Open Networking Foundation

REFERENCE DESIGN

SDN Enabled Broadband Access (SEBA)

ONF TS-100

Version 1.0 | March 2019

About ONF Reference Designs

Reference Designs (RDs) represent a particular assembly of components that are required to build a deployable platform. They are "blueprints" developed by ONF's Operator members to address specific use cases for the emerging edge cloud.

RDs are the vehicles to describe how a collection of projects can be assembled into a platform to address specific needs of operators. By defining RDs, ONF's operator members are showing the industry the path forward to solutions they plan to procure and deploy.

Each RD is backed by specific Operator partner(s) who plan to deploy these designs into their production networks and will include participation from invited supply chain partners sharing the vision and demonstrating active investment in building open source solutions. The RD thus enables a set of committed partners to work on the specification and a related open source platform.

Assembling the set of selected components defined by the RDs into a platform enables a proof-of-concept to allow the test and trial of the design. These platforms are called Exemplar Platforms and each of them will be based on a Reference Design and will serve as reference implementations. These platforms are designed to make it easy to download, modify, trial and deploy an operational instantiation and thereby speed up adoption and deployment.

About the Open Networking Foundation

The Open Networking Foundation (ONF) is an operator led consortium spearheading disruptive network transformation. Now the recognized leader for open source solutions for operators, the ONF first launched in 2011 as the standard bearer for Software Defined Networking (SDN). Led by its operator partners AT&T, China Unicom, Comcast, Deutsche Telekom, Google, NTT Group and Turk Telekom, the ONF is driving vast transformation across the operator space. For further information visit http://www.opennetworking.org

Disclaimer

THIS DOCUMENT HAS BEEN DESIGNATED BY OPEN NETWORKING FOUNDATION ("ONF") AS A **FINAL SPECIFICATION** AS SUCH TERM IS USED IN THE ONF INTELLECTUAL PROPERTY RIGHTS POLICY.

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. WITHOUT LIMITATION, ONF DISCLAIMS ALL LIABILITY, INCLUDING LIABILITY FOR INFRINGEMENT OF ANY PROPRIETARY RIGHTS, RELATING TO USE OF INFORMATION IN THIS SPECIFICATION AND TO THE IMPLEMENTATION OF THIS SPECIFICATION, AND ONF DISCLAIMS ALL LIABILITY FOR COST OF PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES, LOST PROFITS, LOSS OF USE, LOSS OF DATA OR ANY INCIDENTAL, CONSEQUENTIAL, DIRECT, INDIRECT, OR SPECIAL DAMAGES, WHETHER UNDER CONTRACT, TORT, WARRANTY OR OTHERWISE, ARISING IN ANY WAY OUT OF USE OR RELIANCE UPON THIS SPECIFICATION OR ANY INFORMATION HEREIN.

No license is granted herein, express or implied, by estoppel or otherwise, to any intellectual property rights of the Open Networking Foundation, any ONF member or any affiliate of any ONF member.

A license is hereby granted by ONF to copy and reproduce this specification for internal use only. Contact the ONF at http://www.opennetworking.org for information on specification licensing through membership agreements.

WITHOUT LIMITING THE DISCLAIMER ABOVE, THIS SPECIFICATION OF ONF IS SUBJECT TO THE ROYALTY FREE, REASONABLE AND NONDISCRIMINATORY ("RANDZ") LICENSING COMMITMENTS OF THE MEMBERS OF ONF PURSUANT TO THE ONF INTELLECTUAL PROPERTY RIGHTS POLICY. ONF DOES NOT WARRANT THAT ALL NECESSARY CLAIMS OF PATENT WHICH MAY BE IMPLICATED BY THE IMPLEMENTATION OF THIS SPECIFICATION ARE OWNED OR LICENSABLE BY ONF'S MEMBERS AND THEREFORE SUBJECT TO THE RANDZ COMMITMENT OF THE MEMBERS. A COPY OF THE ONF INTELLECTUAL PROPERTY RIGHTS POLICY CAN BE FOUND AT: https://www.opennetworking.org/organizational-documents/

SDN Enabled Broadband Access (SEBA)

Reference Design v1.0

Document Revision Date: March 18, 2019

Document Revision Number: v1.0

Document Release Status: Final Specification

This reference design specification was authored by an operator-led Reference Design Team (RDT) composed of experts from:

Operator Group:

AT&T, Deutsche Telekom, Google Access, NTT, Turk Telecom/Netsia

ONF Supply-Chain Partners:

Adtran, Ciena, DellEMC, Edgecore, Juniper and Radisys

Write to rdspec@opennetworking.org with comments or questions.

Contributors

AT&T: Eddy Barker (RDT Chair), Tom Anschutz, Sumithra Bhojan, Michael Gasser, Tom Moore, Shawn Ying

Deutsche Telekom: Hans-Joerg Kolbe, Manuel Paul, Lothar Hoepken, Thomas Haag

Google Access: Cedric Lam

NTT: Keita Nishimoto, Tomoya Hatano, Kota Asaka, Jun-ichi Kani

Turk Telekom/Netsia: Bora Eliacik, Serhat Özdeveci, Cemil Soylu

Adtran: Andy Ruble

Ciena: Sergio Slobodrian

DellEMC: Tim Epkes

Edgecore: Dinesh Belwalkar, Jason Huang

Juniper: Marcel Wiget

Radisys: Shaun Missett

ONF Liaison: Aseem Parikh

Document Revision History

Date	Revision	Description
9/12/2018	Draft 0.1	Initial Draft, Partner Confidential
9/21/2018	Draft 0.2	Working Draft, Partner Confidential
10/1/2018	Draft 0.3	Working Draft, Partner Confidential
10/18/2018	Draft 0.4	Draft for Partner Vote to proceed to TLT for Member Review. Partner Confidential.
10/19/2018	Draft 0.5	Working Draft for TLT Review feedback Partner Confidential

10/31/2018	Draft 0.6	Draft for TLT Approval. Partner Confidential.
2/6/2019	Draft 0.7	Draft for SEBA RDT review, following General Membership review, ending on February 1, 2019. Header page - Add Turk Telekom / Netsia Implementation Stream contact for SEBARD-6 Section 2.2 Figure 2 updated for SEBARD-22 (Towards shareable FTTx infrastructure) Section 2.2.2 Tenant Approach – updated for SEBARD-22 (Towards shareable FTTx infrastructure) Section 2.3.2.1 for SEBARD-7 (Decomposing BNG) and SEBARD-21 (Design Guideline for Decomposed BNG) Section 2.3.2.4 for SEBARD-15 (Shared Redundant Database) Section 3.1.8.1 for SEBARD-18 (Abstract OLT) Section 3.1.12 Telemetry – expanded to address Telemetry, Monitoring and Logging, Analytics and Policy Functions for SEBARD-14 and SEBARD-19 (Telemetry)
3/18/2019	1.0	Final Specification for public release

TABLE OF CONTENTS

TABLE OF FIGURES

1 INTRODUCTION

This Open Networking Foundation (ONF) Reference Design describes the SDN-Enabled Broadband Access (SEBA) exemplar platform. SEBA is intended to support network and feature needs of multiple operators with a common architecture. The SEBA Reference Design provides a high-level template or architecture for supporting broadband access with a minimal prescription of technology choices.

The approach allows for multiple implementation streams to meet the SEBA requirements in whole or in parts as a set of modules and compositions that allow for a mix of SDN, NFV and also legacy PNF components to be used as compositional elements in a deployment. In addition to the SEBA Reference Design, ONF will potentially develop exemplar implementations and implementation streams which derive from the exemplar platform.

1.1 PURPOSE & SCOPE

SEBA is created to provide an architecture pattern for developing solutions for carrier broadband access. The purpose is to define a common infrastructure component that would be considered non-differentiating both for operators as well as suppliers. The commonality helps create efficiency in the development of open source and white-boxes, and then commercial products and support for those entities. To drive toward this purpose, the operator group involved in developing this Reference Design is expected to make use of the resulting implementation stream work product in some form or fashion beyond just lab or field trials.

The scope of the SEBA Reference Design is intended to cover a broad set of wireline and fixed wireless access technologies and related Service Edge capabilities. These include, but are not limited to: PON, XGS-PON, NG-PON2, EPON, future PON technologies, Gfast, Ethernet, fixed wireless, DOCSIS and xDSL. The scope should allow easily adapting new technologies, new silicon supporting these technologies and new devices that incorporate these technologies and silicon into deployable elements. This should be possible without re-writing major sections of the subcomponents that make up SEBA and should not require new fundamental interactions northbound to carrier automation platforms. The Reference Design supports the POD approach which means that all necessary components for service access delivery are covered. This includes beside plain access nodes also leaf spine architecture. This enables Aggregation and Service Edge functions which are supported by a switching fabric. The Reference Design supports requirements of operators who want to deliver broadband services without Service Edge capabilities as well as operators delivering IP services including service edge capabilities within the POD

Direct support of wireless mobility access, like that defined by 3GPP, is not in scope of this Reference Design, however many in the operator group have voiced the desire to have a common infrastructure layer and common components support both, so a keen eye is given to coordinating with a future wireless ONF project(s).

