802.1CF text review

Date: 2015-07-13

Authors:

Name	Affiliation	Phone	Email
Max Riegel	Nokia Networks	+49 173 293 8240	maximilian.riegel@nokia.com

Notice:
This document does not represent the agreed view of the OmniRAN TG It represents only the views of the participants listed in the 'Authors:' field above. It is offered as a basis for discussion. It is not binding on the contributor, who reserve the right to add, amend or withdraw material contained herein.

Abstract

This document contains a compilation of text of the P802.1CF specification as generated by assembling in FrameMaker contributions on Network Reference Model, Network Discovery and Selection, and SDN Abstraction. The document is aimed for editorial review and consolidation of the presentation of the technical content.

2. Normative references

The following referenced documents are indispensable for the application of this document (i.e., they must be understood and used, so each referenced document is cited in text and its relationship to this document is explained). For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments or corrigenda) applies.

IEEE Std 802.1D™, IEEE Standard for Local and metropolitan area networks—Media Access Control (MAC) Bridges.

IEEE Std 802.1Q™, IEEE Standard for Local and metropolitan area networks—Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks.

3. Definitions

For the purposes of this document, the following terms and definitions apply. *The IEEE Standards Dictionary Online* should be consulted for terms not defined in this clause.[1]

To do: Extract nondictionary terms and definitions from document text

4. Acronyms and abbreviations

AN	Access Network
ANC	Access Network Control
ANI	Access Network Identifier
AR	Access Router
ARC	Access Router Control
ARI	Access Router Interface
BH	Backhaul

[attention here]

CIS	Coordination and Information Service
CIS	Coordination and Information System
CN	Core Network
CNS	Core Network Service
CNSI	Core Network Service Identifier
EUI48	48-bit Extended Unique Identifier
NA	Node of Attachment (e.g., AP)
NAI	Network Access Identifier
NRM	Network Reference Model
SS	Subscription Service
SSI	Subscription Service Identifier
TE	Terminal
TEC	Terminal Control
TEI	Terminal Interface

[1] *The IEEE Standards Dictionary Online* subscriptions are available at http://www.ieee.org/portal/innovate/products/standards/standards_dictionary.html.

5. Conformance

5.1 Requirements terminology

5.2 Conformant components and equipment

5.3 Protocol Implementation Conformance Statements (PICS

6. Network Reference Model

6.1 Basic architectural concepts and terms

NOTE— This section is essentially adopted from IEEE 802.1AC Chapter 7 with some figures added from IEEE 802 for illustration.

The architectural concepts used in this and other IEEE 802.1 standards are based on the layered protocol model introduced by the OSI Reference Model (ISO/IEC 7498-1) and used in the MAC Service Definition (IEEE Std 802.1AC), in IEEE Std 802, in other IEEE 802 standards, and (with varying degrees of fidelity) in networking in general. IEEE 802.1 standards in particular have developed terms and distinctions useful in describing the MAC Service and its support by protocol entities within the MAC Sublayer

6.1.1 Protocol entities, peers, layers, services, and clients

The fundamental notion of the model is that each protocol entity within a system exists or is instantiated at one of a number of strictly ordered layers, and communicates with peer entities (operating the same or an interoperable protocol within the same layer) in other systems by using the service provided by interoperable protocol entities within the layer immediately below, and thus provides service to protocol entities in the layer above. The implied repetitive stacking of protocol entities is bounded at the highest level by an application supported by peer systems, and essentially unbounded at the lowest level. In descriptions of the model, the relative layer positions of protocol entities and services is conventionally referred to by *N*, designating a numeric level. The *N* service is provided by an *N* entity that uses the (*N* 1) service provided by the (*N* – 1) entity, while the *N* service user is an (*N* + 1) entity.

Figure 1 illustrates these concepts with reference to the layered protocol model and service access points of IEEE 802 end stations.

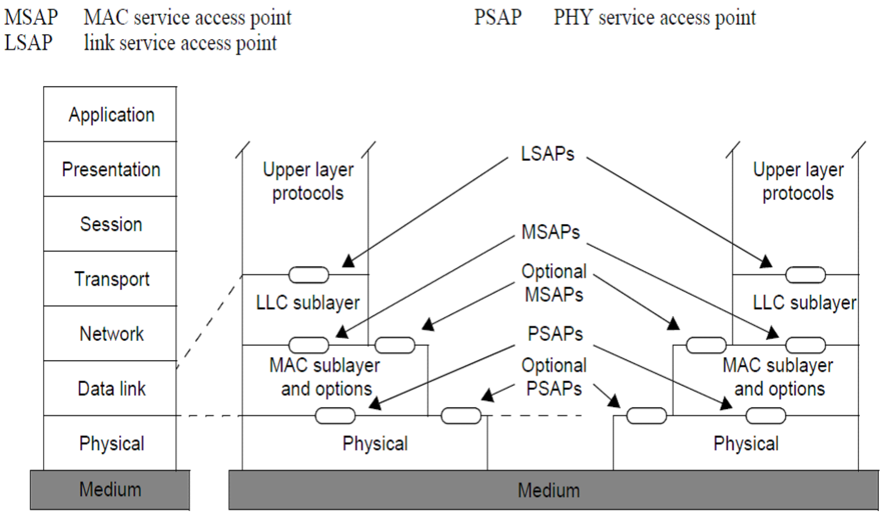

Figure 1—Title?

6.1.2 Service interface primitives, parameters, and frames

Each *N*-service is described in terms of service primitives and their parameters, each primitive corresponding to an atomic interaction between the *N*-service user and the *N*-service provider, with each invocation of a primitive by a service user resulting in the service issuing corresponding primitives to peer service users. The purpose of the model is to provide a framework and requirements for the design of protocols while not unnecessarily constraining the internal design of systems: primitives and their parameters are limited to, but include all of the information elements conveyed to corresponding peer protocol entities or required by other systems (and not supplied by protocols in lower layers) to identify (address) those entities and deliver the information. The parameters of service primitives do not include information that is used only locally, i.e., within the same system, to identify entities or organize resources.

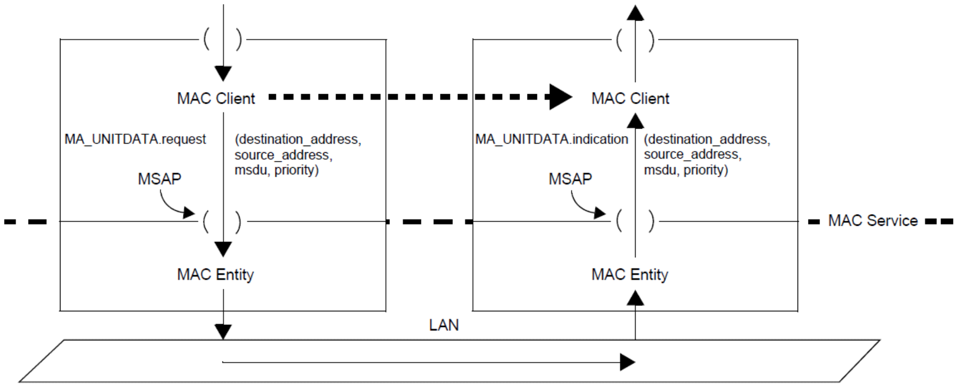
MAC Client
MAC Client
MA_UNITDATA.request
(destination_address, source_address, msdu, priority)
MA_UNITDATA.indication
(destination_address, source_address, msdu, priority)
MSAP
MSAP
MAC Service
MAC Entity
MAC Entity
LAN

Figure 2—MAC entities, the MAC service, and MAC service users (clients)

Figure 2 illustrates these concepts with reference to the MAC Sublayer, which contains MAC entities that provide the MAC Service at MAC Service Access Points (MSAPs), to MAC Service users.

