and also as each client modifies its service demands. The actual state of the resources includes measures of load, changes due to failure and repair, and administrative actions. Actual state may be polled on an ongoing basis or updated asynchronously by notification. The controller evaluates the difference between desired and actual state in light of the optimization policy (note), which may also change from time to time, and attempts to modify the state of the resources accordingly. The optimization policy may include the ability to create (or request the creation of) new resources, scale or migrate existing resources, etc. State convergence may include negotiation with neighboring domains.

Note – It is expected that the optimization policy would include compliance with all of the client SLAs. Nevertheless, there could be provider policies that govern how the effects of network overload or failure are distributed among clients of various types.

If desired and actual states cannot be reconciled within the bounds of the policy, the SDN controller issues an exception, either immediately by rejecting a client request or as a notification during ongoing operation.

In addition to its core feedback control function, an SDN controller may invoke arbitrary supplementary functions, support arbitrary collections of additional features, any number of interfaces and protocols, any number of applications of arbitrary type and complexity, resources of any and all categories, etc., according to its particular deployed purpose.

Further reading: orchestration, resource groups, complexity, delegation, notifications

6.2 Orchestration

In the sense of feedback control, orchestration is the defining characteristic of an SDN controller. It is the selection of resources to satisfy service demands in an optimal way, where the available resources, the service demands and the optimization criteria are all subject to change.

When the optimization function indicates that a better optimum can be achieved, the orchestration function adjusts the state of the resources under its control to move toward that optimum.

Note – Policy may or may not prevent in-service reassignment of resources.

Depending on its environment and the optimization criteria, the orchestration function may be complex or relatively simple. Candidate orchestration algorithms may be evaluated on their ability to deal with complexity, as well as on their real-time responsiveness to change.

Orchestration includes at least the following functions:

- Validating service requests from clients against client-specific policy, denying requests that fall outside policy.
- Configuring resources and subordinate services to satisfy client service requests according to the provider's controller-wide policy and the specific policy and SLA associated with the client.
 - This includes configuring data plane entities to enforce policy themselves, for example by policing ingress traffic rates and priorities, by filtering ingress traffic address fields.
 - This includes creating and configuring data plane mechanisms for dynamic sharing such as prioritized weighted fair queueing engines.
 - This includes the configuration of encapsulation, address translation, or other means to ensure mutual isolation of client traffic.
 - If permitted by client and provider policy, this includes requesting the instantiation, scaling, migration or deletion of resources, as may be appropriate.

- This includes configuring and operating network support functions such as protection switch engines, CFM, STP, performance monitoring, if specified by provider policy or specified or implied by the client SLA.
- These operations may need to be recursively invoked from one SDN controller downward to another before they reach actual traffic-processing engines.
- Coordinating service requests and service changes with neighbor domains, be they SDN domains or otherwise, and whether the other domains are peers or hierarchically subordinate.
- Publishing notifications of interest to particular client subscribers, in the terms of the pertinent client context, and notifications of global interest to global subscribers (the administrator of the SDN controller).

Virtual resources in the CC are dedicated to the given client, and exist because the client needs to express its service demands in terms of some kind of entities, usually at least its points of presence on the provider's network. The satisfaction of a client service request often requires the orchestration function to select additional resources from the underlying pool available to the server, i.e., resources that are not dedicated to the given client, and invoke services against those resources. The selection of such shared resources is affected by policy and resource traffic load or congestion.

Because virtual resources, service requests, and notifications exist in specific client contexts, it is necessary that the orchestration function collaborate intimately with virtualization.

Further reading: Control as feedback, complexity, virtualization, client context, policy, notifications

6.3 Virtualization

As stated, orchestration is the core active component of an SDN controller. Virtualization is the process that populates and maintains the resource and naming environments that link orchestration with clients. Orchestration and virtualization interoperate inextricably; for clarity, they are described as separate functions, without meaning to imply anything about separability of implementations.

Recall that virtualization is the abstraction of resources allocated to a particular client, application, or service. Because virtual resources are conceptually contained in client contexts, a client context is a precondition for a virtualization.

The initial state of an SDN controller is a default CC for the administrator. This default CC conceptually contains all of the underlying resources. Each non-default client context receives a custom virtualization, in the form of allocations from the underlying resource pool.

Not all of a client's virtual resources are necessarily visible to the client. Of those resources that are exposed to the client, not all attributes are necessarily visible, and not all exposed administrator-writable attributes are necessarily writable by the client. Namespace translation is needed only for resources that are visible to the client.

A virtualization may be created and populated in any way suitable for the circumstances.

Implicit creation and population is exemplified by a residential subscriber who supplies a street address to a portal as part of signing up for service, from which a fixed access point of presence can be derived as the necessary initial virtual resource. Additional client resources may also be populated from equipment and inventory records, such as DSL modem or PON ONU, and a physical or virtual residential gateway (RG). The services delivered to such a client are otherwise free to use whatever resources may be selected by the orchestration function.

Another method is dynamic assignment. In satisfying client service requests, the orchestration function may subdivide, combine or otherwise abstract underlying resources as necessary. The resulting resource would be recorded in the appropriate client and service context, though it would not be visible to the client.

A third method is explicit negotiation between customer and provider, for example to agree on the detail of a number of UNIs of a corporate customer, or to specify an entire topology of a sub-network that may be leased by the provider to a re-seller. Explicitly pre-negotiated resources permit the client to exercise greater detail in its service requests, while correspondingly reducing the number of choices available to the orchestrator in satisfying service requests.

Changes to resource environment, business agreements, or other factors, necessitate the ability to incrementally update virtualizations in place. Updates must be coordinated with the orchestration function.

Further reading: client context, roles, virtualization definition discussion, orchestration

6.4 Resource sharing

The principle of centralized control implies that no given resource is controlled by more than one external entity. If several clients contend for the resource, they must contend as clients of an SDN controller, which is responsible to arbitrate amongst their claims. There can be no multiple-writer problem in an architecturally compliant SDN.

Resources may be allocated to a given client or service in their entirety, in which case there is no possibility of contention. This allocation may be static during CC creation, or by the orchestration function according to a schedule or on demand. Examples of such wholly allocated resources – at some level of virtualization – include physical points of presence (UNI, NNI data plane ports), dedicated wavelengths, dedicated spectrum, dedicated time slots.

Resources may also be shared dynamically, i.e., packet by packet. This requires that the infrastructure to be able to support at least some aspects of prioritized weighted fair queueing, that the SDN controller (potentially in a hierarchy) recognize attributes such as priority, committed, and extended information rate in service requests, and that it combine such requests in accordance with the provider's policy on oversubscription.

Arbitrary dynamic sharing on all network resources is complicated by the need to identify specific underlying resources in a mix of arbitrary client topologies and at least a partial mesh in which arbitrary flows may be multiplexed on arbitrary links, changing the mix at every node. This complexity encourages the allocation of large flows to fixed capacity paths through the core network, with the core network engineered to achieve statistically satisfactory performance for smaller flows.

Another resource that can be shared is address space. Depending on business agreement, a client is often free to use all of the (e.g., IPv4, C-VID) address space that exists at its hand-off network layer; the client relies on the server to isolate traffic. This may be done by physical separation, by encapsulation, or by network address translation (NAT). Alternatively, addresses or address blocks may be administratively allocated (and represented in policy form) when the client context is created, or may be requested as needed during operation.

Further reading: orchestration, virtualization, complexity

6.5 Delegation

In the simple SDN model, a resource is a passive entity that exposes actions and attributes for query and manipulation by an SDN controller or application. Nothing requires resources to be passive, however, and there are strong reasons for certain resources or resource groups to have active functionality.

The boundary between active and passive resources is not necessarily strict. For example, otherwise passive resources may be instrumented to collect performance statistics and either upload them autonomously or generate threshold crossing notifications.

Accordingly, the architecture provides for active resources. Active resources must be visible to and controllable by the SDN controller or application, to a degree that satisfies the engineering design or service requirements of the particular situation. The algorithmic intelligence of an active resource may be built in or may be downloaded and installed on demand, either by the SDN controller itself or by out of band means that are not explicit in the architecture.

Active resources may publish notifications, to which the SDN controller or application can subscribe. Active resources may engage in peer interactions, for example those of STP or CFM. If an active resource qualifies as a control function, in the sense of driving a feedback loop, the deployment engineer is responsible to ensure that all feedback loops, including that of the SDN controller itself, are configured such that they do not destabilize one another.

Further reading: control as feedback, orchestration, resource groups

6.6 Client context

A client context (CC) models everything that needs to exist in an SDN controller (note) to support a given client. To a given SDN controller, a client exists only during the lifetime of an associated CC. If the relationship with the client terminates, all trace of the client may be removed by deleting its CC and performing suitable resource reclamation.

Note – By the principles of recursion and abstraction, SDN access to a hardware network element still requires a virtualization of its resources and a client context.

A client business relationship and its services are required to persist even in the absence of a management-control association (session) between client and server. When the client re-establishes a management-control association, its policy, resource, and service views must be restored. This includes, for example, notification subscriptions and client-specific profiles. Logs may be needed, in part to permit clients to view events of interest that occurred during disassoci-

ation. Much of the CC must therefore be persistent. As the repository of client-sensitive information, backup and restoration of a CC must be secured from unauthorized access.

A session is an association between a specific instance of a client organization and an SDN controller. It may be modeled in terms of user login. A session normally begins with an exchange of identity and security credentials, followed by agreement on an initial state, much of which may be restored from prior sessions. A session continues with the exchange of information. Each information exchange can be attributed to that session, for example in an audit log. A session may continue indefinitely, or end with an explicit logout, a failure, or a timeout. In many cases, the client will automatically restore a lost session, and it may be expected that the SDN controller also be able to restore a lost session.