1.2 ASSUMPTIONS, DEPENDENCIES, PROCESS VARIANCES, OUT-OF-SCOPE SUMMARY

This document assumes the following relationships among Reference Design, target, and time to market solutions.

The Reference Design is described in terms of one or more solution elements. The particular implementation stream document will define how to deliver the elements and how to produce control, data and management plane. The most important of these steps is the target or end state. Target describes the most desired Reference Design in the context of the preferred future environment. The Reference Design may also define a series of well-known intermediate elements that might be needed to go to market sooner or with near term constraints that prevent moving directly to the target. These are defined as Time-to-Market solutions.

Given these relationships, the target Reference Design describes both a set of assumptions and dependencies in addition to the body of the design itself. Because the design may change over time, it is also expected that the assumptions and dependencies may also change with the Time-to-Market solutions.

The assumptions for the SEBA Reference Design are:

1. Operators are interested in initial commercial deployments in 2018
2. The existing carrier automation platforms include both legacy OSS as well as new orchestration systems, e.g. ONAP.
3. There are no strictly-greenfield deployments envisioned. This means that SEBA will need to be able to work within larger, already existent networks, services, and operational models.
4. There are requirements for multiple options to support Broadband Network Gateway (BNG) functions and these include functional aggregation in an exterior legacy device (PNF) and disaggregated placement of functions within the SEBA POD - either as Software of SDN VNFs.
5. The Reference Design will reduce operational complexity by hiding it in layered abstractions, as is typical of IT systems. This also means that the design will incorporate self-sufficiency and automation for its assembly, failure recovery, and performance.
6. SEBA will be deployed in infrastructure aggregates, often called PODs. A point of delivery, or POD, is a "module of network, compute, storage, and application components that work together to deliver networking services. The POD is a repeatable design pattern, and its components maximize the modularity, scalability, and manageability of data centers" (reference: Cisco Nexus 2000 Series Fabric Extenders Data Sheet).
7. SEBA will be constructed using containers run in Kubernetes as the cloud underlayer. It may use existing projects like Akraino or Airship to provide these functions and is expected to be loosely coupled to such layers.

It is recommended that implementations of the SEBA Reference Design are built using:

1. Working Kubernetes Environment
2. CI/CD Tools (e.g. Jenkins, etc.) for development as well as deployment instances

SEBA is related to several mature projects at ONF, including ONOS, VOLTHA, XOS, and R-CORD. Because this work interacts with existing released work in active communities at ONF, it is likely that some of the processes defined for normal new Reference Design work may need to be adjusted to ensure that the existing communities and their work do not become disenfranchised.

1.3 SEBA AND EXISTING MANAGEMENT AND ORCHESTRATION PLATFORM

Most carriers will need to develop adapters or agents that allow interworking between SEBA and existing management and orchestration platforms. To the extent that this work affects requirements and common aspects within SEBA, such work will be adopted in a common Network Edge Mediation module (NEM described in more detail below). However, aspects that are unique to a single carrier, product or deployment are considered out of scope for this Reference Design.

1.4 AUDIENCE

The ONF partners are the current audience of this pre-Alpha version of the document, per the ONF Reference Design Process, and at this stage this document should not be shared outside of the ONF Partners defined at the top of the ONF membership page.

Upon reaching the criteria for an Alpha stage RD, the ONF TLT at its discretion will send drafts to the full ONF membership list.

Following the ONF RD process for the timeframe for members to review and comment, and following review of comments by the TLT, the TLT will provide decisions about revision of the document and when to release the RD as a Final Specification.

1.5 DOCUMENT RELATIONSHIP

The SEBA RD is a standalone document in the SEBA process.

The ONF site provides a SEBA wiki that provides references to the designs, code, workflows, JIRA board, meeting times, meeting recordings, developer meeting list and Slack channel. Solution development represented in the artifacts at the SEBA wiki should be considered as a more detailed snapshot of implementation(s) for SEBA.

1.6 SOFTWARE RELEASES

The SEBA project should define software releases as a solution set for the software components, including but not necessarily limited to Network Edge Mediator/Edge Cloud Orchestrator, SDN controller, Access Node driver, and Aggregation and Service Gateway driver.

The SEBA software release documentation should provide the solution set information for these software releases.

The SEBA software release documentation should also provide the lifecycle management of the compatible releases between these components, in order to define flexibility and dependencies for coordinated upgrades of the components.

The hardware from vendors may also include embedded software for controlling, monitoring and abstracting low level functions of the hardware, including BIOS, firmware, board support drivers and board management controllers (BMCs). The vendors shall identify the required versions of these embedded software components, and how to upgrade these embedded software components using open software lifecycle management procedures.

1.7 HARDWARE RELEASES

The carriers define the hardware solution, including vendors, models, and releases. ONF suppliers do provide value to identify hardware for an ONF Reference Design, and to update the carriers with roadmaps and new product information for enhancements and improved cost.

2 REFERENCE DESIGN TARGET (24-36 MONTHS)

SEBA is designed as a set of container elements run in a Kubernetes environment. The system is modularized per typical microservice system architectures, and there is a hierarchy of modularity used to allow flexible compositions at different scales.

As is shown in Figure 1, SEBA is comprised of a few high-level software modules, including:

- Network Edge Mediator (NEM)
- SDN Control
- Control Applications
- Access Node (AN) Driver
- Aggregation and Service Gateway (ASG) Driver

Figure 1 also shows some of the typical hardware equipment that comprises SEBA, including:

- Access Nodes: PON-OLT, PON-ONT, DPU, other
- Aggregation and Service Gateway: one or several switches/routers in a setup that supports options for layer 2 aggregation, layer 3 service aggregation, Service Edge/BNG (and/or S/P-GW) functions and supports composing an SDN-controlled leaf-spine fabric. The fabric design can be VxLAN based as well as MPLS based. It provides a mesh for the localized ANs to access the network, and optionally a management network between the compute functions, the ANs, AN drivers and ASG drivers.
- Compute: servers that host the AN driver and ASG driver, as well as control and management plane modules

Finally, Figure 1 below also shows various interfaces among the physical and software entities and also the Technology Profiles, labeled TP, that provide abstraction hints for controlling technologies that are not similar to Ethernet.

The RD expects downstream implementation stream documents to develop and maintain specific implementation details, both from choices of components and also from instances or releases of those components.