The primitives of the MAC Service comprise a data request and a corresponding data indication; each with MAC destination address, MAC source address, a MAC service data unit comprising one or more octets of data, and priority parameters. Taken together these parameters are conveniently referred to as a frame, although this does not imply that they are physically encoded by a continuous signal on a communication medium, that no other fields are added or inserted by other protocol entities prior to transmission, or that the priority is always encoded with the other parameters transmitted.

6.1.3 Layer management interfaces

A given *N*-entity can have many associated management controls, counters, and status parameters that are not communicated to its user's peers, and whose values are either not determined by its user or not required to change synchronously with the occurrence of individual *N*-service primitives to ensure successful (*N* + 1) -protocol operation. Communication of the values of these parameters to and from local entities, i.e., within the same system, is modeled as occurring not through service primitives but through a layer management interface (LMI). One protocol entity, for example an SNMP entity, can be used to establish the operational parameters of another. Communication of the results of authentication protocol exchanges to entities responsible for controlling and securing access is one of the uses of LMIs in this standard.

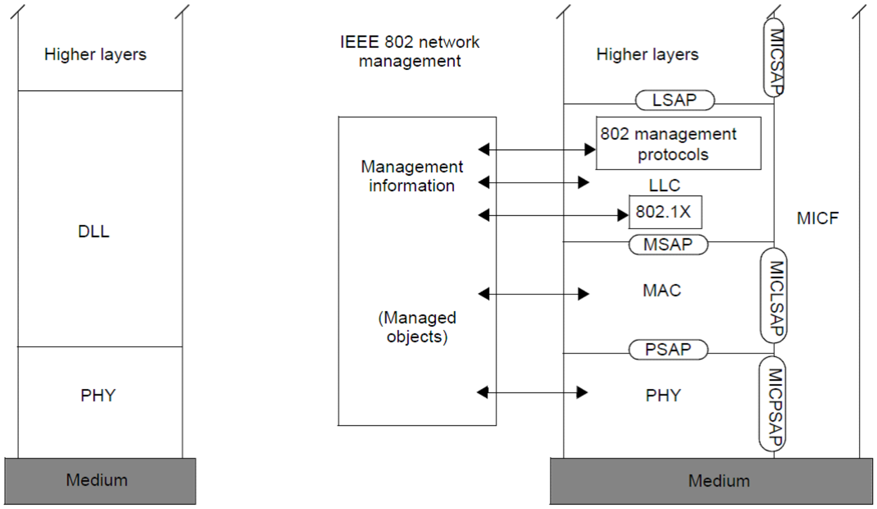

Figure 3—IEEE 802 reference model with end-station management

Figure 3 illustrates the layer management interfaces allowing access to controls, counters, and status parameters inside a protocol entity.

6.1.4 Service access points, interface stacks, and ports

Each service is provided to a single protocol entity at a service access point (SAP) within a system. A given *N*-entity can support a number of *N*-SAPs and use one or more (*N* - 1)-SAPs. The service access point serves to delineate the boundary between protocol specifications and to specify the externally observable relationship between entities operating those protocols. A service access point is an abstraction, and does not necessarily correspond to any concrete realization within a system, but an *N*-entity often associates management counters with the SAP and provides status parameters that can be used by the (*N* + 1)-entity using the SAP. Examples include the MAC_Operational and operPointToPointMAC status parameters provide by MAC entities.

The network and link layers of the reference model accommodate many different real networks, subnetworks, and links with the requirements for bandwidth, multiplexing, security, and other aspects of communication differing from network to network. A given service, e.g., the MAC Service, is often provided by a number of protocols, layered to achieve the desired result. Together the entities that support a particular service access point compose an interface stack.

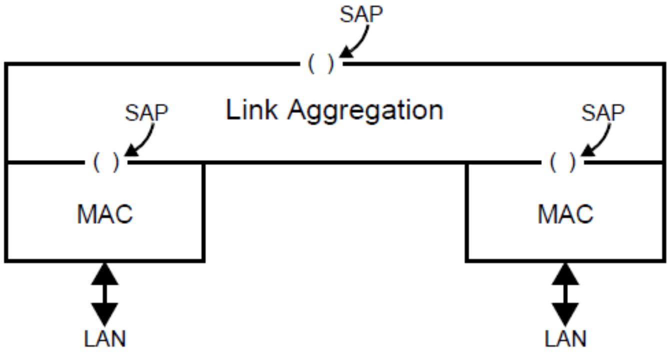

Figure 4—n interface stack

Figure 4 provides an example, showing the use of Link Aggregation (IEEE Std 802.1AX).

The term *port* is used to refer to the interface stack for a given service access point. Often the interface stack comprises a single protocol entity attached to a single LAN, and port can be conveniently used to refer to several aspects of the interface stack, including the physical interface connector for example. In more complex situations—such as that illustrated in Figure 4, where the parts of the interface stack provided by the IEEE 802.3 MAC entities effectively compose two ports that are then used by link aggregation to provide a single port to its user—the port has to be clearly specified in terms of the particular service access point supported. Port-based network access control secures communication through that service access point.

6.1.5 Media independent protocols and shims

Some protocols, such as those specified in IEEE Std 802.3, IEEE Std 802.11, and other IEEE 802 standards, are specific to their LAN media or to the way access to that media is controlled. Other protocols and functions within the MAC sublayer, such as link aggregation and bridging, are media independent—thus providing consistent management and interoperability across a range of media.

IEEE 802.1 standards use the term *shim* to refer to a protocol entity that provides the same service to its user as it uses from its provider (see 3.168 of IEEE Std 802.1Q-2011). Shims can be inserted into an interface stack to provide functions such as aggregation (e.g., IEEE Std 802.1AX), security (e.g., IEEE Std 802.1AE), or multiplexing.

6.1.6 MAC Service clients

The protocol entity that uses the service provided at a MAC Service access point (MSAP) is commonly referred to as the client of the MAC Service or of the entity providing the service. Within a Bridge, the MAC Relay Entity is a client of the Internal Sublayer Service (ISS), and the Logical Link Control (LLC) Entity is a client of the MAC Service. The LLC Entity is described in IEEE Std 802 and provides protocol identification, multiplexing, and demultiplexing to and from a number of clients that use a common MSAP. The clients of LLC are also often referred to as clients of the MAC.

6.1.7 Stations and systems

A LAN station comprises a single media access method specific entity, operating the MAC procedures specified in the applicable IEEE 802 standard, together with other protocol entities mandated by those standards (e.g., an LLC Entity) or commonly used in conjunction with that entity.