A business agreement may allow for any number of user logins – sessions – from one or several client platforms. In a minimal client context that supports only one login, session identification may be implicit.

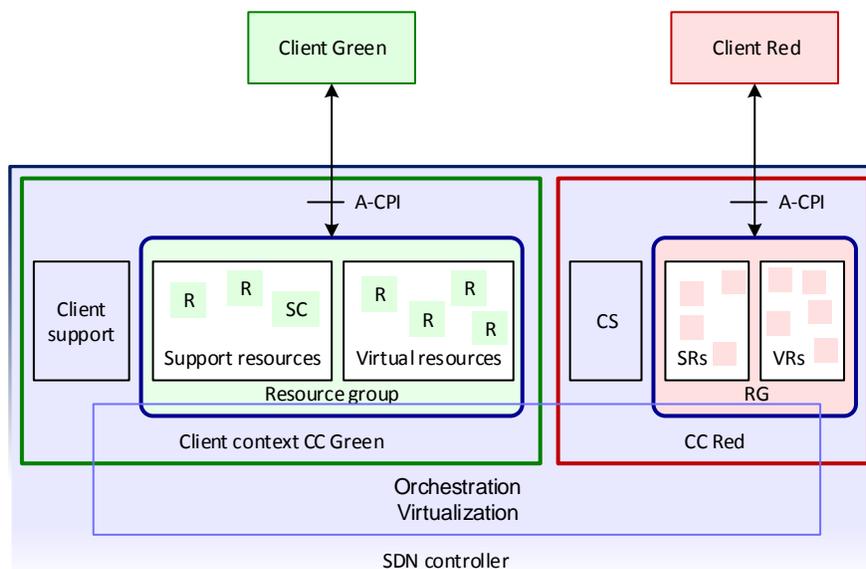


Figure 6 – Client context

Figure 6 illustrates that, logically, a client has visibility and control of the resources in its resource group. Semantically, the A-CPI between client and controller is at the boundary of the RG. The CC also conceptually contains support information and functionality that are needed to support the client but that are not exposed to the client.

The functional contents of a CC are:

- One resource group. The RG defines the protocol interface (A-CPI) exposed to the client. Two sub-groups are identified.
 - Virtual resources represent infrastructure resources that are created from the SDN controller's underlying resources through the process of virtualization, and that are exposed to the client by way of a mapping function. The most pertinent example is the client's point(s) of presence.

- Support resources, which represent functions hosted in the SDN controller itself. Their purpose is to enable or facilitate interaction with the client. Examples of support resources include client login profiles, notification subscriptions, logs. As suggested by the resource called `SC`, current service contexts are also retained here, including the client's SLA expectation.
- Client support, CS. This function supports the client in as many ways as may be necessary. It may include both code, ephemeral data, and persistent data. While the client's RG contains resources that are at least visible to the client, though not necessarily writeable, the CS block may contain additional resources that are not exposed to the client at all, or private views of resources dedicated to the client. Content of the CS includes information to map names and actions between client and server, policies on what the client is allowed to see and do (client view definition), how client traffic is to be isolated (e.g., NAT, encapsulation, reserved or on-demand addresses or ranges), and provider logs or usage measurement for billing or other purposes. Part of the policy function deals with security, representing everything necessary to allow the client to connect to the SDN controller, for example certificates, keys, management-control channel encryption policy, client logins, policies on single or multiple login.

As will be apparent from the overlapping descriptions above, the blocks in figure 6 are not intended to be sharply bounded discrete components with defined interfaces, merely to clarify the functions of the CC. Additional functions may be included, subject to the definition that the CC represents all things and only things that support a given client.

When an SDN controller is created, a default CC is required, with unrestricted privileges and views. As well as configuring the SDN controller itself, and importing underlying resources, this permits the SDN administrator to create and populate CCs for additional clients.

Further reading: multiple logins, policy, resource groups, virtualization, realization, server context, service context, persistence

6.7 Service context

NB – Not to be confused with server context

Services usually involve data plane information exchanges with the client domain (figure 1). These data plane exchanges typically persist during intervals when no active management-control session exists between the SDN controller and the client. The existence and characteristics of a service are represented by a persistent service context object in the client context of its owner.

A service context conceptually contains at least all of the attributes of a service as requested by the client, and server-specific information necessary to map service attributes into the realization of the service. Notifications, logs, and profiles may or may not be associated with a service. There is no implication that a given service be owned by any particular user login, but also no prohibition.

A service context is populated by the server, driven by events that may include any or all of:

- Client context creation by server administrator, to provide prenegotiated or default service
- Service creation operation by client, specifying service type and attributes
- Service update operation by client, modifying existing service
- Changes in network state or policy on the part of the server that alter the static attributes of orchestration and virtualization necessary to realize the service.

When a service context is deleted, all resources dedicated to that service are released and no further trace of the service remains in the server's environment.

The identifier of the service context allows for unambiguous exchange of service-related control, query and notification information between client and server. In the general case, a client can create, read, update and delete (CRUD operation) services at will. If policy restricts a client to a single service, the service context may be created by the server administrator, and not subject to create-delete operations by the client.

At any given time, zero or more services may exist for any given client.

Further reading: client context, multiple logins, server context

6.8 Server context

NB – Not to be confused with service context

Recall that a client context contains everything needed by an SDN controller to support an application. By symmetry, a server context (SC) contains everything necessary for an application to support the client end of that association. The existence of a properly matched SC is necessary and sufficient for a client application to communicate with an RG. As with the CC, an administrator creates SCs as needed.

As with the CC, an SC is at least partly persistent. Figure 7 illustrates an application client that is an SDN controller, with server contexts that relate it to underlying services and resources, two of which are within the same Blue trust domain as the controller itself.

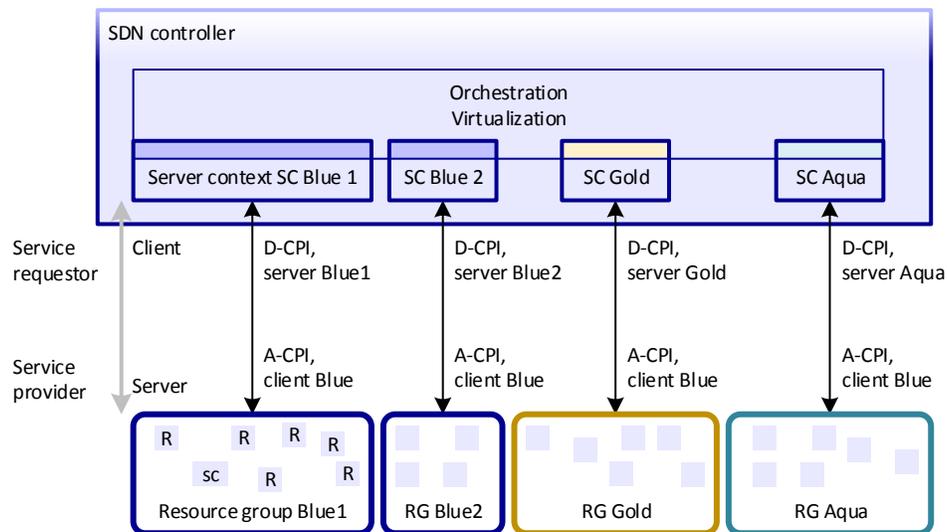


Figure 7 – Server context

An application's SC must contain information for each RG to which it may connect. This includes login and security information that will be acceptable to the server's RG environment, and whatever stack parameters are needed to achieve semantic compatibility. If multiple logins from the same client SC are contemplated (e.g., parallel OpenFlow, OF-config, SNMP sessions), this information would be accordingly multiplied. Local logs may be required in the SC.

Customized software may be included, for example to adapt the application's internal APIs to the protocol and information model of the subordinate RG. While the A-CPI references on figure 7 imply that the underlying RGs are themselves aligned with the SDN architecture, this is not necessarily the case. They may well be non-SDN entities of any arbitrary nature, as long as the application's SC is capable of importing their capabilities.

Note – OpenDaylight plug-ins may exemplify such customized software.

Dynamic information retained in the SC would include off-normal state information, such as alarms or notifications of reduced capacity.

The underlying RG is mirrored into the client-side RDB, specifically in the SC.

The SC may also contain information such as logs. Performance monitoring uploads may be retained here, at least temporarily until deleted or offloaded onto analytics systems. Other functions are other storage are not precluded.

It is for further study whether a client application requires policy enforcement to protect itself from the server.

Further reading: service context, client context, resource database

7 Applications

The classical SDN model describes an application plane, populated by instances of applications. An application is an entity that requests a service of some kind from an SDN controller. The idea

is recursive; an application can itself be an SDN controller. Sometimes applications are described in terms of end users, but from the perspective of a server SDN controller, there is no way to know. An SDN controller offers services; an application-plane entity on the other side of its A-CPI invokes the services.

Recognizing that an application may be an SDN controller, it may be useful to list some characteristics to clarify the terminology.

- From the viewpoint of an SDN controller, anything that invokes services from one of its A-CPIs is an application.
- If an entity orchestrates more than one resource group across its D-CPIs, it is an SDN controller.
- As the intelligent node in a feedback loop, an SDN controller is likely to engage in fine-grained state monitoring and to have a wide repertoire of nuanced responses to deviations from the desired state. A service-consuming application is more likely to express only its purpose (intent) against a single server, leaving the satisfaction of that purpose to underlying intelligence. Feedback to a non-SDN controller application is more likely to be coarse-grained: accept-deny or alarm on failure, with little expectation that the application can remedy the situation. A non-SDN controller application is more likely to have no ongoing management-control sessions and minimal or no notifications subscriptions.
- If an entity exposes its functionality to human users, it is probably an application. Because humans always communicate with software through mediation devices, the depth and sophistication of the mediation affects the interpretation of this criterion.