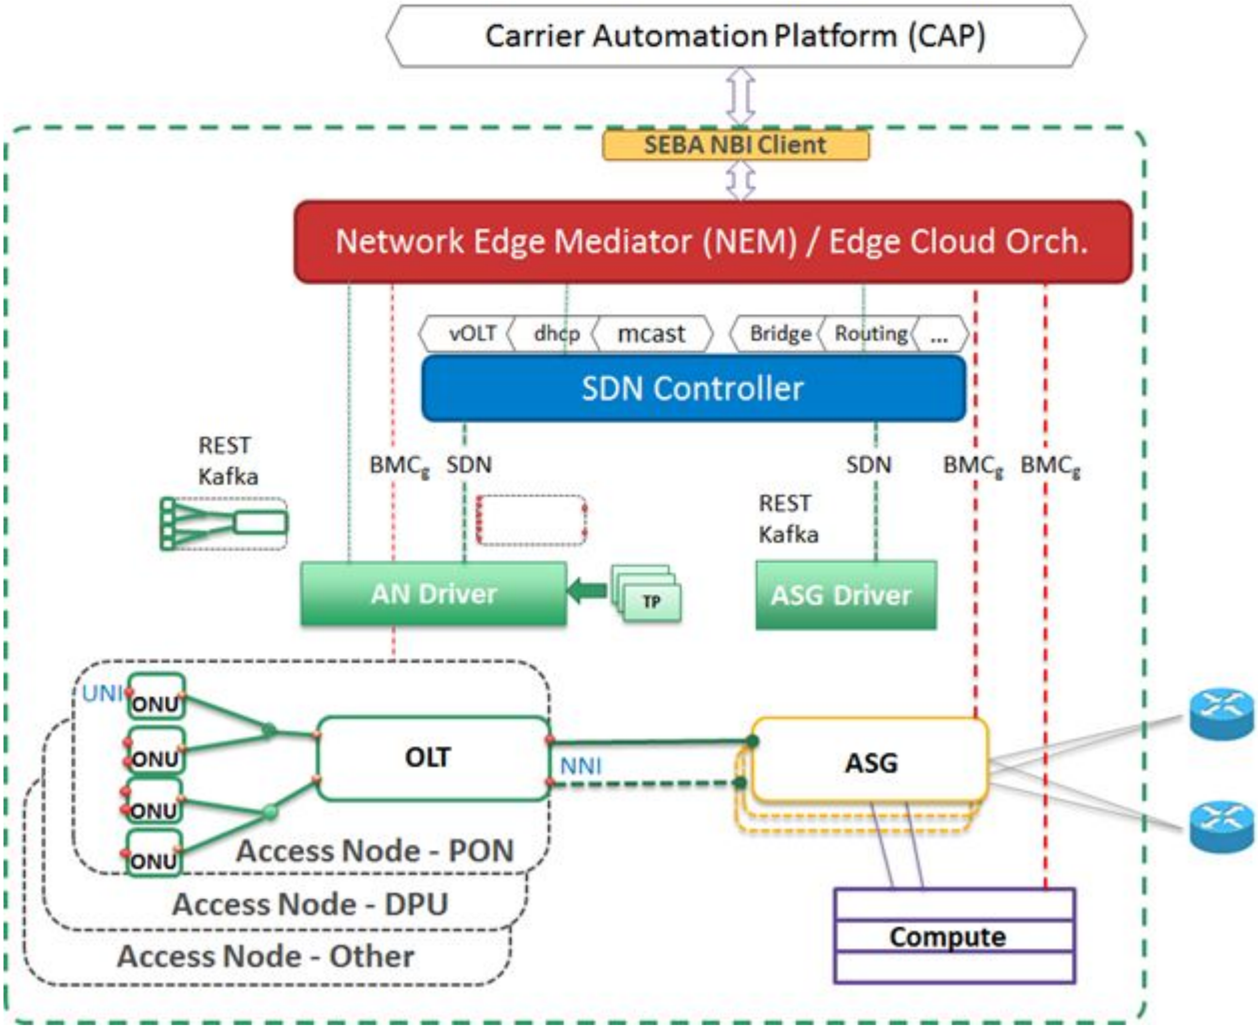

Figure 1 High Level Target Architecture

The target architecture diagram requires some principles and definitions to be noted:

- The Infrastructure Layer is not denoted but it includes the physical components - the Access Nodes, the Aggregation switches, and Compute.
- The Service Layer defines the binding of the components in the Infrastructure layer to deliver a service.
- The SDN controller maintains autonomy of the control structure to each component of the Infrastructure Layer involved in a service.
- ASG is only a functional block that supports aggregation, switching and routing of data plane, control plane and management plane traffic within a POD, and supports Service Edge capabilities. Use of multiple ASG devices, as depicted in figure 1, is a deployment option an operator may

select, but the ASG setup may also be non-redundant.

- The Board Management Controller (BMC) is a functional definition and industry term for an interface to equipment management functions.

2.1 SALIENT CHARACTERISTICS OF THE END STATE

Over the past several years, the SDN and NFV ecosystem has moved beyond skepticism and doubts into actionable strategies being adopted and deployed by operators around the world. For L4+ services (e.g., IMS, DNS, etc) and L1-3 (e.g., Optical and IP/MPLS, etc.) services, approaches, while not all optimal, are all fairly well understood and achieve many of the key benefits of decoupling, common off the shelf systems, and disaggregation of SW. While the industry is not at an ideal state, it is well on its way and headed in a common direction, thanks in a large part to organizations like ONF who were the early trailblazers. Moving forward, operators' access networks and service delivery platforms stand to benefit from such a focus. Access plant remains highly proprietary and capitally intense, representing a large component of capital outlay of most large operators. Access networks also come with unique constraints, such as environmental, regulatory, space, and power that tend to favor small edge compute platforms and highly distributed open-spec peripherals.

In a sense, the "easy work" has been done and the largest operational and capital outlays for operators remain fertile ground for transformation. The SEBA effort will drive access networks across the globe to deliver on the promise of open, software-driven systems in new key areas such as multi-gigabit fiber access networks and will ultimately look to reuse the concepts presented here in emergent wireless access.

Each of the key characteristics below represent key industry drivers for highly performant, secure, flexible common solutions that aim to represent a lowest TCO infrastructure for operators across the globe.

1. Automated (Zero-Touch), secure, reliable
2. Affordable and transparent – Startup cost & operations
3. Highly Modular HW and SW, including peripherals and acceleration.

4. Architecture of HW and SW that enables small start and scalability
5. Open Systems First, Disaggregation by Design, Modularity by Design, Loose Coupling.
6. May be sourced using a variety of business and packaging models – Operational Abstractions is a Priority.
7. Performant. Real-time, low latency
8. Integrally considers and intersects with automation and management systems and approaches
9. Multi-access by design
10. Multi-Cloud/Hybrid Cloud support
11. Is not defined by physical location (e.g., CO)

2.2 TARGET REALIZATION APPROACHES

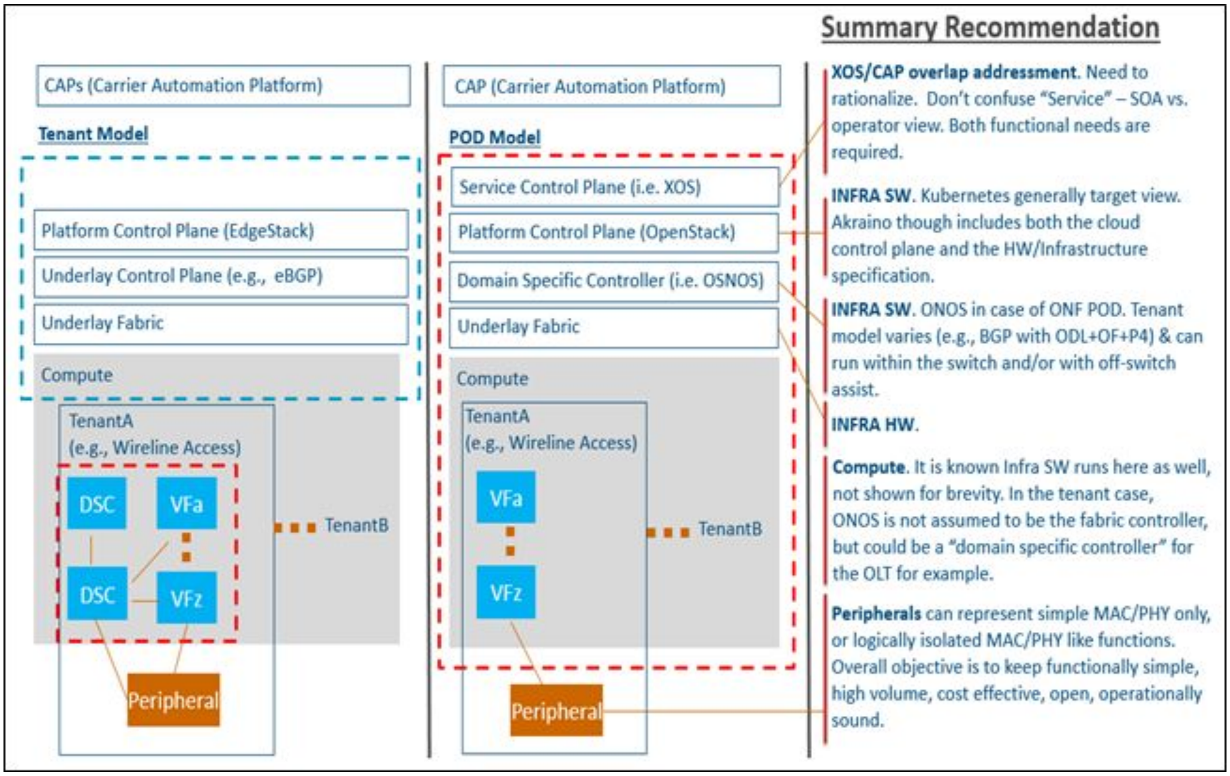
CAPs (Carrier Automation Platform)
Tenant Model
Platform Control Plane (EdgeStack)
Underlay Control Plane (e.g., eBGP)
Underlay Fabric
Compute
TenantA
(e.g., Wireline Access)
DSC
VFa
DSC
VFz
TenantB
Peripheral
CAP (Carrier Automation Platform)
POD Model
Service Control Plane (i.e. XOS)
Platform Control Plane (OpenStack)
Domain Specific Controller (i.e. OSNOS)
Underlay Fabric
Compute
TenantA
(e.g., Wireline Access)
VFa
VFz
TenantB
Peripheral
Summary Recommendation

XOS/CAP overlap addressment. Need to rationalize. Don't confuse "Service" – SOA vs. operator view. Both functional needs are required.