A system is a collection of hardware and software components whose intercommunication is not directly externally observable and outside the scope of the IEEE 802 standards that specify the system behavior as a whole. Management of a system, when supported, is typically provided through a single management entity. A system (such as a bridge) can contain many media access method specific entities, of the same or a variety of types, attached to different LANs. A system can therefore be said to include one or more LAN stations.

6.1.8 Connectionless connectivity and connectivity associations

The MAC Service supported by an IEEE 802 LAN provides connectionless connectivity; i.e., communication between attached stations occurs without explicit prior agreement between service users. The potential connectivity offered by a connectionless service composes a connectivity association that is established prior to the exchange of service primitives between service users (see RFC 787). The way in which such a connectivity association is established depends on the particular protocols and resources that support it, and can be as simple as making a physical attachment to a wire. However simple or complex, the establishment of a connectivity association for connectionless data transfer involves only a two-party interaction between the service user and the service provider (though it can result in exchanges between service-providing entities in several systems) and not a three-party user-service-user interaction as is the case for connection-oriented communication. With the continual increase in the number of ways that IEEE 802 LAN connectivity can be supported, it is no longer useful to regard a LAN as a definite set of physical equipment. Instead, a LAN is defined by the connectivity association that exists between a set of MSAPs.

6.2 Overview of IEEE 802 Network Reference Model

The network reference model defines a generic foundation for the description of IEEE 802 access networks, which may include multiple network interfaces, multiple network access technologies, and multiple network subscriptions, aimed at unifying the support of different interface technologies, enabling shared network control and use of software-defined networking (SDN) principles.

It adopts the generic concepts of SDN by splitting the network model into a data plane and a control and management layer with well defined semantics for interfacing with higher layer management, orchestration, and analytics functions. Additionally the model deploys a clear separation of functional roles in the operation of access networks to support various deployment models including leveraging wholesale network services for backhaul, network sharing, and roaming.

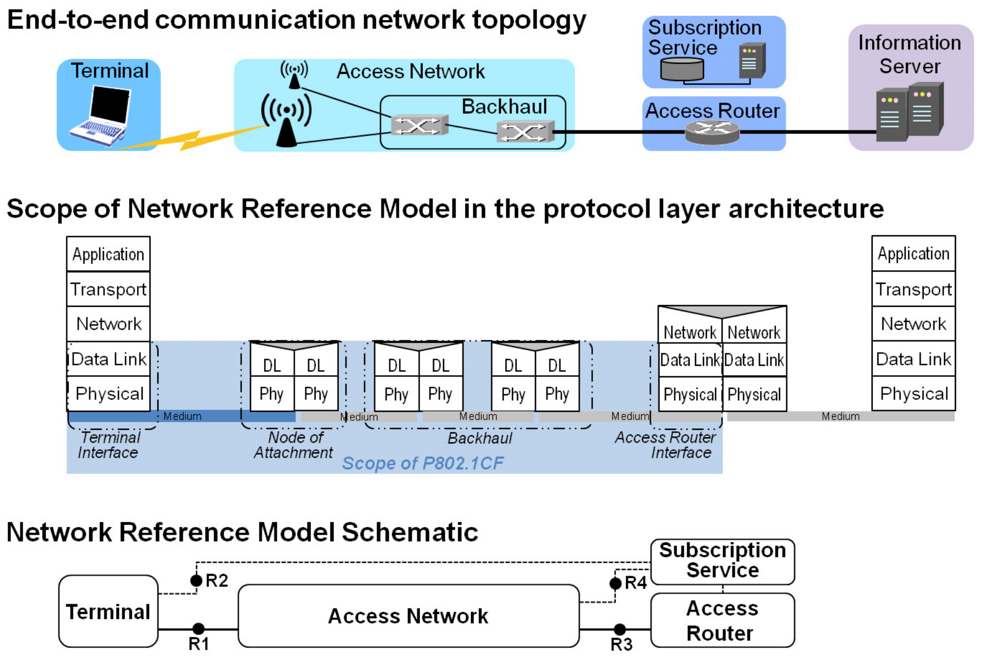

Figure 5—NRM overview

Within the bigger picture of an end-to-end network model for providing access to IP services, the NRM deals in particular with the link layer communication infrastructure between the host in the terminal and the access router in the core network as depicted in Figure 5.

In IEEE 802 access networks, the user data is forwarded according to the destination MAC address in the Ethernet frames, which represent the endpoints of the link in the access network. Avoiding a functional separation of the user plane from the transport plane, the specification provides an integrated model for backhaul connectivity combined with subscriber specific connectivity functions as facilitated by modern IEEE 802.1 bridging technologies. At first glance, the network model for an IEEE 802 access network consists of the terminal, the access network comprising the node of attachment and the backhaul, the access router and the subscription service, which provides authentication, authorization, accounting as well as policy functions specific for particular user accounts and terminals. Beyond the access router and out of scope of this specification is the infrastructure providing IP-based information services to the terminals.

Communication interfaces between the entities are denoted by R1 for the interface between the terminal and the node of attachment, by R2 for the authentication procedures between terminal and subscription service, by R3 for the interface between access network and the access router, and by R4 for the authentication, authorization, accounting, and policy functions between the access network and the subscription service.

6.3 Functional Entities

6.3.1 Terminal

The terminal is a mobile device that seeks connectivity to a communication infrastructure to get access to communication services. The terminal comprises a terminal interface building the physical port for connectivity, and eventually deploys a terminal controller for dealing with particular parameters and configurations conveyed by the control and management interface.

6.3.2 Access Network

The access network consists of the Nodes of Attachment providing the physical ports towars the terminals and the Backhaul for connecting the Nodes of Attachment toward the access router. The access network may deploy a dedicated access network controller for configuration and management of the elements inside the access network as well as exchange of control and management information with both the terminal and access router.

6.3.3 Access Router

The access router terminates the layer 2 connectivity to the terminal by realizing the anchor for network layer communication toward the terminal side. The access router comprises an access router interface that establishes the physical port of the connectivity toward the access network, and may eventually comprise a dedicated access router controller that handles and exchanges layer management information and configurations. With a dedicated access router controller,the access router becomes a logical functional unit with various implementation options for the controller and the packet forwarding engine attached to the access router interface.

6.3.4 Subscription service

The subscription service provides authentication, authorization, and accounting services (as well as user-specific policies) to the terminal, the access network, and the access router. The subscription service usually comprises a database containing all the subscription-specific information. Multiple subscription services may be interlinked with each other for roaming users, i.e. for subscribers, who make use of network resources not belonging to their own business.

6.3.5 Coordination and Information Service

The coordination and information service in an entity that coordinates the use of common resources and exchange of operational parameters among multiple access networks.