Further reading: control as feedback, orchestration, principles of SDN

8 Putting it together: the integrated architecture

Figure 8 reiterates the earlier figure 2. It shows the controller as the central element in an SDN, and with its conceptual content of contexts and the orchestration-virtualization function. In this illustration, SDN controller Blue offers services to clients Green, Red and itself, Blue. CC boundaries and resource groups are colored to indicate dedication to clients, while RG border colors identify the underlying resource owner and the accompanying isolation and policy enforcement barrier. An administrator in the Blue organization has unrestricted view and privilege over the SDN controller itself, along with all client and server contexts.

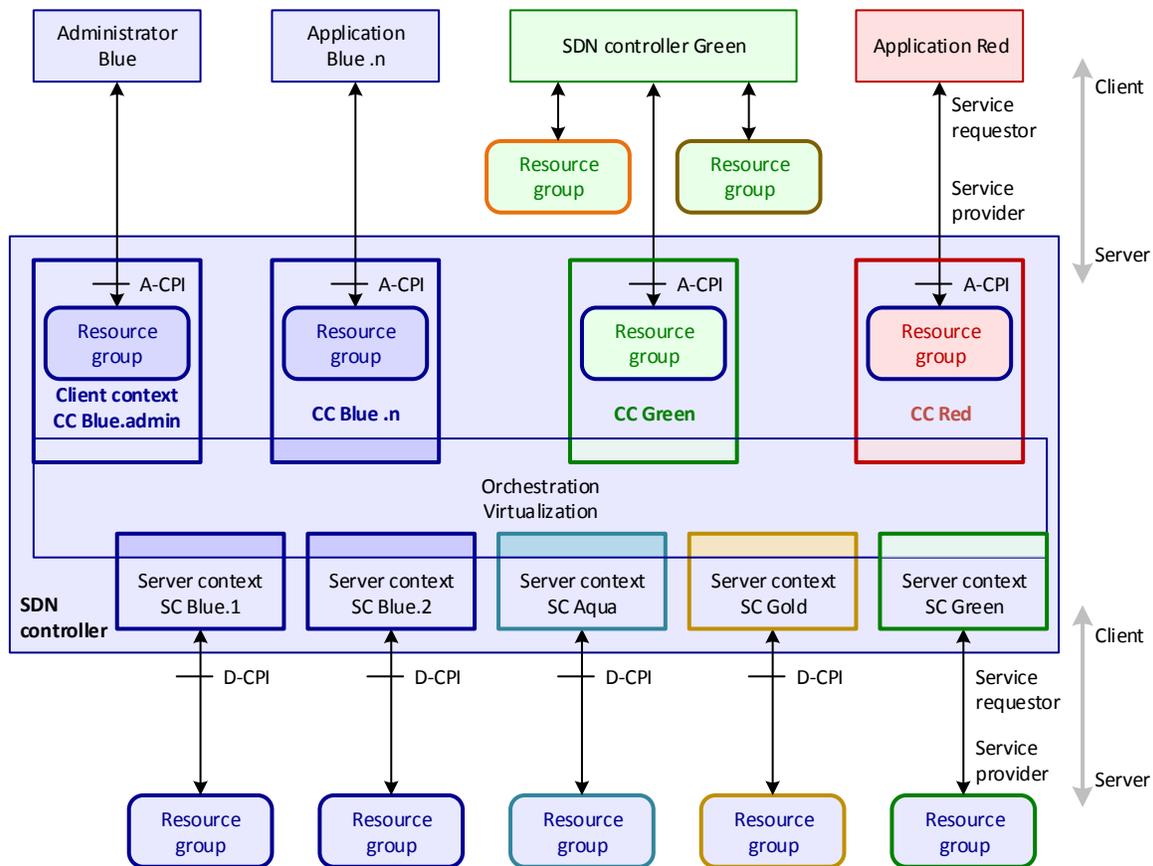


Figure 8 – SDN controller illustrating contexts

The SDN controller consumes underlying resources offered by Gold and Aqua, as well as two resource groups owned by itself, Blue. The Green client application is shown as an SDN controller in its own right, orchestrating services and resources across several domains, one of which is Blue. At the D-CPI, the Blue SDN controller is shown in the same way, and in fact as a service requestor client to Green. Reciprocal requestor-provider roles allow Blue and Green to operate as peers in providing services that may have been requested by some client above either domain.

The resource group in the administrator client context includes the default unrestricted policy, and the information necessary to manage (note) the SDN controller and the various client and server contexts. In addition, it contains support objects (resources) for use by the administrator login, for example logs, subscriptions, and profiles.

Note – Because the idea of a CC or an SC is that they contain all of the information related to a client or server, the administrator’s RG is modeled as containing only the information needed to manage such CCs and SCs, but not their contents. Whether this distinction actually matters is an open question.

8.1 Interfaces

8.1.1 Controller plane interfaces

The data-controller plane interface (D-CPI) is the interface between an SDN controller and the resources under its direct control. The application-controller plane interface (A-CPI, also called NBI) is the interface between application and SDN controller, across which an application invokes services from the SDN controller.

By recursion, A-CPI and D-CPI are instances of a common interface, as viewed from the server and client sides, respectively. The semantic content that may be conveyed across the interface is unbounded, though the architecture models the entities as client and server. A client should be able to query or discover the underlying environment and subscribe to notifications from a server. A client commands; a server responds.

From semantics down through the various levels of the stack, any given association must support at least a subset of some common information view and should deal with misalignment gracefully. Exact or partial compatibility across an interface may be known a priori, discovered or negotiated when the client-server association is established, custom-configured, or any combination of those.

The idea of a priori knowledge deserves explanation. In real-world deployments, functionality, information model compatibility and stack compatibility will be known and specified from the initial engineering of the network through vendor and product selection, through product qualification, and into coordinated deployment, along with integration into related systems including business processes such as product catalog development, advertising and billing. Random ad-hoc associations of applications, controllers and network resources are of no concern.

8.1.2 Additional interfaces

An SDN controller may be implemented as a collection of distributed components, the interfaces between which are internal, and out of scope. This clause describes external interfaces to the controller that are not explicitly shown in the architecture. These fall into three broad categories, which overlap widely.

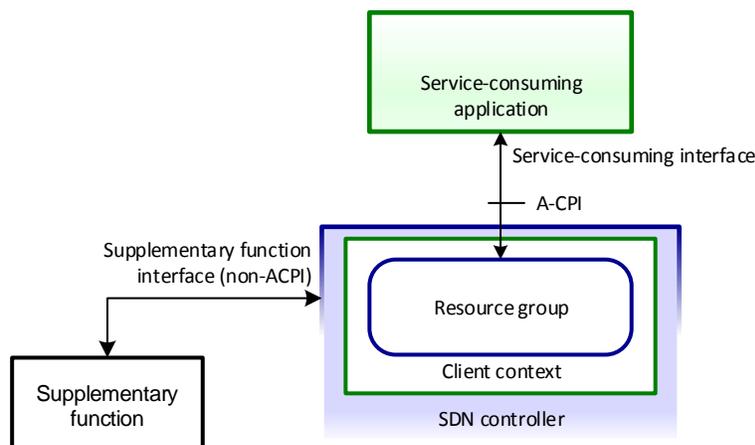


Figure 9 – Supplementary function example

Supplementary functions

An SDN controller may use other interfaces and other functions, and in particular may invoke supplementary functions as exemplified in figure 9. If a supplementary function modifies the state of network resources or services, the modified state must be flagged to the SDN controller that invoked the function, for example via notification or function return parameters.

Supplementary functions may also be internal to the SDN controller. An SDN controller may perform a given function by executing its own code, including library code, or by invoking supplementary functions that reside externally. For example, scope and clean partitioning suggest that path computation will almost certainly be external to an SDN controller, as will services such as AAA, DHCP, DNS.

The architecture makes no distinction between internal and external functions; it assumes that the SDN controller can achieve the necessary outcome.

An SDN controller is architecturally allowed to instantiate, destroy and scale resources, some types of which may be subsumed under the concepts and tools of the ETSI NFV discipline. Because of the current prominence of NFV, it is worth emphasizing a strong preference that the outcome be an invocation by the SDN controller of NFV-related tools, either from a library or across a protocol interface, but not by way of reinvention.

Note – If an NFV-related entity invokes SDN functions, it should align with the usual SDN server-client model at an A-CPI.

Directories and databases

An SDN controller requires considerable information for its normal operation, which is architecturally modelled as being available in a conceptual resource data base. It may also be desirable that some of this information reside elsewhere, accessible to the SDN controller when needed.

Such information might include widely shared subsets of a common information base, such as inventory, common service profiles, client contexts. The SDN controller may update such common information, and may require transactional interfaces to them to assure consistency. The SDN controller may also need to subscribe to change notifications from such external information.

Information to support mappings between client and server views of resources and services may reside externally to the SDN controller. Reasons for external hosting may include security aspects of read-write access shared between the server (provider) administration and various clients.

Backup and restoration datasets are additional examples of pertinent external information stores, as are logs and performance monitoring archives, and software management repositories.

Other FCAPS functions

An SDN controller is responsible for establishing and maintaining services for its clients. A number of additional functions fall in the FCAPS definition (fault, configuration, ac-

counting, performance, security), some of which may be partly or entirely outside the scope of a given SDN controller.