INFRA SW. Kubernetes generally target view. Akraino though includes both the cloud control plane and the HW/Infrastructure specification.

INFRA SW. ONOS in case of ONF POD. Tenant model varies (e.g., BGP with ODL+OF+P4) & can run within the switch and/or with off-switch assist.

INFRA HW.

Compute. It is known Infra SW runs here as well, not shown for brevity. In the tenant case, ONOS is not assumed to be the fabric controller, but could be a "domain specific controller" for the OLT for example.

Peripherals can represent simple MAC/PHY only, or logically isolated MAC/PHY like functions. Overall objective is to keep functionally simple, high volume, cost effective, open, operationally sound.

Figure 2 Target Realization Approaches

SEBA may be realized in a variety of different ways, depending on the operator and/or situation. Most important in this regard is the clear decoupling and separation of "service" from "infrastructure". This also provides for the realization of SEBA on a variety of different infrastructure platforms, including, but not limited to those that are CORD based. Modularity through good functional decomposition should provide for extensive reuse across the ONF solutions providing for high volume, consistent, SW assemblies possible from the ecosystem.

2.2.1 POD Approach(es)

Generally, this would be inclusive of infrastructure and service layer and be operated as SDN based access platforms. Controller can be operated inside POD or as "cloud in a box" (e.g. based on cloud control plane functions embodied within the POD).

2.2.2 Tenant Approach

This would be the case where an operator already has a cloud environment specified, and SEBA is either a tenant set of workloads and peripherals or is itself hosting a tenancy environment, with the peripherals managed and operated primarily at the service layer.

Both POD approach and Tenant approach do not limit how the broadband access services are deployed. SEBA supports the Access network infrastructure provider to deliver the broadband access services, or the infrastructure provider to deliver the network to third party broadband access service providers under a wholesale agreement.

2.3 MAJOR FUNCTIONAL COMPONENTS

2.3.1 SEBA Infrastructure

2.3.1.1 Hardware

Hardware includes physical components - the ANs (including OLTs, ONUs), switches, compute servers, physical interface plugins, fibers, cabling and powering.

2.3.1.2 Software

The SEBA project delivers a set of software components specified in the high-level architecture, including but not necessarily limited to NEM/Edge Cloud Orchestrator, SDN controller, Access Node (AN) driver, Aggregation & Service Edge (ASG) driver.

Note that Service Providers specify the SEBA NBI client software for the interface to their Carrier Automation Platform (CAP).

The hardware from vendors may also include embedded software for controlling, monitoring and abstracting low level functions of the hardware, including BIOS, firmware, board support drivers and board management controllers (BMCs).

The infrastructure services for instantiating a POD and remotely maintaining the lifecycle of the POD requires open and automated management approaches.

2.3.2 SEBA Service Layer

2.3.2.1 BNG

The BNG (Broadband Network Gateway) is a key component in fixed broadband access networks. It resides at the demarcation point between the access network (usually based on L2 tunnels) and the routed IP/MPLS network. The BNG provides per-subscriber services and is the highest-tier network element that has the full per-subscriber context. Beyond this point, traffic can be correlated to a subscriber using the IP address, but the complete view is gone. The functional requirements on BNG types are described by the Broadband Forum (BBF, TR-178 and related documents). The core functions of a BNG at the access side and towards the core, but not limited to, are:

- Aggregation of L2 access tunnels / VLANs / MPLS PWs
- Termination of network attachment (PPPoE or IPoE tunnels), authentication, (dynamic) policy enforcement, AAA client
- Tunnel switching and termination (e.g., L2TP LAC)
- Traffic filtering and shaping

- Lawful interception
- Anti-Spoofing
- Split horizon rules
- Per-subscriber OAM (e.g. using keepalives)
- Accounting

The BNG also acts as a router as it is part of the IP core network of the service provider. This covers, amongst others, the following functions:

- Routing protocols (IGPs, EGPs)
- MPLS control and user planes

The BNG directly interfaces with Policy control systems and AAA services.

In SEBA, there are multiple ways to deal with the required BNG functionality:

a) Serve an external BNG. In this case, SEBA acts as smart aggregation network
b) Embed the BNG as part of the POD, either as PNF or VNF
c) Decompose the BNG into Service Edge (SE) and Router and deploy as PNF or VNF with the PNF possibly being embedded in the AN or ASG

While the first two options basically represent state-of-the art thinking and keep existing centralized functionality in closed systems, option c) goes beyond.

De-composing a BNG

Looking at a BNG, we can de-compose it to the following control plane and user plane building blocks as shown in Figure 3.

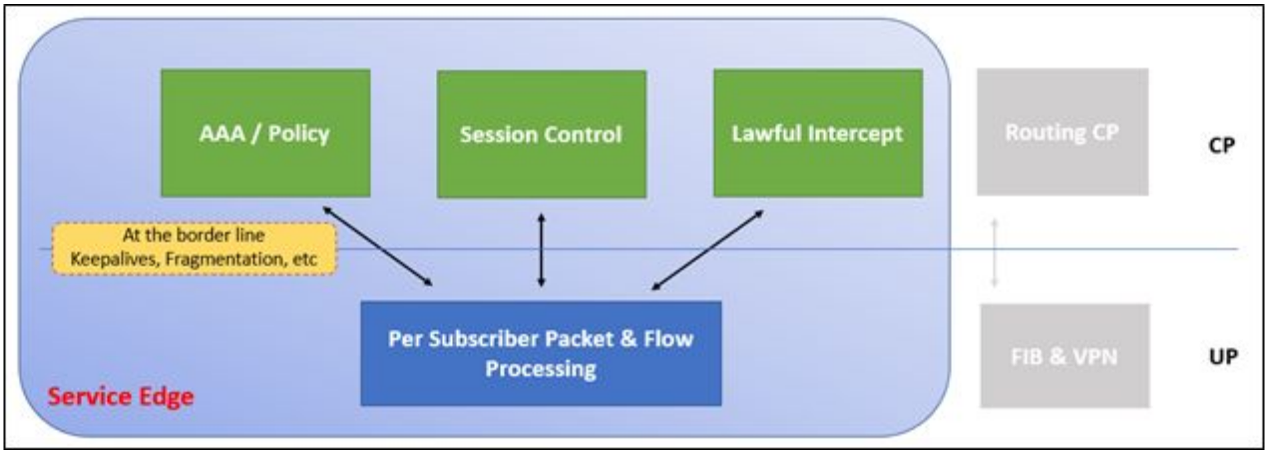

Figure 3: BNG Building Blocks

The Routing Control Plane (CP) as well as User Plane (UP) are not dependent on the per-subscriber functions and can thus become separated. By doing so leads to a look at a BNG as being split into a subscriber-facing functions which we call Service Edge (SE), and a router.

First step of disaggregation: Split BNG into router and service edge.

Further steps of disaggregation are possible. A natural next step is splitting control and user plane (similar to 3GPP CUPS model) and allowing e.g. to run CP on servers while UP stays on specialized hardware. This further enables the service providers to scale both layers independently of each other or centralize CP.

The user plane can even further be decomposed. This may be needed in deployment scenarios where multiple physical entities serve as distributed processing points.

De-composed BNG Deployment options

When going for a de-composed BNG, a natural choice is to embed the routing function (CP and UP/FIB) into the switching fabric which is constituted of Aggregation and Service Gateway (ASG) devices. The CP may run on a centralized SDN controller such in Trellis, it may also run as a distributed routing process on the ASG.

The SE may run:

- Embedded in ASG devices by making use of programmable silicon
- Embedded in AN devices by making use of programmable silicon
- Embedded but spread across ASG and AN devices by making use of programmable silicon
- On a "black box" mounted into a rack in the POD as PNF on dedicated hardware or as VNF on generic servers

The decision where to place functions strongly depends on the actual requirements of the service providers. SEBA shall cater for all the above, leaving sufficient flexibility for placement.

An important aspect is the case where components of the POD, e.g. the ANs are located in geographically further distributed locations.

All cases, including the one with a non-decomposed BNG directly attached to a POD have the need for a traffic steering mechanism in common. The SDN control function needs to steer the L2 tunnels of the customers to the service edge. Once such a steering mechanism is in place, the service provider can even steer customers to different SEs. Slices can be built. One can e.g. imagine dedicated SEs for enterprise or IoT customers. Those can even be implemented using different technologies (e.g. virtualized on x86 for IoT, on programmable switching silicon for enterprise customers).