6.4 Basic Network Reference Model

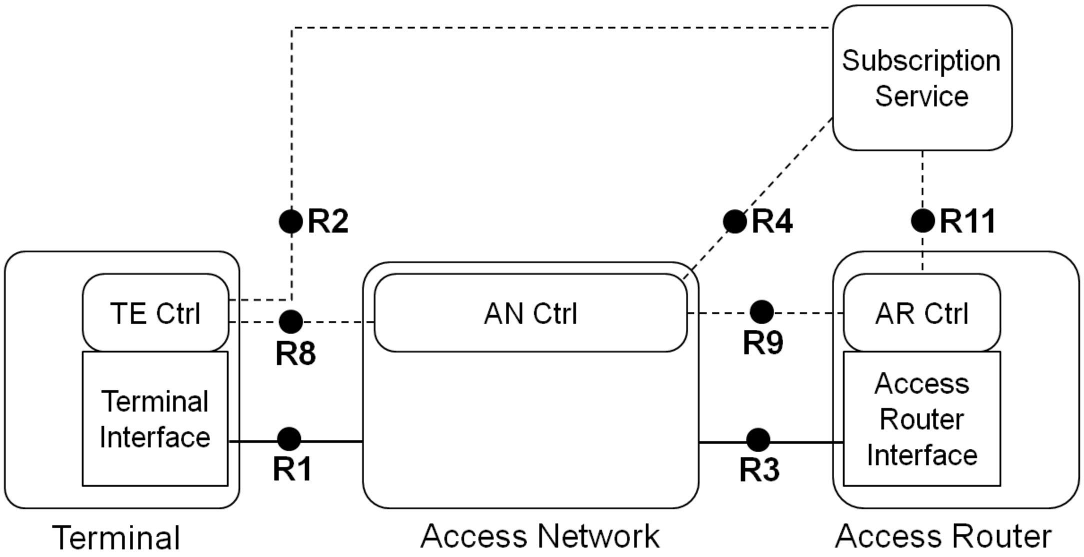

Figure 6—Basic Network Reference Model

Figure 6 presents the Basic Network Reference Model. Solid lines represent the interfaces representing the data plane and connecting ports, while dotted lines show the flow of control and management information. This NRM is the foundation for further refinements and includes the basic differentiation between functional entities and the reference points for their communication. The Basic NRM is composed of four main elements: i) the Terminal (TE), ii) the Access Network (AN), iii) the Access Router (AR), and iv) the Subscription Service (SS).

As depicted in Figure 6, the TE, the AN, and the AR each comprise a control entity, which is denoted Controller (Ctrl). Each of the three elements has its specific Controller.

Note—The access router is a logical functional unit with various options for implementation depending of the design and architecture of the access router controller.

Note—Please note that currently no assumptions are made regarding the ownership of the functional units. Access Network, Subscription Service, and Access Router may belong to the same operator, or may be distributed among three distinct operators.

6.4.1 Reference Points

R1 represents the reference point for the PHY and MAC layer functions establishing the physical port, as specified in numerous IEEE 802 standards, between terminal and access network

R2 represents a control interface between terminal and the subscription service, e.g. for authentication.

R3 represents the physical port for the communication between the access network and the access router.

R4 represents a control interface communicating subscription-specific information elements between the access network controller and the subscription service.

R8 represents the control and management interface between the AN and the TE, which terminates in Access Network Controller and the Terminal Controller, respectively. The functionalities of this reference point are related to the configuration of the physical port in the terminal and the control of the data flows in the terminal. In addition, the reference point may include some additional configuration parameters to influence the behavior and configuration of the terminal.

R9 represents a control and management interface between the access network controller and access router controller.

R11 represents a control interface communicating subscription-specific information between the subscription service and the access router controller.

6.5 Network Reference Model including Coordination and Information Service

Figure 7—NRM with Coordination and Information Service

Some deployments include a Coordination and Information Service (CIS) to provide advanced services such as spectrum management, coexistence, and information services for mobility. The reference model includes the option for CIS by providing a reference point to communicate the information between CIS and the AN Ctrl, possibly propagated further by the AN Ctrl to the TE Ctrl and AR Ctrl over the R8 and R9 interfaces, respectively.

6.5.1 Reference Points

R10 represents a control and management interface between the Access Network Controller and the CIS.

6.6 Comprehensive Network Reference Model

The Network Reference Model comprises further details of interfaces inside the Access Network.

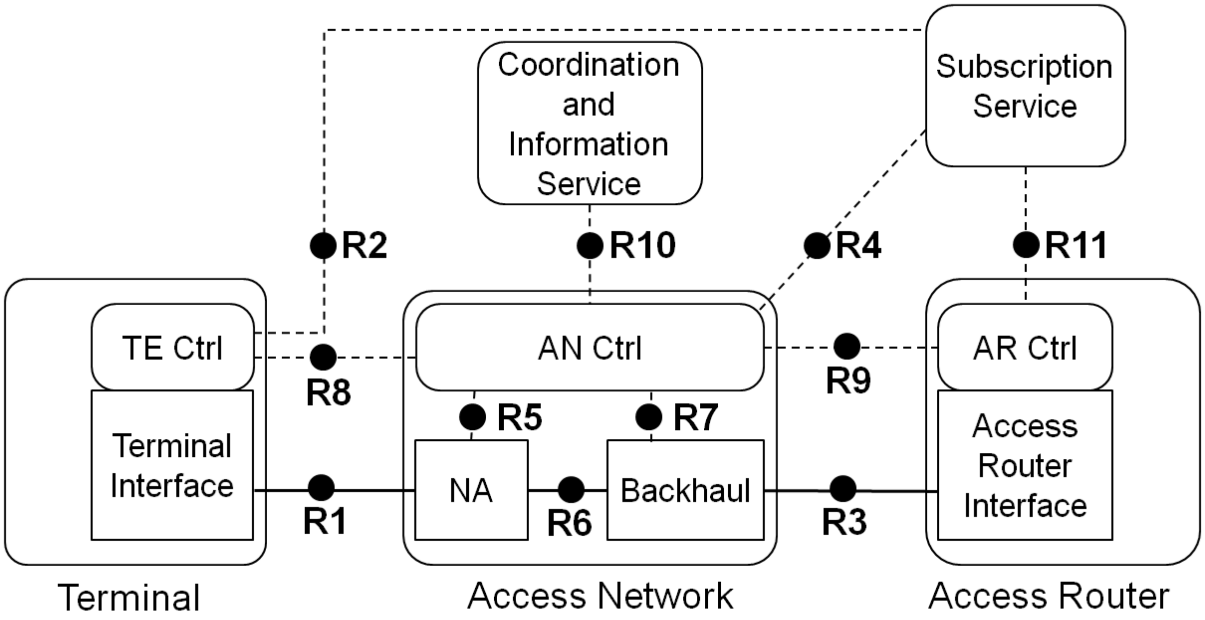

Figure 8—Network Reference Model exposing Access Network details

In Figure 8 the access network is decomposed into a Node of Attachment (NA) and the Backhaul (BH). The NA represents the entity providing the link to the terminal, the interface to the backhaul, and the data forwarding function between these two. The connections between NA, backhaul, and AN control are described by reference points R5, R6, and R7.

6.6.1 Reference Points

R5 represents a control-only interface for the configuration and operation of the node of attachment. It includes information elements for the configuration of the R6 port toward the backhaul, the R1 port toward the terminal, and the data-forwarding functions inside the Node of Attachment.