A prime example is the exception thrown by an SDN controller when it is unable to converge desired state to actual state. This may trigger diagnostic, fault isolation, and repair activity that is not part of the SDN controller's normal functions. Similarly, functions such as equipment installation may be beyond the normal scope of an SDN controller's duties.

The reason for a possible non-CPI interface to the SDN controller is that such activities typically involve the SDN controller, at least to the extent of overriding its normal operation in relation to certain resources, and often as proxy for external activities such as testing or diagnosis.

Note – In some such circumstances, an ordinary administrator access may suffice.

Restoration from backup and software upgrade are related functions that may need to be performed on SDN controllers that are at least partially operational during these processes.

Further reading: Persistence, initialization, Resource data base

8.2 Notifications

Notifications are an important link in the feedback mechanism that is the core of an SDN controller. All SDN entities are expected to follow a publish-subscribe notifications model. Subscriptions and profiles are associated with particular client logins and are persistent.

A client context includes a definition of the notifications available to that client, and is thus part of the virtualization of underlying resources configured by the server administrator.

The notifications available to any given client must be discoverable by that client, and notifications management capabilities must be supported. These include subscription management and, if alarms are subscribed, severity assignment profiles, active alarms lists, and alarm reporting control (ARC).

Notifications from an underlying layer may be of interest to clients, but must be virtualized into the client context. This may require more than a namespace change; for example, an attribute value change notification from the underlying resources may become an object creation notification to a client. Notifications delivered to a client should also be available to the administrator or a log for troubleshooting, correlated with the triggering notification from the infrastructure.

Particularly because a client need not maintain a continuing login session with the SDN controller, client-visible notification logs may also be required.

The administrator role can control notification subscriptions on behalf of itself or any of its clients. This is recursive to a client that plays the administrator role in its own CC.

ITU-T M.3702 [10] is the recommended baseline for notifications management. Further detail of the SDN notification framework appears in a Notifications Framework document, in progress.

Further reading: control as feedback, client context, multiple logins

8.3 Peer controllers

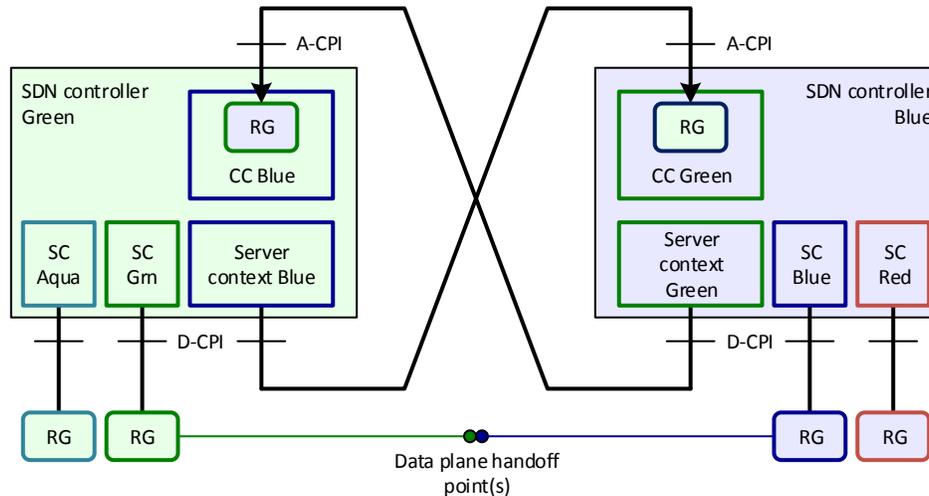


Figure 10 – Peers as symmetric requestors and providers

Figure 10 illustrates two SDN controllers in a peer relationship, in which either may act as client to invoke services from the other as server. Each requires a client context and a server context for the other. SLA, enforced policy and information hiding are applicable between peers just as between hierarchical provider and customer. The CC and SC for a peer may be separate, as shown, but they could combine certain aspects such as association credentials and security policy.

9 Specific perspectives on the architecture

9.1 Security

Data plane connections face the same security concerns in SDN and in current networks. Traffic in and out of the network must be traceable to a known client, policed for SLA compliance as well as against DOS, spoofing, or other attacks, and isolated from other traffic that may use the same address space.

Many other security concerns are also the same as in current networks, for example the procurement and deployment of executable software, the security of the management communications network, the integrity of directories and network services such as DNS, and the need to safeguard against insider errors and attacks.

What is different in SDN is the widespread exposure of management-control interfaces to third parties, and the wide range of services that may be offered by way of these interfaces. With more parties involved across more interfaces that offer greater functionality, the difficulty of assuring proper behavior across each interface increases manifold, whether misbehavior is due to error or malice. Interfaces traceable to a common information model are an important way to ensure API integrity from the beginning.

The SDN architecture specifies that a management-control association between a client application and an SDN controller be properly authenticated and secured. Not least because a client ap-

plication may not be fully trustworthy, the architecture provides for limited information exposure by way of the virtual resources offered to the client, and policy enforcement on all of the client's management-control operations. The combination of access security, tailored virtual resources and policy enforcement is called a client context.

Flaws in the client context may result from software or configuration error, from insider tampering, or from network environment shifts that expose new possibilities beyond those covered by policy rules. Probing and auditing of client contexts may be valuable in discovering such flaws, with audit logs available for post facto analysis.

It may also be permitted that a client download executable code onto an SDN controller. Even if such code is certified and secured, it behooves the SDN controller to bound its execution environment within at least the same client context, information hiding, and policy enforcement constraints.

Possibly as a new, or exaggerated, feature of SDN, backups from SDN controllers and applications may contain commercially sensitive information, for example about customers and partners, and must be protected from leakage, not only during storage but also upon re-installation or subsequent analysis. Historically, much of this information would have been localized in semi-isolated BSS/OSS environments.

Further reading: persistence

9.2 Reliability, availability

In the data plane, reliability and availability considerations are largely unchanged when management-control is logically separated from traffic processing infrastructure. Protection switching and load sharing concepts remain applicable, and the ability of an SDN controller to delegate functions into the infrastructure helps to avoid unacceptably slow responses to events.

Although SDN does not mandate a move of existing hardware functions into software, such a move is facilitated by SDN, and in particular NFV, and is under way today. From the reliability and availability viewpoint, a network implemented in software, as distinguished from a network controlled by software, requires re-examination.

- Software data plane functions may be more finely decomposed and therefore require more network assets, and more dynamic network assets, to chain them together, than was the case before.
- The dynamics of failure discovery may differ. For example, loss of signal on a physical link can be signaled immediately to concerned parties, while loss of heartbeat may be recognized only after several heartbeat intervals elapse. This will affect the agility of response, the proper kinds of response, and the commitments that can be made to customers.
- Virtual network functions reside on general purpose substrates. New instances may be readily created as an alternative to repair of a failed instance, or hot spares may be available for immediate use, subject to correspondingly agile updates to connectivity.
- Because failures may lose state, restoring state from checkpoints or continuously synchronizing spare instances will be important.

As to control, the possible separation of an SDN controller from traffic processing resources implies a reduction in availability. This may be addressed in several ways.

- Traffic processing engines should be designed to continue their function unchanged in the absence of an active controller connection (note). This will affect partitioning and delegation of functionality from the SDN controller into data plane elements.

Note – This is another way of pointing out that services normally continue to operate during the absence of a management-control session.

- SDN controllers may be implemented in load-sharing or redundant configurations, with near real-time synchronization. Subordinate systems will need to support multiple controller sessions or fast establishment of new controller sessions.

The notion of recursion implies that a service request may need to propagate through a number of hierarchical or neighbor SDN controllers before it is satisfied. If a controller fails or is otherwise unable to perform its expected functions, it will be necessary to unwind intermediate resource commitments, either for a new attempt or for a failure indication to the originating client. It may be helpful to re-use concepts such as resource reservation in the forward direction and commitment in the reverse, or acknowledgement, direction.

Further reading: persistence

9.3 Identifiers

The core information model [3] includes material on identity, names, labels, and addresses. This architecture document refers to them generically as identifiers. Identifiers are essential to determine what is under consideration and where it is to be found, both of these within some space that may be implicit.

In the SDN architecture, identifiers are significant because clients view resources and services in their own terms, which must be mapped into the server's identifier space, as well as possibly translated into quite different resources.

It is not specified how identifiers in either space are initially established.

- They may be administratively configured and known by manual provisioning, discovery, directory query or other means. Recognition of a wireless client by SIM card identity is one example; another is translation from a customer's service address to a provider's equipment or cross-connect panel appearance.
- They may be proposed by one side and accepted by the other during negotiation.
- Other mechanisms are not precluded.

When a client's virtual resource is derived as a one to one view of some underlying resource, a simple identifier mapping suffices. A virtual resource may also be based on a complex view of several underlying resources. As an example, a single subnetwork identifier exposed to a client may map onto a set of underlying resources that represents arbitrary functions, technologies, administrations, geographic locations, etc. Creating and traversing such mappings are major func-

tions of virtualization and orchestration. Because the orchestration function generally has a choice of underlying resources, the mapping may change over time.

9.4 Realization considerations

The administrator role creates and modifies client and server contexts and the SDN controller's optimization criteria. Any or all of these could be done through offline tools, with the result installed in the SDN controller as a subsequent step.

An SDN controller requires a significant amount of persistent information, particularly in client and server contexts. Backup and restoration may be appropriate for such information. It will be important to understand the gold reference for each fragment of persistent information and to ensure timely consistency in mirrors, copies or derivations. These issues exist in current networks; whether there are new aspects in an SDN is a matter for further study.