2.3.2.2 Network Edge Mediator (NEM)

The Network Edge Mediator (NEM) serves as the mediation layer between the edge/access system and the service provider backend and global automation frameworks. NEM will provide the interfaces and components to support FCAPS functionalities required by the service provider for managing the access network components and broadband service subscribers the SEBA POD is designed to offer and support. A variety of operator OSS/BSS and global orchestration frameworks can be integrated northbound for specific deployment needs.

2.3.2.3 SEBA NBI Client

The SEBA NorthBound Interface (NBI) Client provides an application layer for management interfaces between the Carrier Automation Platform and NEM.

The SEBA NBI client is tightly coupled with the Carrier Automation Platform and so is specific to the Service Provider.

2.3.2.4 SDN Control

SDN comes with three major capabilities:

- A means to take control functions out of a dedicated box and centralize them and create applications for such purposes
- A means to dynamically program data paths through the network
- A means to directly program packet processing on a chipset

The first two play a key role when steering subscriber traffic. The last one enables programming user plane packet processing for a service edge onto programmable silicon.

When looking at a customer / CPE attachment process, there are two major stages

Stage 1: Device Attachment and Recognition

A device is powered on and attaches to the access node in the service provider domain (usually the AN). Layer 1 comes up and the AN needs to enable the L2/L3 connection set up. To do so, an AN can create an event like a 'Port Up' message that is processed by the SDN control framework. Based on implementation, a network path can be created to enable step 2, where the device attaches to the service edge / BNG. CORD does this using 802.1x port authentication.

Stage 2: Subscriber Session Establishment

Prerequisite of this step is the reachability of the BNG/SE. Depending on the service provider policy, this additional step is required here for CPE authentication and access protocol establishment (e.g. PPPoE termination and PPPoE/L2TP).

The actual SDN and BNG implementation may benefit from external state databases. For the traffic steering mechanism in stage 1, it is obvious that these created flows need to be stored in a central database inside or attached to the SDN controller. This may also include implicit authentication states such as e.g. via line IDs.

For the subscriber session state (stage 2), the options depend on the BNG option chosen. For external BNGs, storing the subscriber session state is out of scope for the SEBA POD.

In case the BNG resides in the SEBA POD, independently on whether it is embedded in an ASG element or a standalone PNF/VNF, session state may be kept internally to this instance or, in order to e.g. allow for fast failovers, be stored in a centrally accessible state database inside the POD.

2.3.2.5 Aggregation and Service Gateway (ASG)

Aggregation and Service Gateway (ASG) devices (switches) support Layer 2 or Layer 3 network aggregation, switching, and routing of data plane, control plane and management network connectivity within the POD as well as to external data networks, and supports Service Edge/BNG capabilities.

There may be one or more ASG devices, and setups as switching fabric, depending upon the implementation.

2.3.2.6 AN Driver

The AN Driver shall be a collection of loosely coupled services which provide an abstract interface from the SDN controller to target device hardware. Different AN drivers can support many technology types such as PON, XGS-PON, NG-PON2, Gfast, or DOCSIS.

A PON AN Driver shall be developed to support the XGS-PON technology. It is expected able to support similar technologies such as GPON, EPON, or NG-PON2. OLT/ONT hardware can be delivered in many forms. Vendor OLTs/ONTs, whitebox OLTs/ONTs, and pluggable OLTs are all supported.

Adapters provide an interface from the core AN Driver to the specific hardware implementation. The PON AN Driver hides PON level details (T-CONT, GEM ports, OMCI etc.) from the SDN controller, and abstracts each PON as a pseudo-Ethernet switch easily programmed by the SDN controller.

The AN Driver has the responsibility of establishing the data plane connections through the hardware by interpreting service requests from the

SDN controller and transforming them into requests to be fulfilled by the appropriate adapter.

The PON AN Driver has the responsibility of forwarding control plane requests to the SDN controller. Control plane requests corresponds to authentication protocols (802.1x, PPPoE, DHCP) and multicast service such as IGMP, and OMCI messaging.

The PON AN Driver provides to the SDN controller the ability to manage and control the ONU through OMCI messaging.

A DPU AN Driver provides to the SDN controller to manage and control the aggregation functions of the DPU and the Gfast access interface functions of the DPU.

2.3.2.7 ASG Driver

The ASG Driver provides the management and control functions for the ASG devices. Functions for user plane aggregation include create, delete, update and retrieve L2 or L3 connections between access ports and uplink ports, and to monitor these connections. Functions for management control plane connectivity include (as applicable) to create, delete, update and retrieve management and control paths between ASG devices and compute servers, between certain ANs and compute servers, and from ASG devices to external BNGs or routers.

Note that an ASG device does not necessarily attach to all types of ANs. ASG may attach for example to an OLT AN, and the OLT AN in turn may connect to an ONU AN.

2.3.2.8 Access Nodes (AN)

The Access Nodes will be a specific implementation the broadband access technology, such as PON technology. Vendors can produce AN boxes providing drivers to interface the required adapter. White box ANs based on industry standard chipsets are used. Drivers provide a bridge between hardware supplied SDKs (software development kits) and the required adapter.

AN devices provide the physical layer termination of the network access ports and the aggregation of the traffic to the ASG switch. The number of ports provided can vary based on the hardware implementation, such as 8, 16, and 24 access ports and beyond.

2.3.2.9 Profiles

A Technology Profile (TP) helps to define a subscriber service. It contains AN specific parameters specific to a technology such as GPON, XGS-PON, NG-PON2, EPON, future PON technologies, DOCSIS, Fixed Wireless, Ethernet, xDSL etc. Thus, the profile is specific to the technology.

A device adapter interprets the technology profile. Multiple technology profiles can be defined for a specified technology type. These profiles define the service level characteristics. A residential service could use a weighted 4 queue model while a business service could require a strict priority 8 queue model.

A speed profile will define the service parameters related to the bandwidth achievable by a subscriber. Depending on the type of service being offered different parameters may be defined. A residential service could define a minimum and maximum speeds. A business service could define minimum, maximum, and guaranteed speeds.

Different technologies may implement speed profiles differently. Some may use simple meter bands (XGS-PON) while others may manage the physical line (Gfast sync rates).

2.3.2.10 Access Technology - PON

SEBA is expected to support various kinds of PON-related technology (e.g. GPON, XGS-PON, EPON, Gfast etc..) and physical devices. In this architecture, these PON-specific features and devices (i.e. OLT/ONUs) are abstracted by AN Driver into a pseudo-Ethernet switch whose ports correspond to ONU-UNIs and OLT-NNIs. This abstraction provides operators with various options to deploy PONs that have different equipment structures while managing them in a common SDN architecture. For instance, it is possible to use the both types of OLT: Box-type OLT (refer to Figure 1, "High Level Target Architecture") and pluggable module-type OLT (refer to Figure 4). In the latter case, the control messages sent from AN Driver to OLT go through ASG, while the messages are directly sent to OLT in the former case.

It is also possible to run a TC (Time Critical) functions (e.g. DBA (Dynamic Bandwidth Allocation) function) apart from hardware (refer to Figure 5 (right) instead of running these functions inside the hardware (refer to Figure 5 (left)). In any case, the PON-specific features are managed under AN Driver.