R6 represents a reference point representing[?] the physical ports between the node of attachment and the backhaul.

R7 represents an interface used to control and configure the user plane within the backhaul. The backhaul interconnects the NAs with the Access Router.

7. Functional Design and Decomposition

7.1 Access Network Setup

7.1.1 Dynamic Spectrum Analysis

7.2 Access Network Discovery and Selection

7.2.1 Introduction

Access network discovery and selection describes the process of exploring the surrounding environment for detection of available access networks, followed by retrieval of information about each of the access network, and finally the evaluation of the collected information in order to determine the most appropriate Node of Attachment to connect to.

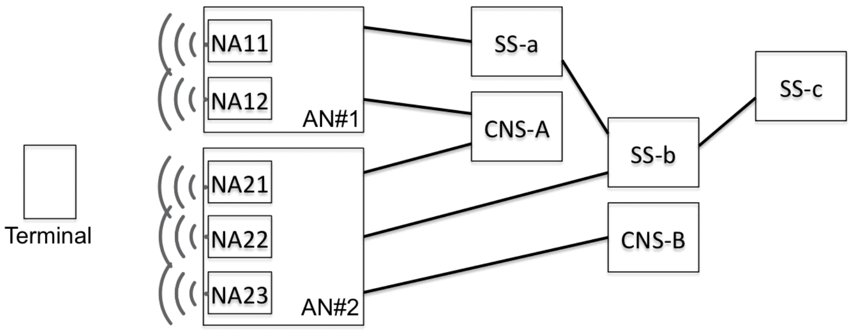

Figure 9—Network discovery scenario with multiple SSs and CNSs

The process is usually executed either when a terminal performs its initial network entry after power on, or when a terminal lost or is going to lose its network connectivity and prepares for re-entry at another node of attachment, or when a terminal moves across an access network coverage area built by multiple nodes of attachment and the terminal relocates the link to another point of attachment to maintain best possible network connectivity during the move.

7.2.2 Roles and identifiers

7.2.2.1 User

User represents the unique identity of a subscription. A user may have subscriptions with one or more subscription services. Unique subscription identifiers are created by an username amended by the identity of the subscription service.

ID of User: Subscription Identifier {NAI} + Subscription Name {String}

7.2.2.2 Terminal

Terminal represents the physical device communicating with the core network service making use of an access network to establish the link. A unique identifier is assigned to each of the terminals.

ID of Terminal: {EUI48} or {EUI64}

7.2.2.3 Node of Attachment

Node of attachment is the physical device at the edge of the access network creating the communication link to the terminal. Different NAs may have different capabilities.

ID of Node of Attachment: {EUI48} or {EUI64}

7.2.2.4 Access Network

Access network denotes the infrastructure consisting of one or more Nodes of Attachment and the related backhaul for providing the communication links between the nodes of attachment and one or more interfaces to connected core network services.

ID of Access Network: ANI {EUI-48} + AN Name {String}

7.2.2.5 Subscription Service

The subscription service is the entity establishing and maintaining user-specific configuration and usage data. For security reasons, the subscription service performs authentication of the corresponding terminal to ensure that usage and modification of user-specific information is really caused by that user. Subscription service is commonly known as termination point of AAA.

ID of Subscription Service: SSI {FQDN} + SS Name {String}

7.2.2.6 Core Network Service

Core network service denotes the termination point of the user plane of a terminal. Multiple terminals may connect to the same core network service, but there may be several core network services available by an access network. Selection of the corresponding core network service may happen through authorization by the subscription service eventually amended by signaling from the terminal.

ID of Core Network Service: CNS Identifier {??? - ffs} + CNS Name {String}

7.2.3 Use Cases

Network discovery and selection is a prerequisite for a mobile terminal to establish and maintain network connectivity. A terminal initiates the network discovery and selection process for the following four reasons.

7.2.3.1 Initial AN access

Initial AN access describes the case when a terminal is powered up or the network interface of the terminal is enabled and network connectivity initially does not exist without any prior knowledge about the availability of NAs.

In this case the terminal usually performs a complete network discovery process to learn about all reachable NAs before executing the selection process from the root.

7.2.3.2 AN re-entry

In this case the terminal has lost, or has not yet established, network connectivity, but has some stored information about the last AN and the last NA to which it was connected. When selection policies prefer to re-establish connectivity to the last used AN, the terminal will try to execute an abbreviated NDS process by directly checking for the reachability of the last used NA. This process optimization makes particular sense when the access technology allows for active scanning, resulting in much faster network connectivity establishment.

When AN re-entry is not possible due to movement of the terminal completely out of the previously used coverage area, the terminal will perform an initial AN access process. Statistically, however, performing a AN re-entry trial before falling back to an initial AN access provides benefits, even when the worst case lasts longer than going straight into an initial AN access process.

7.2.3.3 NA transition

The network discovery and selection process is initiated not only when network connectivity is missing but also when the terminal detects degradation of network connectivity that endangers loss of connectivity. In this case the terminal provisionally searches for another NA offering better link conditions than the NA to which it is currently connected.

When another NA of the same AN with better link conditions exists, the terminal will initiate a relocation of its ongoing network connectivity to the other NA while maintaining all upper-layer connectivity states. Such a transition is commonly denoted as seamless handover.

7.2.3.4 AN transition

When connectivity is in danger but seamless handover to another NA of the same AN is not possible, the terminal will carry through a discovery process for other ANs allowing for network connectivity. Usually the transition of ongoing connectivity to another AN will cause some disruption. How long connectivity is broken, and whether upper-level connection state can be maintained, depend on the particular AN arrangements and implementations.

Usually interruption of connectivity during AN transition is much longer than during NA transition, but often much less severe than for initial AN access, which completely resets the whole communication stack.

7.2.4 Functional requirements

The following requirements apply to the NDS procedures.

7.2.4.1 Support for multiple access technologies

The NDS procedures SHOULD be able to handle, within the same terminal, various access technologies with different characteristics.

7.2.4.2 Support for multiple different access networks supporting the same or different subscription services

The NDS procedures SHOULD to able to handle multiple different access networks based on the same or different access technologies serving the same or different subscription services.

The NDS procedures SHOULD support access networks served by multiple subscription providers.

7.2.4.3 Support for multiple subscriptions on the same access technologies.

The NDS procedures SHOULD support multiple different subscriptions on the same access technology and/or the same access network as well as the usage of the same subscriptions on multiple different access technologies.

7.2.4.4 Extensibility to support specific service requirements

The NDS procedures SHOULD support upper-layer service-specific attributes to enable different treatment of various access technologies and access networks depending on service requirements.

7.2.4.5 Discovery of access network capabilities

The NDS procedures SHOULD NOT require establishing *a priori* knowledge within the terminal about offered services of the existing access networks to perform the selection process.

The discovery procedures SHOULD allow retrieving service-specific attributes.