Given suitable definition of which instances of client application are permitted to connect with which instances of SDN controller, at least part of the CC may be amenable to common external storage, accessible through a directory service, such that it can be downloaded to any of several possible SDN controllers that allow association with the given client. As an example of this flexibility, if a roaming user were to connect to a foreign SDN, a partial CC for that user could be dynamically downloaded and installed on the SDN controller from the user's home location, whereupon the SDN controller could dynamically instantiate services as specified by the included service contexts.

An SDN controller may be implemented in a distributed computing form, with redundancy or functional separation or both. Maintaining synchronized state during operation and during exception recovery is important.

As a new network entity of arbitrary and varying complexity, it may be expected that an SDN controller will very often be implemented as a virtual network function (VNF). If the SDN controller has responsibility within the same NFV domain, it will be important to plan a controllability tree that avoids the SDN controller unintentionally disrupting its own operation.

Further reading: persistence

9.5 Initialization

The preparation of network resources for service requires operations such as code installation and configuration of initial parameters. Ultimately, these may come down to craft terminals and faceplate connectors, factory-installed software that is at least capable of downloading and installing operational software, and the like. Some equipment may be able to configure itself and appear autonomously to its neighbors as a new resource without manual intervention. Certain levels of test and equipment management may also be required during initial turn-up. These functions are not new to SDN. It is for further study whether the SDN architecture has anything special to say about them.

The SDN architecture presupposes an initial SDN controller environment that includes working software with an administrator client context. The administrator client context permits the creation of additional client and server contexts. As the administrator creates server contexts and the controller contacts the server, the resources exposed by that server are made available in the

RDB, from which they may be statically allocated by the administrator to particular CCs or consumed dynamically by the orchestration function.

Note – Architecturally, transparent pass-through of underlying resources is thought of as a degenerate case of virtualization. Just as with similar situations in information hiding, strong security, and policy enforcement, transparent virtualization may explain why many existing approaches to intra-administration SDN do not recognize the need for these functions.

Given the necessary physical or virtual raw material, an SDN controller may instantiate or cause the instantiation of, new resources from scratch, e.g., VNFs on VMs or on bare metal compute servers.

Further reading: Persistence, resource data base

9.6 Complexity

Orchestration, the core function of an SDN controller, is a multi-dimensional real-time optimization problem over a complex and potentially large space. Further, it must simultaneously work across the underlying resources and all of their various virtualizations. The complexity of the controller is a legitimate concern. Several approaches are available; all of them are likely to be useful in various combinations and according to circumstances.

Continuing improvements in computing technology assist in the ability of an SDN controller to handle larger workloads, especially when the controller is implemented as a VNF that can be distributed, scaled, or migrated.

Subdividing the problem space is a common scaling technique, whether by defining a greater number of smaller SDN domains or by allocating separate functional areas to separate SDN controller components. Existing BSS/OSS/NMS functionality may be re-used en bloc or incrementally.

An important opportunity for separation of concerns is the ETSI NFV focus on life cycle maintenance of certain resources, and the SDN focus on services and the use of resources for service delivery. To illustrate the point, a soft switch may be a virtual network function (VNF), created by NFV entities and made available as a network node to an SDN controller. The SDN controller would then dynamically configure layering and forwarding of particular flows or flow bundles through the switch.

The optimization problem may be simplified if the SDN controller's feedback criteria define a range of results that is considered to be good enough, rather than seeking to converge on a precise optimum. It may both simplify the problem and improve customer quality of experience if resources are not reallocated while they are in use.

Although its functionality is reduced, an SDN controller is simplified if it deals only with the allocation of pre-existing resources, and does not consider the possibility of scaling resources or creating new resources. In such a case, some other entity (e.g., related to ETSI NFV) might have responsibility for dynamic resource inventory, updates to which could be made known to the SDN controller.

The architecture assumes that some part of the underlying resources can be efficiently virtualized into disjoint subsets that satisfy the needs of various clients, including QoS and availability commitments, while isolating each from the others, and instrumenting the lot for monitoring, fault management, billing, and network engineering. The virtualization is assumed to remain optimal, or near optimal, in the face of customer churn, network build-out, and other such factors.

Some aspects of this problem are comparatively easy to address, for example isolating customers in a common address space, or the policing and prioritizing of traffic near the point of network ingress. Other aspects may be harder to automate.

Whether intentional or otherwise, over-engineering reduces the complexity of resource management. This is particularly true in the core, where elephant services (few, large, long-lived) are worth managing explicitly, while mouse services are best treated statistically. The cost of under-utilized resources is visible, but when compared to the comparatively hidden cost of complexity, especially the development and maintenance of algorithms, a certain amount of over-engineering may well be cost-effective.

As suggested, the greatest cost may be that of human effort in developing and – not to be underestimated – validating resource allocation and optimization algorithms. Machine learning or artificial intelligence may play an important role in this regard.

9.7 Persistence

Information has a useful scope and a useful lifetime. Information also exhibits primacy characteristics, in the sense that some information is primary and other information is derived.

Note – Global business agreements exemplify primary information; the interpretation of a global business agreement into SLA or policy for a specific SDN controller or service exemplifies derived information. An accepted service request is primary, its implementation is derived. Except for planned resources that do not yet exist, it is widely accepted that the network itself is the primary source of resource state, from which additional information may be derived.

Policy for redundant and persistent information storage should trade off primacy and reconciliation against the cost of derivation. Factors to be considered include security, synchronization, storage, bandwidth, code, practices development, training, and administration. The choice of hosting site for redundant and persistent information must also consider the likelihood and consequences of various exception cases. For a given SDN entity, several policies and hosting sites may be appropriate for different classes of information.

It is expected that much of the content of client and server contexts and the provider's optimization policy will justify remote persistent storage and possibly local redundancy to reduce the effect of SDN controller failure. When several business entities are involved in a relationship, each bears primary responsibility for its own information storage, but nothing precludes one player from offering synchronization, backup and restoration services to other players.

Further reading: security, realization considerations, resource database, client context, server context, service context, notifications, reliability and availability, initialization

9.8 Migration and coexistence

When new resources are developed and installed, for example virtual network functions from concepts developed by ETSI NFV ISG, it is natural to deploy them in an SDN domain and use them in alignment with this SDN architecture. However, the broad scope of SDN and the abstract nature of its architecture are also intended to facilitate its incremental adoption and coexistence with current networks.

An SDN controller may be instantiated above a non-SDN domain, abstracting the views presented by non-SDN network elements, element management systems (EMSs), or network management systems (NMSs) into a common information model that can be virtualized and used to deliver services to clients.

A subnetwork, which may be as small as a single network element, may be an SDN island domain, in which the SDN controller adapts its underlying resources to control and management by non-SDN entities, be they peer or superordinate entities. An SDN controller may act as a peer in multi-domain information exchanges such as via BGP.

An SDN controller may have responsibility for certain aspects of resource control, but not others. In the near term, much existing FCAPS and OSS/BSS functionality will remain with current OSS/BSS/NMS entities. When distinct entities control different aspects of the same resources, it will be important to ensure that they do not disrupt each other.

In short, SDN principles can be applied incrementally to peer, subordinate, and superordinate entities, as opportunity and business justification arises, and need not be viewed as all or nothing, or as a strong push into asset replacement.

9.9 Relationship of SDN and NFV

The ETSI NFV ISG focuses primarily on the creation of virtual network functions (VNFs). SDN can help organize VNFs into network services (NS). VNFs and NSs then become resources for use in constructing services for clients, which is the focus of SDN. To a significant extent, these disciplines thus complement each other. The SDN architecture fully recognizes the value in the work done in ETSI NFV ISG, and encourages its use.

ONF TR-518 [2] is a more detailed comparison of SDN and NFV perspectives. ETSI NFV ISG has also published a less abstract view of the topic [7].

It may be expected that all OSS/BSS/NMS/EMS and SDN controllers will be or become VNFs, along with utility services such as DNS, DHCP, etc. These are already software entities, and may already be able to run on x86 cores, so the migration should be straightforward. But much or all of the “data plane” exposed by these entities is network management-control traffic, so the levels of abstraction will need to be carefully considered.

10 Appendices

Having intentionally minimized the size of the main document text, it is recognized that a number of topics justify wider discussion. Some of them may be developed as stand-alone Architec-

ture Notes documents over the course of time. These appendices contain material of sufficient interest to warrant inclusion in the main document.

10.1 Discussion of definitions

A definition should be concise, precise and pedantic. However, informative material is often helpful in interpreting the definition, drawing out its implications, and relating it to other concepts. Clause 4 contains the definitions used in this document; this appendix contains their discussions.

Ed note: definition text is hyperlinked from clause 4. Edit it there, not here!

10.1.1 Abstraction

Definition: The representation of an entity or group of entities according to some criterion, while ignoring aspects that do not match the criterion.

Discussion: Even when the underlying object is completely physical, it is always viewed through a filter of some sort. The filter necessarily loses some aspects of the original and thus necessarily provides an abstract view. Some filters are natural, for example human visual or tactile perceptions, but those of interest to SDN are software representations chosen to expose aspects of interest. Any number of filters (abstractions) can exist for a given underlying object.

It is sometimes considered that, at the bottom of a recursive hierarchy, the lowest level SDN controller controls a bare metal switch, for example through OpenFlow. It must be understood that, even in this case, the hypothetical bare metal switch necessarily exposes an abstraction, namely an OpenFlow logical switch. This observation changes only in detail if it is asserted that the bare metal switch is directly visible to OF-config, rather than to OpenFlow, or if OpenFlow is used to program chips rather than network elements.