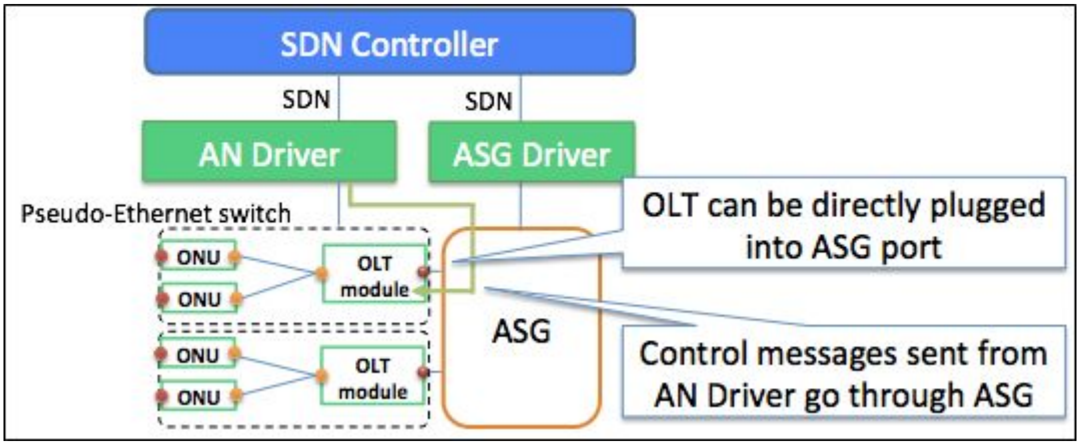

Figure 4: Architecture of pluggable module-type PON-OLT

Refer to Figure 5 for Box-type PON-OLT architecture

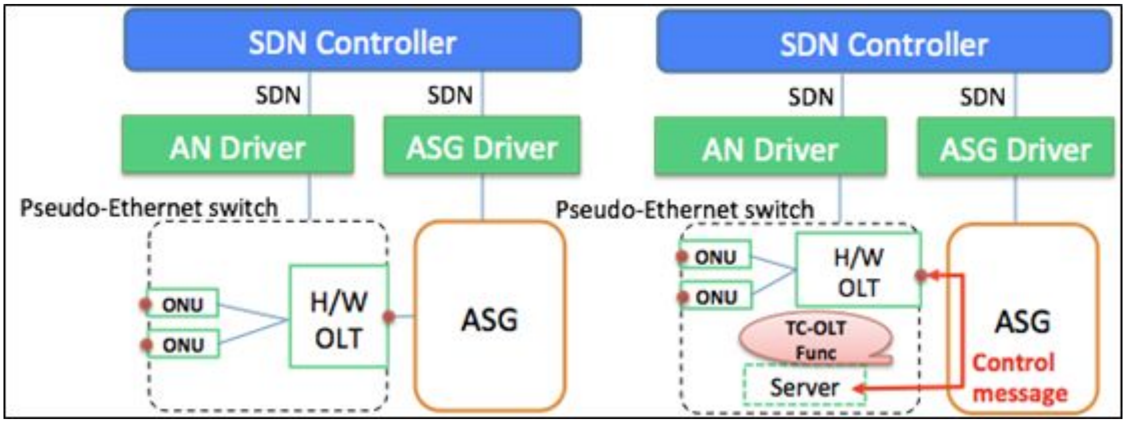

Figure 5: Options for PON-OLT architecture

(Left) Single H/W-type PON-OLT architecture.

(Right) Functional decomposition-type PON-OLT architecture.

2.3.2.11 Infrastructure Requirements

Scaling

The scalability requirement is deployment dependent. It depends heavily on the implementation network architecture and services offered from the POD. The SEBA community shall conduct the scalability study and provide the deployment guidelines which meet the service provider implementations.

SEBA POD must support scaling the POD elements horizontally to accommodate from the initial deployment to the larger footprint by adding the ANs and resources.

Vertical scaling of a POD means adding more resources to VNFs or into operating hardware. Horizontal scaling is preferred in clouds to simplify operations. The application of vertical scaling must be justified as an alternative to horizontal scaling.

Resource management

Resource management includes the definition of compute, memory, storage and network connectivity and bandwidth required for the initial POD installation, and the modeling and capacity management plan as additional utilization grows by ANs and subscribers.

Data Collection

The status and the health of the SEBA POD, including all physical and local components and services, must be monitored and provide on-demand and periodic reporting mechanism to the CAP through the SEBA NBI Client.

Operation configurable Threshold Crossing Alert (TCA) must be utilized for raising the event to inform the operator about the state of the POD elements.

Overlay networking requirements

SEBA POD must provide secure communication channels internally within the POD and externally to the provider's management network. The data path channel must not be compromised and provide unlawful access to the management and control channels into the POD infrastructure and provider's management network.

Speeds/feeds, performance (e.g., compute and I/O)

The physical connectivity between the POD to the access network and backbone must maintain the SLA of the services offered of the POD. The connectivity requirement may vary based on the overall access network architecture design and the type of the services offered. The SEBA exemplar implementation of a provider will determine the requirement. The requirement shall define whether it is Layer 2 and/or Layer 3 connections.

Defined by the implementation, the oversubscription over the physical bandwidth may be required.

The SEBA project shall provide the compute and storage performance based on the provider's implementation and scalability requirements.

2.3.3 POD Assembly

The assembly of a POD includes the definition of the virtualization approach of the software to the hardware resources.

The exemplar platform begins with a set of container elements run in a Kubernetes environment to optimize the utilization of compute resources.

Large-scale orchestration and lifecycle management of PODs from an operator's cloud environment may lead to approaches to standardize on other variants of infrastructure environments.

A VIM or Multi-VIM approach may enable SEBA to run on other infrastructure environments.

A Virtual Infrastructure Manager (VIM) controls and manages the NFV infrastructure (NFVI) compute, storage and network resources, usually within one operator's infrastructure domain. VIMs can direct a multidomain environment or optimize to a specific NFVI environment. A VIM coordinates the physical resources to deliver network services.

2.3.4 Use Cases and Flow

The workflow branch under the SEBA project wiki page hosts the SEBA use cases and workflows from multiple Service Providers and will become the basis for the SEBA implementation streams.

2.4 COMPLIANCE WITH END STATE

This section defines guidelines that the Reference Design (RD) should follow to achieve the desired end state. Include Key Performance Indicators (KPIs) to achieve the technical and business goals.

Guidelines for the end state of the Reference Design include:

- Suppliers can use this RD to build the solution
- Providers can verify compliance by the suppliers to the RD
- Reliability above a defined number of “9s” for availability of the solution

3 TIME TO MARKET SOLUTIONS

There will be multiple options, some trivial and some substantial. The motivation is for an operator to realize a substitutional model for hardware and software selection to enhance features and reduce costs. There may also be variants amongst the operators in the operator group. It is highly desired that operator variants be minimized to the functional components rather than the interfaces between the components.

3.1 SOLUTION ELEMENTS

3.1.1 Major Functional Elements

The following sections define the functional building blocks.

3.1.1.1 Carrier Automation Platform (CAP)

The Carrier Automation Platform that is external and northbound of the SEBA is a robust design framework that allows specification of the service in all aspects – modeling the resources and relationships that make up the service, specifying the policy rules that guide the service behavior, and the applications, analytics and closed-loop events needed for the elastic management of the service. ONAP is an example of a CAP.

The orchestration and control framework (Service Orchestrator and Controllers) is recipe/policy-driven to provide automated instantiation of the service when needed and managing service demands in an elastic manner.

The analytic framework closely monitors the service behavior during the service lifecycle based on the specified design, analytics and policies to enable response as required from the control framework, to deal with situations ranging from those that require healing to those that require scaling of the resources to elastically adjust to demand variations.

3.1.1.2 Infrastructure Layer

The Infrastructure layer includes the hardware in the solution, including the devices, racks and shelves, powering equipment and connections, and external fibers or electrical cables, and other passive devices.

The devices in the POD include the ASG devices, compute servers, OLTs, or other devices defined by the technology (e.g., other types of devices for Wireless or DOCSIS solutions that become defined for SEBA).

The external fibers are part of the infrastructure as they connect to ports of the devices and carry user plane and/or management traffic, and the devices monitor performance of signals and protocols carried by the fiber media.

The passive devices include PON splitters, patch panels, and other equipment that do not monitor the signals, but which perform essential functions to connect paths, combine signals, split signals or filter signals.

The combination of components in the Infrastructure layer compose the physical inventory of the solution that suppliers plan for delivery, fulfillment solution providers aggregate in a supply chain, and that installers place and validate.

3.1.1.3 Physical Topology

The physical topology presents more detail about the detailed organization and connectivity of the infrastructure layer components into a complete solution.

3.1.1.4 Service Layer

The Service Layer defines the attributes of the service, and the configuration and binding of the components in the Infrastructure layer to deliver a service.

3.1.1.5 Application Layer

The application layer is the Open Systems Interconnection (OSI) layer closest to the user. This layer establishes communications between applications, and to the user. An example SEBA component in the Application Layer is the SEBA NBI client.

3.1.2 Interfaces and Interior APIs

The functional descriptions of the internal interfaces include:

1. API between CAP (Carrier Automation Platform) and the SEBA NBI client
2. API between SEBA NBI Client and NEM
3. API between NEM and SDN Control
4. API between NEM and AN Driver
5. API between NEM and AN (BMC API for system management)
6. API between NEM and ASG (BMC API for system management)
7. API between NEM and Compute (BMC API for system management)
8. API between Edge Cloud Orch and Compute (orchestration and life cycle management)
9. API between SDN Control and AN Driver
10. API between SDN Control and ASG Driver
11. API between AN Driver and AN
12. API between ASG Driver and ASG
13. Interface between DPU and an ONU of the PON AN
14. NNI interface between AN and ASG
15. NNI interface from ASG to external BNG/router

3.1.3 Security

It is proposed to extend and unify the proposed Epics and user stories defined for VOLTHA to SEBA, and to address security requirements for VNFs across the SEBA POD and not just VOLTHA[1]. When referring to "SEBA" only below, the context should be interpreted as for "all of SEBA, that also includes but is not limited only to VOLTHA orchestration".