7.2.5 NDS specific attributes

Each of the entities involved in the NDS process comprises information elements, which are helpful or required when processing the NDS procedures. The following list provides examples of these information elements:

7.2.5.1 User (Subscription)

Access policies

7.2.5.2 Access Network

Supported Subscription Services

    LIST of Subscription Service IDs

    Cost, limitations per

Supported Core Network Services

    LIST of Core Network Service IDs

AN certificate

    CERTIFICATE

Access Network Capabilities

    LIST of Link Layer capabilities

        E.g. MTU, encryption, type of link, privacy

    RECORD of Link Layer performance parameters

        E.g. supported service classes (Throughput up/down, delay, jitter)

7.2.5.3 Subscription Service

Supported Core Network Services

LIST of Core Network Service IDs

SP certificate

CERTIFICATE

7.2.5.4 Core Network Service

Network Layer Capabilities

LIST of Capabilities

E.g. IP versions, configuration, service discovery support

Network Interface performance

LIST of performance parameters

E.g. supported service classes (throughput up/down, delay, jitter)

Offered application services

LIST of application services

E.g. Internet, Voice, Printer, File service

7.2.6 NDS basic functions

7.2.6.1 NA Discovery

NA discovery is the process in the terminal to retrieve the list of nodes of attachment, which can be reached via the physical medium. The discovery process executed is specific for a particular access technology, but a terminal comprising multiple different network interfaces may initiate and perform the process concurrently on all or on a subset of its network interfaces.

NA discovery can be based on either passive scanning or active scanning.

When performing a passive scan, the terminal turns on the receiver path of its network interface and "listens" sequentially to all channels of the medium for messages indicating the existence of an active Node of Attachment. A complete scan may take quite some time depending on the periodicity of the indication messages and the number of channels. When sped up by methods taking *a priori* knowledge into account, the process of passive scanning may deliver specific or initial results earlier, but a complete scan always takes the time of periodicity of indication messages by number of channels. As passive scanning of radio does not emit any radio waves, the approach complies with any radio regulation framework.

Active scanning comprises a trigger sent out by the terminal to initiate directed responses of nodes of attachment. By its nature, active scanning is able to deliver results much faster but requires the terminal to transmit information frames on all channels of the network interface. Before sending out frames the terminal may be required to determine the regulatory domain in which it is operating to ensure that transmissions comply with the applicable regulatory requirements.

NA discovery provides a list of nodes of attachment reachable by the terminal at its particular location.

7.2.6.2 AN Detection

AN detection is the process to determine the identities and the capabilities of the access networks in reach. The terminal retrieves, for each of the detected NAs, the identity of the access network to which the NA belongs.

Further information about capabilities of the detected ANs—like networking and performance parameters, as well as supported subscription and core network services—is derived either from broadcast advertisement information from a preconfigured local database, or from queries to remote databases. Remote databases may be available over specific link procedures in the NAs or access networks, or even over network connectivity anywhere in the network, when some other connectivity exists during the AN detection process.

7.2.6.3 SS Detection

SS detection is the process to determine the subscription services, which can be used for establishment of access to the detected ANs. The process creates a list of all available subscription services, with information about the availability and preference of subscription services for each of the detected ANs.

The information about available subscription services is usually taken from the information collected during the AN detection, but there may be pre-stored information in the terminal together with the authentication credentials which provides the list of ANs to be used with the credential.

7.2.6.4 CNS Detection

CNS detection is the process to retrieve the core network services, accessible through the detected access networks. The process establishes a list of all available core network services, with information about the availability and preference of subscription services for each of the detected core network services.

The information about available core network services is usually taken from the information collected during the AN detection, but there is usually information available in the terminal as part of the subscription, which amends the information derived from the AN detection process.

7.2.6.5 SS and CNS Selection

SS and CNS selection is a multi-dimensional selection process in the terminal making the best choice among the detected subscription services and core network services under the preferences, restrictions, and limitations imposed by the available subscriptions. The selection process may perform a weighted evaluation of all available information down to interface parameters of the physical link to the point of attachment.

The selection process may be either hard-coded in the terminal as part of the operating software, or be configurable by policies provisioned to the terminal.

7.2.7 Detailed procedures

7.2.7.1 First-time use of TE without subscription

The TE performs in steps a) through c) anNDS procedure to find appropriate SS for creation of a new subscription. Online subscription set-up is performed in steps d) through e).
[What about f)?]

a) TE runs NA Discovery and AN Detection, and finds one or more available ANs.

b) TE runs SS Detection and CNS Detection, and finds available SSs and CNSs, and their associations with the ANs.

c) TE performs SS and CNS Selection, and determines an AN and a SS based on defined preference criteria for running the subsequent online subscription set-up.

d) TE performs a special connection procedure with the selected AN for establishment of a subscription.

e) TE creates a trust relationship enabling network access authentication and authorization by the selected SS.

f) TE acquires and stores the subscription of the selected SS.

7.2.7.2 Initial AN access

The TE is equipped with one or more subscriptions, and attempts to establish a network connection after being switched on or moved into a coverage area.

a) TE runs NA Discovery and AN Detection, and finds one or more available ANs.

b) TE runs SS Detection and CNS Detection, and finds available SSs, CNSs, and their associations with the ANs.

c) TE performs SS and CNS Selection according to the provisioned subscriptions, and determines the preferred AN and SS for establishing network connectivity.

d) TE performs a network entry procedure toward the selected AN, making use of the selected SS for authentication and authorization.

7.2.7.3 NA transition

The TE discovers that the link to the current NA is getting weak and decides to pursue a transition to another NA of the same AN to maintain good link quality.

a) TE runs NA Discovery, and finds one or more other NAs belonging to the same AN to which the TE is currently connected.

b) TE selects the NA for transition and performs a network entry procedure to new NA making use of the currently used subscription and SS for authentication and authorization of access. Ifsupported by access technology for faster handover, the TE may pre-establish the connectivity to the new NA through messaging with the AN via the current NA.

c) When connectivity to new NA is established, the TE turns down connection to previous NA.

In the case of failure, TE reverts to initial AN access.

7.2.7.4 AN re-entry

The TE recently lost network connectivity, and finds by NA discovery that NAs of the same AN are accessible. To re-establish network connectivity with same SS and CNS, the TE attempts to connect to NA of previously used AN.

a) TE runs NA Discovery and finds one or more other NAs belonging to the same AN, to which the TE was previously connected.

b) TE selects the NA for connection establishment and performs a network entry procedure to the NA, making use of the previously used subscription and SS for authentication and authorization of access.

c) Depending of duration of the connectivity break, the TE may or may not attempt to resume the previous communication link to the CNS.

In case of failure, TE reverts to initial AN access.

7.2.7.5 AN transition

The TE discovers that the link to the current NA is getting weak and decides to pursue a transition to the NA of another AN, which provides service to the same SS and CNS as currently used.

a) TE runs NA Discovery and AN Detection, and finds that there is another AN, with service to the same SS and CNS, that would provide better link quality.

b) TE decides to transition network connectivity to the other AN for continuation of service to the current SS and CNS.

c) TE selects the NA for transition and performs a network entry procedure to new NA, making use of the currently used subscription and SS for authentication and authorization of access, and requesting connectivity to the currently used CNS.

d) Depending of the capabilities of the TE and the CNS, the TE may or may not attempt to resume the communication link to the CNS.