A similar case can be made about any other management-control interface, and about even further descent into the internal structure of managed entities.

The inclusion of groups of entities in the definition is recognition that entities may be merged or combined in an abstraction.

10.1.2 Client

Definition: An entity that receives services from a server.

Discussion: See clause 10.1.10 Server.

10.1.3 Client context

Definition: The conceptual component of a server that represents all information about a given client and is responsible for participation in active server-client management-control operations.

Discussion: Refer to clause 6.1 SDN controller as feedback node and multiple logins

10.1.4 Domain

Definition: A grouping of entities according to some criterion.

Discussion: Common domains of interest to SDN include

- Administrative domains, the complete set of resources owned by a given business entity.
- Geographic domains, partitioned within an administration for scaling, or across administrations for other reasons, for example the North American [telephony] Numbering Plan.
- SDN domain, the set of resources directly controlled by a given SDN controller.
- Technology or deployment domains, for example transport, access, wireless, cloud. NFV may be a separate domain. Non-SDN domains would be distinguished from SDN domains.

10.1.5 Management-control continuum (MCC)

Definition: The principle that the functions of management and of control are largely, if not entirely, the same.

Discussion: Ultimately, the existence and behavior of networks can be traced to decisions and actions by humans. Activities performed by humans are often thought of as management. The further away from human involvement, the more likely an activity is thought of as control, perhaps exemplified at the extreme by automatic gain control built into a hardware feedback loop.

In the continuum, a great many activities may be performed at a greater or lesser remove from specific human involvement. The MCC concept observes that, to the managed or controlled target entities, it makes no difference why it is being asked to perform some action, or by whom. A practical consequence is that management-control functional components may be re-used in products or deployments that are thought of as managers, as controllers, or anything in between.

In SDN, both terms may be used where their functions align with common usage, but *control* is the preferred default term, as it better matches the idea of feedback in real time.

Note – It may be observed that the resource life cycle maintenance perspective of the ETSI NFV ISG causes it to use the term *management* by default.

10.1.6 Orchestration

Definition: The ongoing selection and use of resources by a server to satisfy client demands according to optimization criteria.

Discussion: Refer to clause 6.2 Orchestration

10.1.7 Policy

Definition: A rule that guides and constrains subsequent actions of, and interactions between, other parties.

Discussion: The term *policy* is widely used but rarely defined. In some contexts, any rule is a policy, including for example, a match-action entry in a switch. In this architecture, a policy is understood to focus on aspects of management (human) purpose, and/or directing the response to exceptions or violations.

A policy is the result of some thought, not just flipping a coin (note). So an administration anticipates conditions that might arise, decides what to do under the foreseen circumstances, then installs a policy on some engine that is capable of observing the circumstance and taking the specified action. Every time the circumstance arises, the action is the same.

Note – A policy could of course obtain a specified amount of randomness from a specified source, when required by a specified event, and use that randomness in a specified way to determine an action. The policy is nevertheless a fixed rule, and the action to be taken is the same each time.

As a counterexample, if some circumstance arose that had not been previously considered, it might be necessary to make an instant decision, but the decision criteria could hardly be called policy-based. (Of course, if an ad hoc decision proved to be optimal, it could later be formulated as a policy.)

The word *subsequent* captures the ideas of forethought and persistence.

As to the word *other*, an administrator role may install a policy that pertains to itself in a subsequent client or user role, but it is not meaningful to pre-constrain actions in the same role.

The idea of interacting parties implies correctly that policy is enforced at interfaces.

In the SDN architecture, an administrator installs a policy whose enforcement protects the server resources from possible abuse by the client, either intentionally or otherwise. The policy also captures the server's commitment to the client, for example by way of SLAs. Finally, the policy could define details about how the server satisfies client demands. An example might be a policy to encapsulate traffic, physically isolate it, or use some particular subset of a common address space.

10.1.8 Recursion

Definition: The repeated application of a process in which the input to each iteration except the first is derived from the output of the previous iteration.

Discussion: In a computational context, recursive iterations are usually executed by the same process, and the result delivered by the recursive process is realized when recursion ceases. In SDN, however, the usual case is that the output of an iteration is passed as input to a different software entity, usually on a different platform. It is therefore the usual case that each entity incrementally realizes some part of the ultimate added value as a side effect of its iterative step. The entities involved in recursion are often SDN controllers, but especially at the beginning or end of a recursive sequence, they may be other entity types.

In mathematical recursion, the input and output of an iteration are of the same data type, for example floating-point reals. However, especially in hierarchical SDN recursion, successive iterations on successive platforms may involve very different protocol stacks, data models, services, and actions. This is also the case in layered networks: each layer is recursively a client of the layer below it and a server to the layer above it. This is recursive, but the details of inter-layer interfaces may differ substantially from one another.

The principle of recursion requires that every interface and every iteration be the same, but the principle must be understood at a high level of abstraction, not necessarily as a statement about details.

No entity has visibility beyond its immediate neighbors, and therefore has no way to know whether it is part of a deep recursion or a shallow recursion. Knowledge of recursion beyond immediate neighbors is possible only from a perspective of omniscience.

Repeated local iteration is available to SDN entities as needed, but has no special or unique significance. SDN identifies two kinds of recursion that are significant.

Hierarchical recursion is a pattern in which higher-level SDN controllers orchestrate a broad scope of resources and services across one or more lower-level SDN controllers with narrower scopes and less abstract resources. In the inverse direction, resources exposed at lower levels of abstraction are recursively partitioned and recombined into increasingly abstract resources and services at higher levels.

Neighbor recursion is a pattern in which SDN controllers peer to deliver services across domains. All participants would be expected to expose comparable levels of abstraction and service, and any SDN controller could act as either client or server to its neighbors, ad hoc. Services invoked by neighbor recursion are more likely to fit a call model – service invocation, service usage, service release – rather than persisting indefinitely.

10.1.9 Resource

Definition: Anything that can be used to deliver a service.

Discussion: A resource is modelled as an instance of a managed object class. Resources may be subdivided and combined with other resources in any way supported by the information model. When underlying resources are virtualized by an SDN controller, the result is a new resource that may be offered to a client. When exposed resources are manipulated by a client over an A-CPI, possibly with additional resources as marshalled by the server, the result is a service.

Resources are intentionally defined very broadly, in keeping with the broad scope of SDN. The following examples illustrate the range of possible resources:

- A dark fiber
- An equipment plug-in
- A termination point
- A virtual machine
- A virtual network function
- A switch
- A firewall
- A subnetwork
- A catalog of service offerings

10.1.10 Server

Definition: An entity that provides services to a client.

Discussion: The SDN architecture is extensively based on client-server concepts. A server, which may be an SDN controller, exposes resources to zero or more clients (q.v.), which may themselves be SDN controllers or service-consuming applications. The clients invoke various operations on these resources to realize services.

Especially when client and server are in separate administrative domains, the server may be called a provider, and the client may be called a user, customer, or consumer (note). A client may also be called an application.

Note – The industry sometimes uses the term *tenant* in this context. However, the counter-party to a tenant is a landlord, which is usually inappropriate for an SDN context, so the term is deprecated in this architecture.

10.1.11 Server context

Definition: The conceptual component of a client that represents all information about a given server and is responsible for participation in active server-client management-control operations.

Discussion: Refer to clause 6.8 Server context.

10.1.12 Service

Definition: The delivery of value for some time interval by a provider to a consumer.

Discussion: It is understood that a relationship usually involves a mutual exchange of value, for example, but not necessarily, financial compensation. As used in SDN, a service is usually closer to the network resources than the counter-flow of compensation. Temporally, a service invocation is generally the acceptance of an offer, with subsequent compensation.

A given set of entities may offer and invoke services from each other ad hoc. There need not be a fixed server-client relationship.

A service continues to exist during intervals when there is no association between client and server management-control entities. Services such as VoIP may be quiescent during such intervals; other services such as transport may include continuing data-plane activity. In the case of roaming and mobile services, the physical point of attachment may vary over time, even across administrative domains.

In many cases, the actual delivery and use of the service happens across a data plane interface (see figure 1). In such cases, the management-control interface to the SDN controller (A-CPI) is used to invoke, control, modify, and terminate the service, but not to receive the value delivered by the service.

10.1.13 Service context

Definition: The conceptual component of a client context that represents all information about a given service.

Discussion: Refer to clause 6.7 Service context.

10.1.14 Virtualization

Definition: An abstraction of underlying resources on a server, whose criterion is allocation of those resources to a particular client, application, or service.

Discussion: Refer also to clause 6.3 Virtualization.

It must be understood that virtual has a number of meanings in the industry, sometimes implying nothing more than a software implementation on general-purpose computing hardware. This approximates its meaning in ETSI NFV [2], [6].

The SDN meaning is independent of underlying implementation. Further, an SDN virtualization may be several steps of abstraction away from any implementation. These steps of abstraction may subdivide and combine underlying implementations in arbitrarily complex ways. As an example, a VPN may abstract “real” layered networks from a number of administrative domains, with some topology visible and tunnels to conceal further underlying connectivity, and with reserved or on-demand capacity on its various links.

10.2 Observations about models

A model is an abstraction of some underlying reality. The same underlying reality will support any number of models from different perspectives. Different models will be useful in emphasizing different aspects of the underlying reality.

Criteria for the value of a model include

- No model should contradict itself or the laws of physics.
- No two models of the same underlying reality should contradict each other.
- Occam’s razor: a model should be no more complex than necessary.
- Nor should a model be less complex than necessary.
- A model should imply no conclusions that are not true.
- A model should lead to productive insights.