As a managed virtualized environment, SEBA should **derive** and **leverage** security requirements, security architecture, security best practices and open source security solutions from other open communities.

- ONF is a member of Linux Foundation
- In the ONAP project in Linux Foundation, the ONAP requirements for Security, and ONAP architecture and work in progress do provide much of these elements.

Whether or not a service provider deploys SEBA under ONAP, the maturity of the ONAP security requirements, and ONAP security solutions being developed in open source provide a good foundation for security requirements and user stories for SEBA.

- All VNFs within the VOLTHA *(and now also SEBA)* architecture shall comply with **VNF General Security** requirements, as identified and aligned to the current ONAP release (Epic VOL-818)
- All VNFs within the VOLTHA *(and now also SEBA)* architecture shall comply with **VNF Identity and Access Management (IDAM)** requirements, as identified in the current ONAP release (Epic VOL-819)
- All VNFs within the VOLTHA *(and now also SEBA)* architecture shall comply with **VNF API Security** requirements, as identified in the current ONAP release (Epic VOL-820)
- All VNFs within the VOLTHA architecture shall comply with **VNF Security Analytics** requirements, as identified in the current ONAP release (Epic VOL-821)
- All VNFs within the VOLTHA architecture shall comply with **VNF Data Protection** requirements, as identified in the current ONAP release (Epic VOL-822)

[1] Note that the following was originally presented in the VOLTHA community, and VOLTHA epics and user stories now exist based upon this direction.

- VOLTHA transactions that change any configurations or perform any operator actions from any operator or automated interface (CLI, Northbound API, control loop action) should use a Transaction ID (as in ONAP) as a best practice to be able to perform effective security Audit Logging & other Logging.

Some notes for discussion about a "Time to Market" implementation for Security –

- The ONAP requirements do refer to Network Cloud Service Provider (NCSP) and "Network Cloud" functions, such as the NCSP's IDAM API.
- Not all ONAP requirements will apply for SEBA use cases. Some user stories may be analyzed and closed without any required implementation, and noted with the reasons for being "not applicable".
- SEBA should not implement the entire security solution for an NCSP. One NCSP's implementation will vary from another NCSP's.
- The SEBA exemplar platform should incorporate security protocols into the architecture and interfaces to support the core security requirements for a reference architecture only. The interpretation of where a reference architecture ends can be determined in the analysis of the user stories based upon ONAP VNF requirements.
- In general, user stories should define a delineation or interface as needed to the NCSP implementation from the core reference implementation of SEBA

3.1.4 Reliability and Resiliency

Reliability is the definition of the probability of a system or component to function under stated conditions for a period of time. The reliability is also defined in terms of availability of a system or component in terms of number of "9s", such as "five 9s" indicating 99.999% availability.

A reliability analysis is recommended for the components of a POD, and of the entire POD to predict the availability of the POD as a system.

Resiliency is the ability of a server, network component, or POD to recover from a failure (such as a power failure, or equipment failure) and quickly resume operations.

Redundancy or clustering is employed to help improve both reliability of a POD, and resiliency of operations in a POD.

3.1.5 System Performance

System performance measurements in the internal processing functions of the POD are important to measure and understand with respect to system response for control and management transactions, and for scalability of the system to provide a determinate number of operations of a certain type, while the system is operating under a defined load profile (profile of a variety of defined operations at defined frequencies over a defined period of time).

The implementation of monitoring tools to monitor system performance must be considered and selected to minimize their own impact on system performance and resources.

3.1.6 Capacity Management

Capacity management provides analytics and reporting of the resources needed for a solution. It derives capacity measurements from the Performance Management element and from the resources that define the capacity of a solution element.

The implementation of capacity management may occur in the Carrier Automation Platform (CAP) that receives performance measurements from the POD(s), and so the Performance Management collection requires forwarding to the CAP for that function.

A provider may determine that Capacity Management functions are implemented in the local SEBA POD and provided through a local operator interface.

3.1.7 Fault Management

Fault Management applies at a SEBA POD level. A Carrier Automation Platform (CAP) that attaches to the SEBA POD through the SEBA NBI client typically provides two functions at an enterprise level for fault management: (1) Show Current Active Alarms, and (2) Show Alarm & Event History.

For purposes of this discussion, an Alarm can also be a "standing condition" that has a set/cleared state but is mapped to a "not alarmed" severity.

It is desired for NEM to provide a normalized collector implemented in Kafka for all faults for external northbound OSSs/Orchestrators to be able to receive streams of fault history in a deterministic manner. As a streaming platform, Kafka provides the capabilities to publish and subscribe to streams of records, to store streams of records in a fault-tolerant durable way, and to process streams of records as they occur.

While VOLTHA publishes faults to its Kafka bus that NEM can provide to its northbound clients, other faults may be derived from other APIs in a SEBA POD such as Redfish for device management. NEM can develop a "transformer" as necessary to export data from other APIs to a Kafka topic, and thus provide all faults through a normalized collector in Kafka.

Northbound OSSs/Orchestrators can implement and install an agent as a microservice in the SEBA NBI client that subscribes to the normalized Kafka collector and transforms the faults to a desired format (e.g., VES for ONAP, IPFIX for another OSS, etc.) for transport to the Northbound CAP. The SEBA NBI client for a CAP is not part of SEBA, as different CAP implementations may use different management protocols.

As a SEBA NBI client can discover or re-attach to a SEBA POD at any time, it may be able to depend upon cached alarms and events in the SEBA POD and to re-sync with this incremental history to correctly update its "Show Current Active Alarms" function, if and only if the SEBA NBI client can resync from the actual last alarm or event from a prior attachment to the SEBA POD.

The Kafka bus implementation in SEBA should allow the SEBA NBI client to determine if it is able to sync to the last known fault management state of the SEBA POD. If the SEBA NBI client is not able to sync to the last known state, then it will have missing alarms and events. In that event, then the SEBA POD should provide a "Retrieve all Active Alarms from SEBA POD" function.

3.1.8 Configuration

SEBA shall provide abstract configuration interfaces to provide subscriber level services. SEBA can be used in multiple deployment environments. Each implementation should provide simple easy to use APIs providing the minimum required parameters.

3.1.8.1 Abstract OLT

When using an external BNG a model is required to relate the various physical components of a POD into a legacy type OLT. This model is given a name and a structure into which the individual pieces of the POD are mapped. The "Abstract OLT" can be viewed in the figure below.

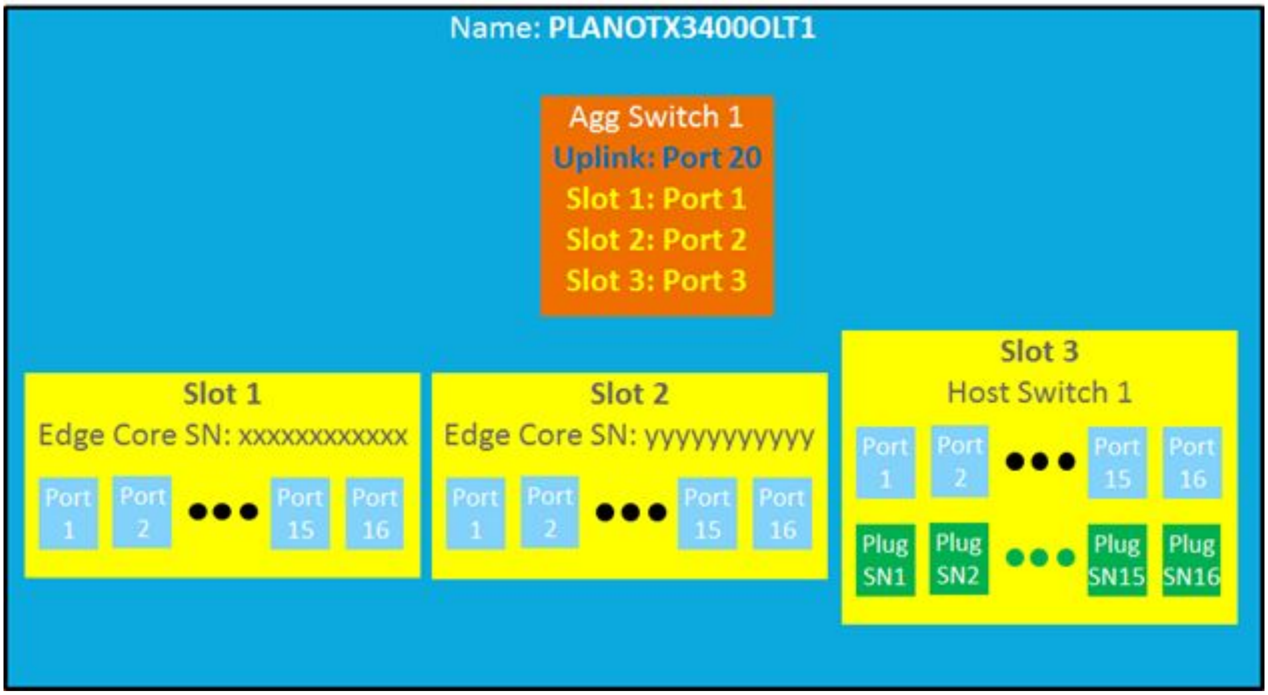

Figure 6: Abstract OLT Example

Individual ANs are mapped into slots of the abstract OLT. Each slot supports 16 XGS-PON ports and each XGS-PON can support 64 ONTs. Each ONT will be given a concrete number directly related to subscriber service provisioning. The ONT number must then be associated with a ONT using the required verification data. This data could be serial number, MAC address, registration ID, or other unique information.