In case of failure, TE reverts to initial AN access.

7.2.8 Mapping to IEEE 802 technologies

Table 1—Title?

		802.3	802.11	802.15.?	802.16	802.22
Identifiers	TE	EUI-48	EUI-48	EUI-64	EUI-48	EUI-48
	NA	EUI-48	EUI-48	EUI-64	EUI-48	EUI-48
	ANI	???	EUI-48	???	EUI-48	EUI-48
	AN-name	256 Char	30 Char	???		
Subscription Type		NAI	NAI/PSK	???/PSK	NAI	NAI
Multiple SSs		Info	ANQP	-	?	-
Discovery process		manual	passive, active	passive, active	passive	passive

7.2.9 Additional capabilities in IEEE 802 technologies

7.2.9.1 IEEE 802.3

For further study.

7.2.9.2 IEEE 802.11

IEEE 802.11 provides a number of functional enhancements to support more complex deployments:

- Access Network Query Protocol
- Pre-Association Discovery Protocol
- Network triggered NDS
    E.g., Directed NA transition

- Online subscription establishment

  E.g., Hotspot 2.0 “Online Sign Up’”

7.2.9.3 IEEE 802.15

For further study.

7.2.9.4 IEEE 802.16

For further study

7.2.9.5 IEEE 802.22

For further study

7.3 Association and Disassociation

7.4 Authentication and Trust Establishment

7.5 Datapath Establishment, Relocation, and Teardown

7.6 Authorization, QoS, and Policy Control

7.7 Accounting and Monitoring

8. SDN Abstraction

SDN Functional Decomposition

8.1 Introduction

Software Defined Networks (SDN) is a new paradigm based on the splitting of control and data planes of networking elements. Basically it works by pushing the intelligence related to the operation of a certain service to a central controller, while the data path (user data packets) is handled based on the orders of the central controller in separate and specialized elements. Within the IEEE 802 set of technologies, there are multiple functionalities that can be designed based on the SDN paradigm. This document presents several use cases showcasing these functionalities:

- Setup of interfaces and nodes
- Detection of node attachment
- Path Establishment
- Path Teardown
- Path Maintenance

- Path relocation
- Affecting the behavior of Coordination and Information System
- Configuration of connection between the Core Network and the Access Network Event handling
- Statistic gathering

8.2 Roles and identifiers

Controller: Application that manages different behaviors (e.g., flow control) in a Software Defined Networking environment.

Data path element: Hardware/Software entity in charge of executing the orders from the controller, affecting the path through which data is forwarded.

8.3 Use Cases

8.3.1 Setup of interfaces and nodes

Through SDN a central controller can implement a control logic enabling it to configure several parameters in the nodes and interfaces of the data path. Within the possible set of configuration parameters there are three main families:

- SDN control configuration
- Short time-scale configuration
- Long time-scale configuration

SDN control configuration refers to the required setup of the parameters ruling the communication between the data path element and the controller. These parameters may include IP address of the controller, kind of protocol, VLAN or interface used for the communication, and certain timers governing the transmission of keepalive messages or the teardown of it [it? the data path?] in case of failure. These configuration parameters must also include the different timers, ports, and protocols used for the communication between the controller and the data path element.

Short time-scale configuration refers to the configuration of parameters that may change in very short time scales. For example, transmission power, MAC QoS parameters, antenna selection, etc.

Long time-scale configuration refers to the long-term configuration of the node or interface; the parameters used by the controller are the typical ones an OAM system may use. Examples of these parameters include the operational frequency, configuration regarding credentials or authentication servers to use, supported authentication modes, VLAN configuration, etc.

8.3.2 Detection of terminal attachment

A very important operation required to use SDN control in the access network is the ability to detect the attachment of new terminals. The user's terminal typically will not include any kind of SDN software or contain the functionality to detect that it is connecting to a network using an SDN controller. Therefore, some mechanism is needed to handle the detection of the terminals while attaching to the network PoAs. In a PoA where the wireless interface (IEEE 802.11) is bridged to a switch, this detection of terminal attachment can be done thanks to the switch sending an LLC SNAP message upon attachment of a new terminal. In other technologies some other mechanisms should be analyzed.

8.3.3 Data Path Establishment

An IEEE 802.1CF network does not include IP layer, hence path establishment mechanisms using above-layer-2 information are outside the scope of this document. In order to establish a path, a controller must be informed of the new flow, including its requirements in terms of capacity, delay, and jitter. After receiving this information the controller can compute the best possible path and communicate this decision to the data path elements. After this, the data path elements will enforce the controller decision on the different packets traversing the data path.

In order to perform the data path establishment, the following information/functionality is required:

- Topology information
- Mechanisms to compute best path based on some criteria
- Communication mechanisms to set specific rules in the data path to decide output port/modifications to frame
- The data path must support some mechanism for packet matching. This mechanism can be arbitrarily complex. It can include simple mechanisms such as input port matching or VLAN tag matching, or complex rules indicating logical combination of parameters, including internal state of the data path element.
- The data path must support the application of forwarding rules to the input traffic. In this way the decisions taken by the controller will be applied and the packet will be sent through the appropriate output port.
- The data path element may support actions over the packets and internal state. The data path element may be able to modify certain parts of the packet and modify internal state variables, such as counters, monitoring variables, etc..

8.3.4 Data Path Teardown

A certain path can be created for a specific flow. Once the flow finishes, there is no need to have the path established any longer, freeing resources allocated to the path. In order for the controller to tear down a path, it first needs to determine that the path can be deactivated. This can be done through monitoring metrics or flow information coming from the flow originator. Once the controller knows that the path can be removed, it can communicate its decision to the data path elements. At that time, all data path elements should remove the stored state corresponding to the path.

8.3.5 Data Path Maintenance

Typically a data path element will configure forwarding rules with a certain lifetime. The rules must be updated within their lifetime, or the data path element will remove them. Upon expiration of the rule, the data path element should inform the controller about the removal of the rule. In this way, the controller can keep records of the current status of the paths in the network.

8.3.6 Control Path Maintenance

In the same way as with data paths, the communication between the controller and the different data path elements must be kept alive through the exchange of some control packets. The actual configuration of the timers to use should be one of the parameters considered in 8.3.1.

[Check ref; no autoupdatinghere]

8.3.7 Path relocation

Due to reasons such as traffic engineering, movement of the terminal, or QoS degradation, it may be necessary to relocate a data path. Relocation isintended to change the data path elements the data path goes through, while keeping the most similar allocation of resources. This functionality can be divided in a sequence of Data Path establishment and Data Path Teardown.

8.3.8 Affecting the behavior of Coordination and Information System

Terminals and network nodes can relay on Coordination and Information Services (CIS) to gather information helping them to make some decision, such as candidate network selection, channel to use, etc. A controller may interact with the CIS in a standalone way, or it may mediate the communication between the terminal and the CIS. This latter approach allows the controller to modify, apply policies, add more information, or simply query different servers based on terminal information such as its user profile.