This document includes a number of mutually consistent perspectives and explains their relationship. The objective is to offer integrated insights that lead to better solutions in all of the multiple dimensions of concern to SDN. To assist in integrating views, many clauses end with cross-references for further reading.

Although a model may be described in pictorial form, no picture tells a complete story. A picture must be accompanied by explanatory text that relates graphic entities to the model and to the underlying reality. As with models themselves, pictures may be deficient in any number of ways, not least clutter and ambiguity.

10.3 Multiple logins

The client context CC of architecture issue 1.1 is a deep dive into the agent of architecture 1.0. Among other things, it recognizes that the client may not have a continuous management-control

session with the controller, but that the client environment and the services must persist nevertheless.

Having recognized that there can be zero active sessions in a given client context at any given time, it is natural to enquire whether there can be more than one, either sequentially or simultaneously. The simultaneous case is more complex for various reasons, and is the basis for these notes.

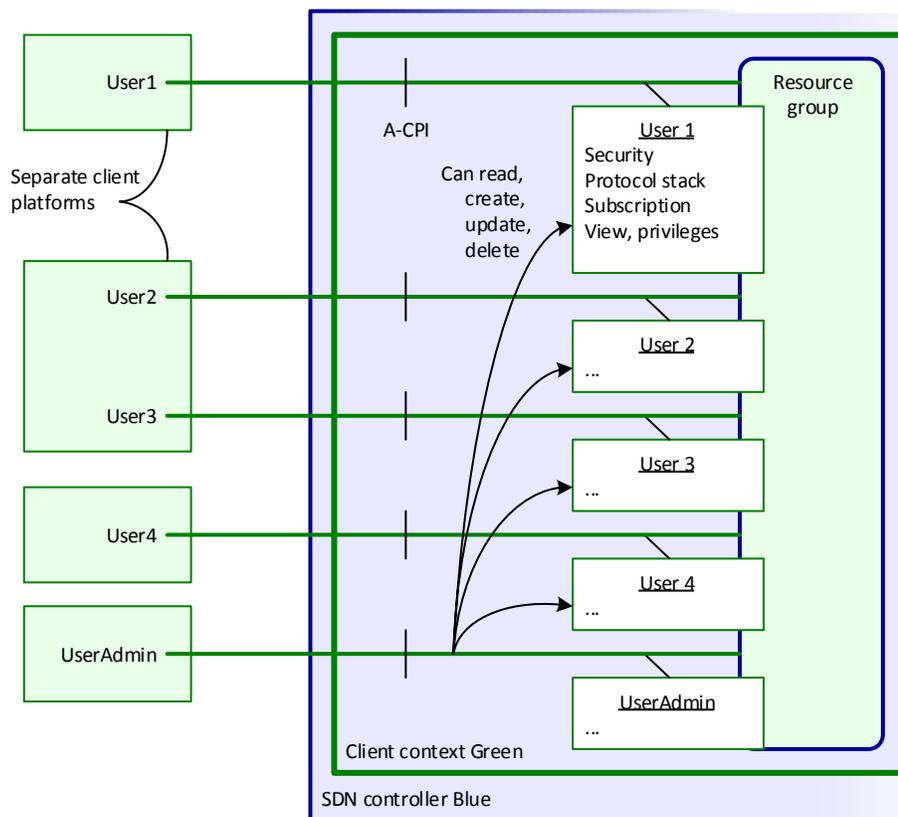


Figure 11 – Multiple users of a client context

Figure 11 serves as the basis for discussion. As to notation: Blue is a provider, the owner of the SDN controller; Green is a client. The Green bound around the client context suggests that everything within is dedicated to Green, as distinct from other clients or general purpose use. Users are hosted by green boxes, and are semantically connected to the resource group in the Green CC. All interactions between Green users and the Blue SDN controller occur within the context of some overarching business agreement, the concrete manifestation of which is the blue bound around the RG, a policy enforcement barrier. A Green user has no visibility or control outside the RG.

White boxes with thin green boundaries indicate user-specific information and code, some of which may be internal to the RG. These may be designated client sub-contexts.

It would be possible to support only a single client sub-context, allowing multiple sessions. In such a case, every user of the single login would have the same credentials, privileges, and profiles, including for example notifications subscription. If Green includes human users, however (or proxies thereof), each user may be expected to have different interests and different responsi-

bilities, and therefore to require different views, different privileges, and different profiles. That is, beyond the global policy applicable to all of Green's operations, each user login is expected to be further constrained by an independent policy.

Figure 11 therefore shows a user administrator login (which would be the default login created by Blue at the time of CC creation). The user administrator has unlimited visibility and privilege (unrestricted sub-context policy), but only within the Green RG. Among its privileges is the ability to create and manage client sub-contexts for other Green users.

Code in the Green CC is responsible to orchestrate and virtualize Green resources and services, and in particular to resolve contention issues within Green's resource and service environment. At this point, it should be apparent that Green's environment is an SDN controller itself (figure 12).

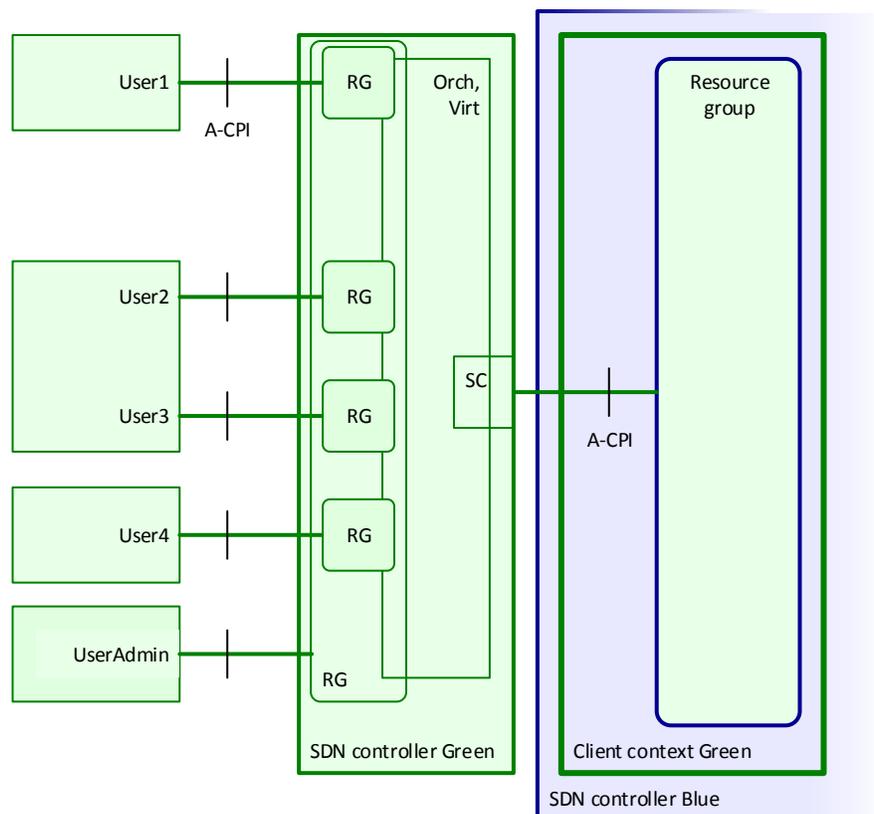


Figure 12 – Separated SDN controller Green

The separate Green SDN controller is now in a position to orchestrate resources across more than one SDN or legacy resource groups (see Green in figure 2). Further, the separate Green SDN controller can contain arbitrary code and use arbitrary external databases and supplementary functions, completely independent of Blue business or security concerns.

This arrangement may offer considerable benefits. Examples include:

1. Enterprise client Green.com, with custom-tailored policies about which of its employees are allowed to see and do what, under the control of a Green administrator. These policies are of no concern to Blue in any event, and Green could well regard them as business

confidential. Specific subdivision of responsibilities might be to a group of service administrators with minimal privilege, and to a separate group of troubleshooters with full visibility of the Green context.

2. Interpreting an intent interface. To say, “Bob is (or is not) allowed to connect to the Internet,” is to raise questions that cannot possibly be answered by Blue. The Green SDN controller may look at directories, policies, or anything else of interest to interpret the intent into terms known to Blue. The Green-Green A-CPI can be whatever is appropriate to the custom service; the Green-Blue A-CPI would be a standardized view of an information model.
3. The Green SDN controller could be a web portal, fielding requests from various users for various Green services, validating and billing the users, and passing the requests into Blue for implementation. In this context, Green could be the external services division of network operator Blue.
4. As a variation on the previous point, Green CCs could accept association requests from anonymous clients. Two options could be offered.

Log in as an existing client. The Green SDN controller might need to refer to external directories or data bases to authenticate the user and discover enough information to redirect the user session to the proper CC on (in this case) the Blue SDN controller.

Negotiate service and business offerings with a potential new customer, for example by exposing a catalog for user browsing and ordering. If an agreement was reached, a CC would need to be created on the spot on (in this case) the Blue SDN controller, and the client session redirected into that CC.

5. The possibility of parallel protocols, for example OpenFlow-switch, OF-config and SNMP logins, having appropriately different views over the same underlying resource base.
6. The possibility of an open wireless interface, available for anyone to connect. Upon connection, an open service could be established with a default SLA and policy for each device. For premium service, it would also be possible to identify the device or user, for example via SIM card or web portal login, thereupon download and install a customized SLA and policy for that device, and collect usage and billing information.
7. Customer self-install of residential access service, using a similar model, either modifying an existing service (e.g., when moving house) or newly subscribing. In this case, subscriber SLA and policy might need to be customized on the spot.