8.3.9 Configuration of connection between the Core Network and the Access Network

TBD

8.3.10 Event Handling

TBD

8.3.11 Statistic gathering

TBD

8.4 Functional requirements

The following requirements apply to the SDN procedures.

8.4.1 Support of a control connection between the different data path elements and the controller

Elements in the network subject to communicating with or being controlled by a controller SHOULD use a secure control connection for the communication. This includes the terminal in case it communicates with the controller.

8.4.2 Support for data path elements with heterogeneous technology interfaces

Controllers SHOULD support the configuration of parameters for multiple technologies. Abstract parameters common for multiple technologies SHOULD be used when possible. The data path element SHOULD provide common controlled behaviors to all the interfaces attached to it, regardless of their technology.

8.4.3 Support of communication mechanisms between the terminal and the controller

The terminal SHOULD use a secure communication channel to communicate with the controller.

8.4.4 Support of per-packet matching, forwarding rules, and actions in the data path element

Data path elements SHOULD include mechanisms for packet matching, forwarding rules, and actions (packet modifications).

8.4.5 Support of state recording in the data path elements

Data path elements SHOULD be able to store operational parameters so they can be retrieved for monitoring.

8.4.6 Support of security associations between the controller and the data path elements (including terminal)

TBD

8.4.7 Support of security associations between controllers belonging to the AN and CN, which communicate

TBD

8.4.8 Support of security associations between AN controllers and CIS servers

TBD

8.5 SDN specific attributes

This section lists possible parameters for the different functions involved in the SDN operation.

8.5.1 Abstract parameters

- Supported Rates
- TxPower, TxPower levels supported
- Operational Frequency
- Statistics: Tx error, Rx error, Number of stations

8.5.2 Terminal Configuration

Terminal Controller

- LIST of control capabilities
- LIST of interfaces and their capabilities
- LIST of protocols to manage interfaces

E.g., interface X supports CAPWAP+OF

Interface

- Abstract parameters
- LIST of parameters to be configured

Technology specific: e.g., BSSID to connect to, RTS Threshold, short retry, long retry, fragmentation threshold, Tx/Rx MSDU lifetime, enable Block Ack, etc.

S Security parameters (technology dependent)

8.5.3 Access Network Configuration

8.5.3.1 Configuration of interfaces

- Abstract parameters
- LIST of parameters to be configured

Technology specific: e.g., BSSID, RTS Threshold, short retry, long retry, fragmentation threshold, Tx/Rx MSDU lifetime, enable Block Ack, etc.

Technology-specific security parameters: WPA/WPA2/WEP, parameters for key management, etc.

- Queue configuration: Capacity, max number packets, rate limitation

8.5.3.2 Configuration of nodes

- Parameters to configure the connection to controller:

Protocol+port

Credentials

Output physical port to use for connection to controller

ID

8.5.3.3 Configuration of data path ports

- VLAN configuration
- Number of tables

8.5.4 Data path establishment

- Matching rule and actions

- Matching rule definition and associated actions

8.5.5 Triggering technology-specific features

- Type for feature
  - E.g., send 802.11v frame, configure 802.11aa groupcast mode
- Content of feature
  - E.g., BSS to attach to, groupcast mode, concealment address, stations to be added

8.5.6 Interacting with CIS

- Parameters to enable the communication
  - Protocol to be used, credentials
- Adding/removing/modifying information at the CIS
  - CIS specific
    - E.g., IE elements to add to ANQP

8.5.7 Communication between CN and AN

TBD

8.5.8 Event handling

TBD

8.5.9 Statistics gathering

TBD

8.6 SDN basic functions

Controller discovery and configuration is outside the scope of IEEE 802.1CF.

8.6.1 Configuration of interfaces

Once a data path element or a terminal is attached to a controller, it can configure the different characteristics of the interface. A typical example of this can be taken from the world of IEEE 802.11, where WLAN controllers configure the different parameters of the technology. The typical parameters that the controller will set up are operational frequency and transmission power. Depending on the technology, the controller will also be able to configure additional parameters such as RTS threshold or ESSIDs/BSSIDs of the different APs. The actual configuration to be installed may come from different sources ranging from fully automatic algorithms computing the best allocation of, e.g., frequency/transmission power to static allocations.

8.6.2 Data path establishment/modification

Once the controller is connected to the data path elements, it must decide the path data flows will follow. Depending on the technology of choice (e.g., OpenFlow, MVRP, SNMP, etc.) the data path will be installed based on static rules such as port allocated VLANs, or it will be installed based on intelligent packet-matching rules. As a result of this operation, each data path element should have a rule stating the forwarding behavior for packets belonging to a certain flow. The computation of the path to be installed depends on the technology of choice for the controller, since there are technologies computing a path in a distributed way and technologies that can run traffic engineering+policying algorithms.

In addition to the proactive instantiation of a path described in the above text, a data path element may reactively interact with the controller due to some event, new packet arrival, or pre-installed rule among others. Following this, the data path element may interact with the controller in order to build a path for a specific new flow that has just appeared and, for any reason, should not be forwarded through the pre-configured paths.

8.6.3 Data path teardown

Once a data path is no longer in use, the rules indicating the forwarding behavior for each data path element may be removed. Generally this is done through lifetime timer expiration, but the controller can choose to remove the rules actively. Note that rules installed in a data path can also be permanent or semi-permanent, not requiring the refreshing of the controller.

8.6.4 CIS communication and controller as proxy for CIS

Nowadays CIS databases are filled with information provided by multiple sources but controlled by the operator. It would be desirable for the AN controller to be able to communicate with the CIS system in order to add,remove, or modify its information based on its knowledge about the network. One example would be to update the list of services being advertised in 802.11aq based on information obtained from the network.

In addition to this, a controlled network can be configured in such a way that the controller is used as proxy for the CIS communication. In this way the controller will get the answer from the CIS and can modify it accordingly to get some expected behavior in the network. For example, a controller may be used as proxy to access a MIIS and, after receiving the response, filter it to remove the surrounding networks that do not belong to a specific operator, in this way enforcing some policy in the terminal.

8.6.5 Triggering technology-specific functionality from the controller

Although SDN controllers have been used typically to just set up data paths in the network and configure characteristics of the interface, the complete possibilities of controlling the specific features of the technologies have not been yet analyzed. The use of technology-specific features by a controller can yield further advantages for the network. For example, a controller may configure the QoS across a mix of wireless and wired domains, by triggering the QoS configuration mechanisms of each technology. Another example may use management frames of IEEE 802.11v to control the point of attachment of the user terminal. To open all this functionality, the controller needs a clear view of the interfacecapabilities and new APIs to trigger it in a remote way.

8.6.6 Event Handling

TBD

8.6.7 Statistics Gathering

TBD

8.7 Detailed procedures

TBD

8.8 Functional Design and Decomposition

Annex A

(normative)

PICs proforma

Annex B

(informative)

Bibliography

Bibliographical references are resources that provide additional or helpful material but do not need to be understood or used to implement this standard. Reference to these resources is made for informational use only.

[B1] ISO/IEC 7498-1

[B2] IEEE Std 802.1AC

[B3] RFC 787