Figure 12 is an omniscient view that shows how the SDN architecture may be extended, but Blue neither knows nor cares whether the Green application is structured as an SDN controller. This illustrates both the broad scope of SDN and also its ability to seamlessly interwork with non-SDN entities.

As to the SDN architecture, this analysis suggests that a single instance of A-CPI per CC suffices, and that multiple logins need not be supported. That said, it might yet be worthwhile for Blue to offer an embedded controller-lite feature that permitted, for example, one writer and multiple readers with individual notification subscriptions. Another possibly worthwhile option might be

for Blue to offer parallel logins for OpenFlow-switch, OF-config, OVSDB, and/or SNMP, with appropriately predefined profiles for each.

10.4 Evolution from issue 1 to issue 1.1

As described in clause 1, SDN architecture issue 1.1 clarifies and extends issue 1 [1]. Experience with issue 1 identified the need for clarification in a number of areas, including some of the terminology. A number of concepts evolved, and several new concepts were brought into the forefront of issue 1.1. The differences in terminology, presentation, and emphasis are sufficient to warrant an explicit comparison. Figure 13 is quoted from figure 3.3 in SDN architecture issue 1.

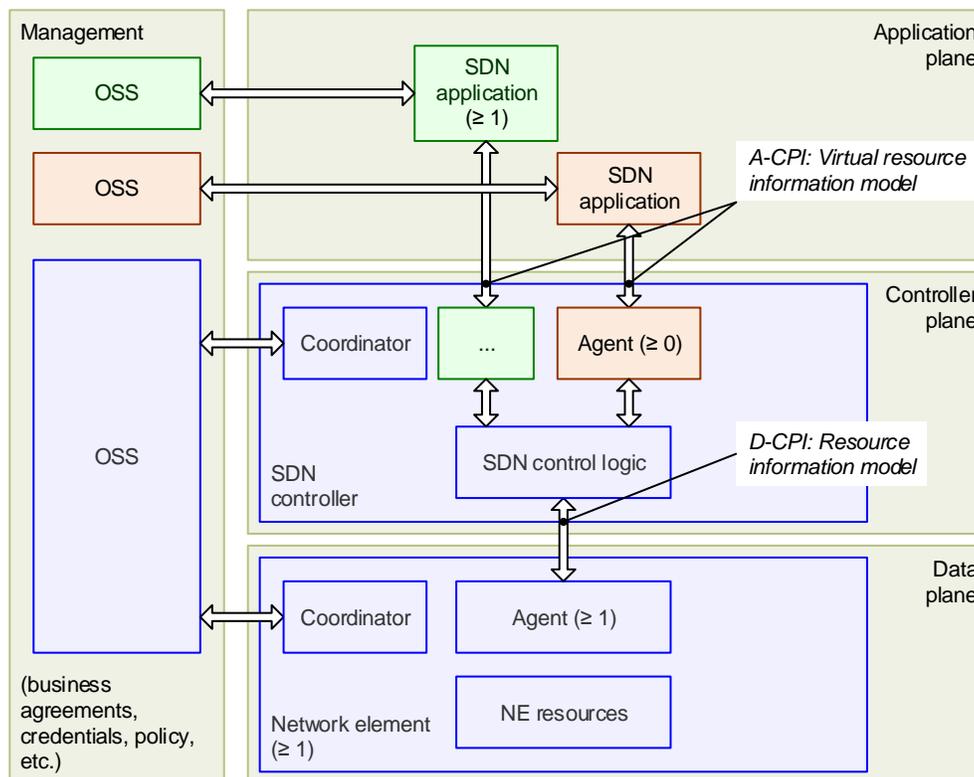


Figure 13 – Issue 1 architecture

Issue 1 shows a management block with a distinct interface and termination in the managed entities. Issue 1.1 incorporates the concept of management-control continuum (MCC), in which an administrator has greater scope and privilege than non-administrator roles, but is not fundamentally different. Both issues expect that the administrator would exist within the same business or trust domain as the entities being managed-controlled.

The data plane in issue 1 is described as a set of (virtual) network elements. The term *network element* has implications within the industry that may be misleading in some cases. In issue 1.1, the issue 1 network element is subsumed as one class of the *resource group* concept.

While issue 1 refers to resources in general, nothing more than a reference to (virtual) network elements is said about their organization. As a matter of clarification, resource grouping has been added to issue 1.1.

The agent in issue 1 is conceived as a window through which a remote SDN controller or application could manage its underlying resources. In issue 1.1, the term has changed to client context, and a great deal of its functionality has been explicated.

Suggested by arguments of symmetry, issue 1.1 introduces a server context concept, which is responsible for interaction with an underlying server instance.

Issue 1 identifies virtualization, orchestration and real-time responsiveness as attributes of an SDN controller, but the controller's role as the intelligence in a feedback loop is described as an option. Issue 1.1 recognizes that feedback intelligence is the essential characteristic of anything that purports to be an SDN controller.

Although it is not apparent from figure 13, issue 1 takes an omniscient view of recursion. Issue 1.1 emphasizes the fact that individual entities see only their neighbors and have no visibility of possible further recursion. While the omniscient perspective is valid, the view from a single given entity is easier to relate to current open-source or vendor-specific product development.

Issue 1.1 introduces the idea that a client may have more than one service active simultaneously (hence the service context concept) and that a client may not have continuous management-control access to the server during the lifetime of its services. Further, a client business may require server logins for a number of its employees. Sessions and persistence are needed for these and other reasons.

Issue 1 describes a resource-based SDN architecture. Work on intent interfaces has been synthesized into issue 1.1 as a complementary service-based view of the architecture. In the general case, the client and server are jointly responsible for developing mappings that translate between the client's frame of reference and the server's frame. The mapping allows for separation between standard information models and arbitrary client- or application-specific models.

Issue 1 refers to resources generically; issue 1.1 explicitly includes the contributions of the ETSI NFV ISG in offering virtual resources to SDN environments.

Further reading: management-control continuum, resource groups, client context, server context, control as feedback, orchestration, service context, multiple logins, persistence, service and resource oriented models

11 Back matter

11.1 Acronyms

AAA	Authentication, authorization, accounting	CS	Client support
A-CPI	Applications-controller plane interface	C-VID	Customer VLAN identifier
API	Applications programming interface	D-CPI	Data-controller plane interface
BSS	Business support system	DHCP	Dynamic host configuration protocol
CC	Client context	DNS	Domain name system
CFM	Connectivity fault management	DOS	Denial of service
CPI	Controller plane interface	EMS	Element management system

ETSI NFV ISG	European Telecommunications Standards Institute – Network functions virtualization – Industry specification group	ONU	Optical network unit
FCAPS	Fault, configuration, accounting, performance, security	OSS	Operations support system
IP	Internet protocol	PON	Passive optical network
LLDP	Link layer discovery protocol	RDB	Resource data base
MAC	Medium access control	RG	Resource group, Residential gateway
MCC	Management-control continuum	RP	Resource provider
MIMO	Multiple input, multiple output	RU	Resource user
NAT	Network address translation	SC	Server context, sometimes Service context
NBI	Northbound interface (A-CPI)	SDN	Software defined networks
NFV	Network functions virtualization	SDO	Standards development organization
NMS	Network management system	SLA	Service level agreement
NNI	Network-network interface	SNMP	Simple network management protocol
NS	Network service	STP	Spanning tree protocol
OF	OpenFlow	UNI	User-network interface
ONF	Open Networking Foundation	VLAN	Virtual local access network
		VM	Virtual machine
		VNF	Virtual network function
		WDM	Wavelength division multiplexing

11.2 References

- [1] ONF TR-502, SDN architecture, 2014, https://www.opennetworking.org/images/stories/downloads/sdn-resources/technical-reports/TR_SDN_ARCH_1.0_06062014.pdf
- [2] ONF TR-518, Relationship of SDN and NFV, 2015, https://www.opennetworking.org/images/stories/downloads/sdn-resources/technical-reports/onf2015.310_Architectural_comparison.08-2.pdf
- [3] ONF TR-512, Core information model (CoreModel), 2015 – Ed: under revision. Will the revised version still be TR-512?
- [4] ETSI NFV, Network Functions Virtualisation, An introduction, benefits, enablers, challenges and call for action, 2012, http://portal.etsi.org/NFV/NFV_White_Paper.pdf
- [5] ETSI NFV GS NFV-INF 001 V1.1.1, Network Functions Virtualisation (NFV); Infrastructure Overview, 2015
- [6] ETSI NFV GS NFV 004 V1.1.1, Virtualisation Requirements, 2013
- [7] ETSI NFV EVE005 – supply proper ref when published
- [8] IEEE Std 802.1AB-2009, IEEE Standard for Local and Metropolitan Area Networks – Station and Media Access Control Connectivity Discovery, [LLDP](#)
- [9] IETF [RFC 3444](#), On the difference between information models and data models, 2003
- [10] ITU-T Recommendation M.3702, Notification management – Protocol neutral requirement and analysis, 2010

11.3 Contributors

Malcolm Betts, ZTE

Chen Qiaogang, ZTE

Luis Miguel Contreras Murillo, Telefonica
Nigel Davis, Ciena
Paul Doolan, Coriant
Dave Hood, Ericsson
Chris Janz, Ciena
Lothar Reith, DT

Sibylle Schaller, NEC Labs
Fabian Schneider, NEC Labs
Stephen Shew, Ciena
Eve Varma, Alcatel-Lucent
Maarten Vissers, Huawei