Note: For small deployments, the host switch for micro-plugs may be collapsed with the AGG switch.

3.1.8.2 Profile configuration

SEBA shall provide the ability to manage profiles including technology profiles and speed profiles and associate these profiles to service types defined by the operator.

3.1.8.3 Service configuration

Service configuration will identify the physical location as defined in the Abstract OLT definition, service type, speed profiles, and VLAN tags. Service types are defined by the operator (RESIDENTIAL or BUSINESS) as their service models require. Likewise, speed profiles are defined by required service tiers (100M, 500M, 1G). VLAN tags define the model expected on the AG switch uplink ports interfacing to the external BNG.

3.1.8.4 Backup and Recovery

The NEM must provide the capability to periodically collect configuration information from each of the POD components and export them to another safe location. Then, in the event of a software or equipment failure it must be possible to restore the SEBA POD and the customer's service using the backed-up system configuration information.

SEBA shall provide an API to backup the configuration information to an external system.

3.1.8.5 Restore from Backups

A service-affecting event may occur in which a part of or all of the SEBA POD components have been corrupted and rendered inoperable. In this scenario the recovery plan would be to restore the POD components using recent backup files.

1. SEBA shall provide an API to restore recent set of backup configurations.
2. The restore process shall be able to be monitored for progress and success.

3.1.8.6 Software Lifecycle Management

SEBA shall provide APIs capable of managing software for the components of the POD. This includes the physical equipment resident in the POD as well as all the ONT devices connected to the XGS-PON ports.

The component of the POD shall be managed independently. Where new services require software upgrades to various components of the POD those services will not configurable until all parts of the POD have been upgraded.

During an upgrade APIs will be provided to track the progress. The APIs shall be able to determine the success of the upgrade activities and identify any fallout activities which are required.

Rollback to previous software components must be available if failure criteria are met.

Activities impacting customer service shall be performed in maintenance periods, usually 2 to 4 hours. Expected interruption to subscriber service shall be less than 5 minutes.

3.1.9 Accounting and Status

To operate the SEBA POD, interfaces are required to determine the operational status of the POD. Real time status shall be provided in the following areas:

- 802.1x Authenticator

- 802.1x Diagnostics
- 802.1x Session
- RADIUS Accounting Server
- ONT Status
- ONT Alarm Thresholds
- ONT UNI Port
- Current Optical Data
- Historical Optical Data
- PON
- PON SFP

3.1.10 Performance Management

VOLTHA publishes performance monitoring data to its Kafka bus that NEM can transform into bulk collections of data (such as organized by a collection interval).

NEM can also poll other performance monitoring data for other functions in a SEBA POD, for example using a Redfish API for device performance monitoring. NEM can develop a "transformer" as necessary from another API to a Kafka topic, and thus provide all performance monitoring through a normalized collector in Kafka.

Northbound OSSs/Orchestrators can implement and install an agent as a microservice in the SEBA POD that subscribes to the normalized Kafka collector and transforms the performance monitoring collections to a desired format (e.g., VES for ONAP, IPFIX for another OSS, etc.) for transport to the Northbound OSS/orchestrator.

The SEBA NB API should provide an API to retrieve all PM data and metrics for the "current" interval(s) - e.g. 15-minute and daily.

If it is possible for a Carrier Automation Platform (CAP) to be out-of-sync with the SEBA POD historical PM interval collection upon discovery or re-attachment to a SEBA POD, then the mechanism to re-sync the CAP from the Kafka bus needs to be determined and implemented.

3.1.11 Inventory

Inventory is the definition of the system components and interfaces. In a life cycle view, the planned inventory includes the expected components and interfaces that need to be operating in the infrastructure layer to deliver the POD functions. The actual or discovered inventory requires discovery and validation against the planned inventory to provide the required infrastructure for services and operations of the POD.

3.1.12 Telemetry, Monitoring and Logging, Analytics and Policy Functions

Telemetry involves automatically recording and transmitting data from remote or inaccessible sources to a management system for monitoring and analysis. For this solution, it is encouraged to direct the collected telemetry data to the performance monitoring management subsystem for common processing.

Note that SEBA provides an optional Monitoring and Logging framework. See the SEBA Monitoring & Logging Infrastructure in the SEBA Design Docs. This optional framework is included via helm statements for monitoring and logging functions in the SEBA startup.

Note that the SEBA project currently does not include or plan to include an Analytics engine within SEBA, such as described in the proposed Analytics for CORD project (A-CORD). In the absence of an Analytics engine from the A-CORD project, an operator may develop its own Analytics applications in the POD that could interface to the SEBA Monitoring and Logging infrastructure.

Note that instead of building analytics applications in the POD, an operator may build centralized Analytics and Policy functions in its northbound CAP. The CAP will receives Faults, Telemetry and Performance Monitoring from the NEM adapters (see also sections above for Fault Management and Performance Monitoring) which subscribe to the Kafka bus and the transformers in the SEBA Monitoring & Logging Infrastructure, in order to provide centralized processing for Fault Management, Performance Management, Telemetry and any Analytics and Policy functions.

3.1.13 Automation and Management (includes Exterior APIs)

In a typical deployment scenario there will be up to thousands of SEBA PODs installed.

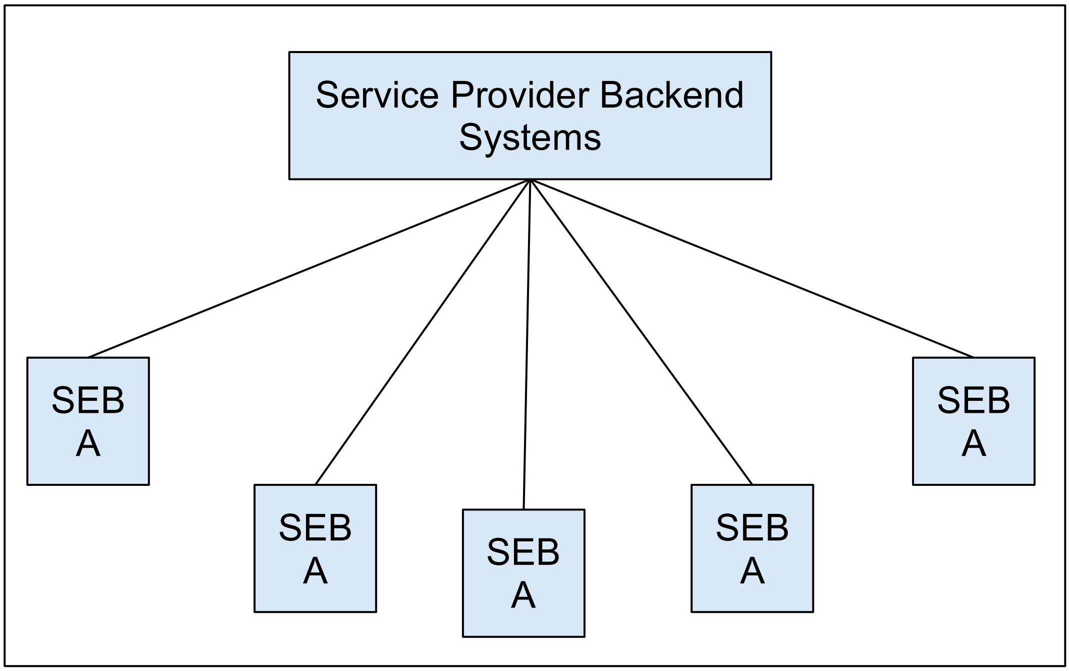

Figure 7: Service Provider Backend Systems to Many SEBA PODs

SEBA should be self-contained and be able to work at any service provider environment. To achieve these goals, we need to have an external adaptation box that sits on top of SEBA Northbound API and may be remotely positioned. An operator may instantiate the SEBA NBI client within a SEBA POD.

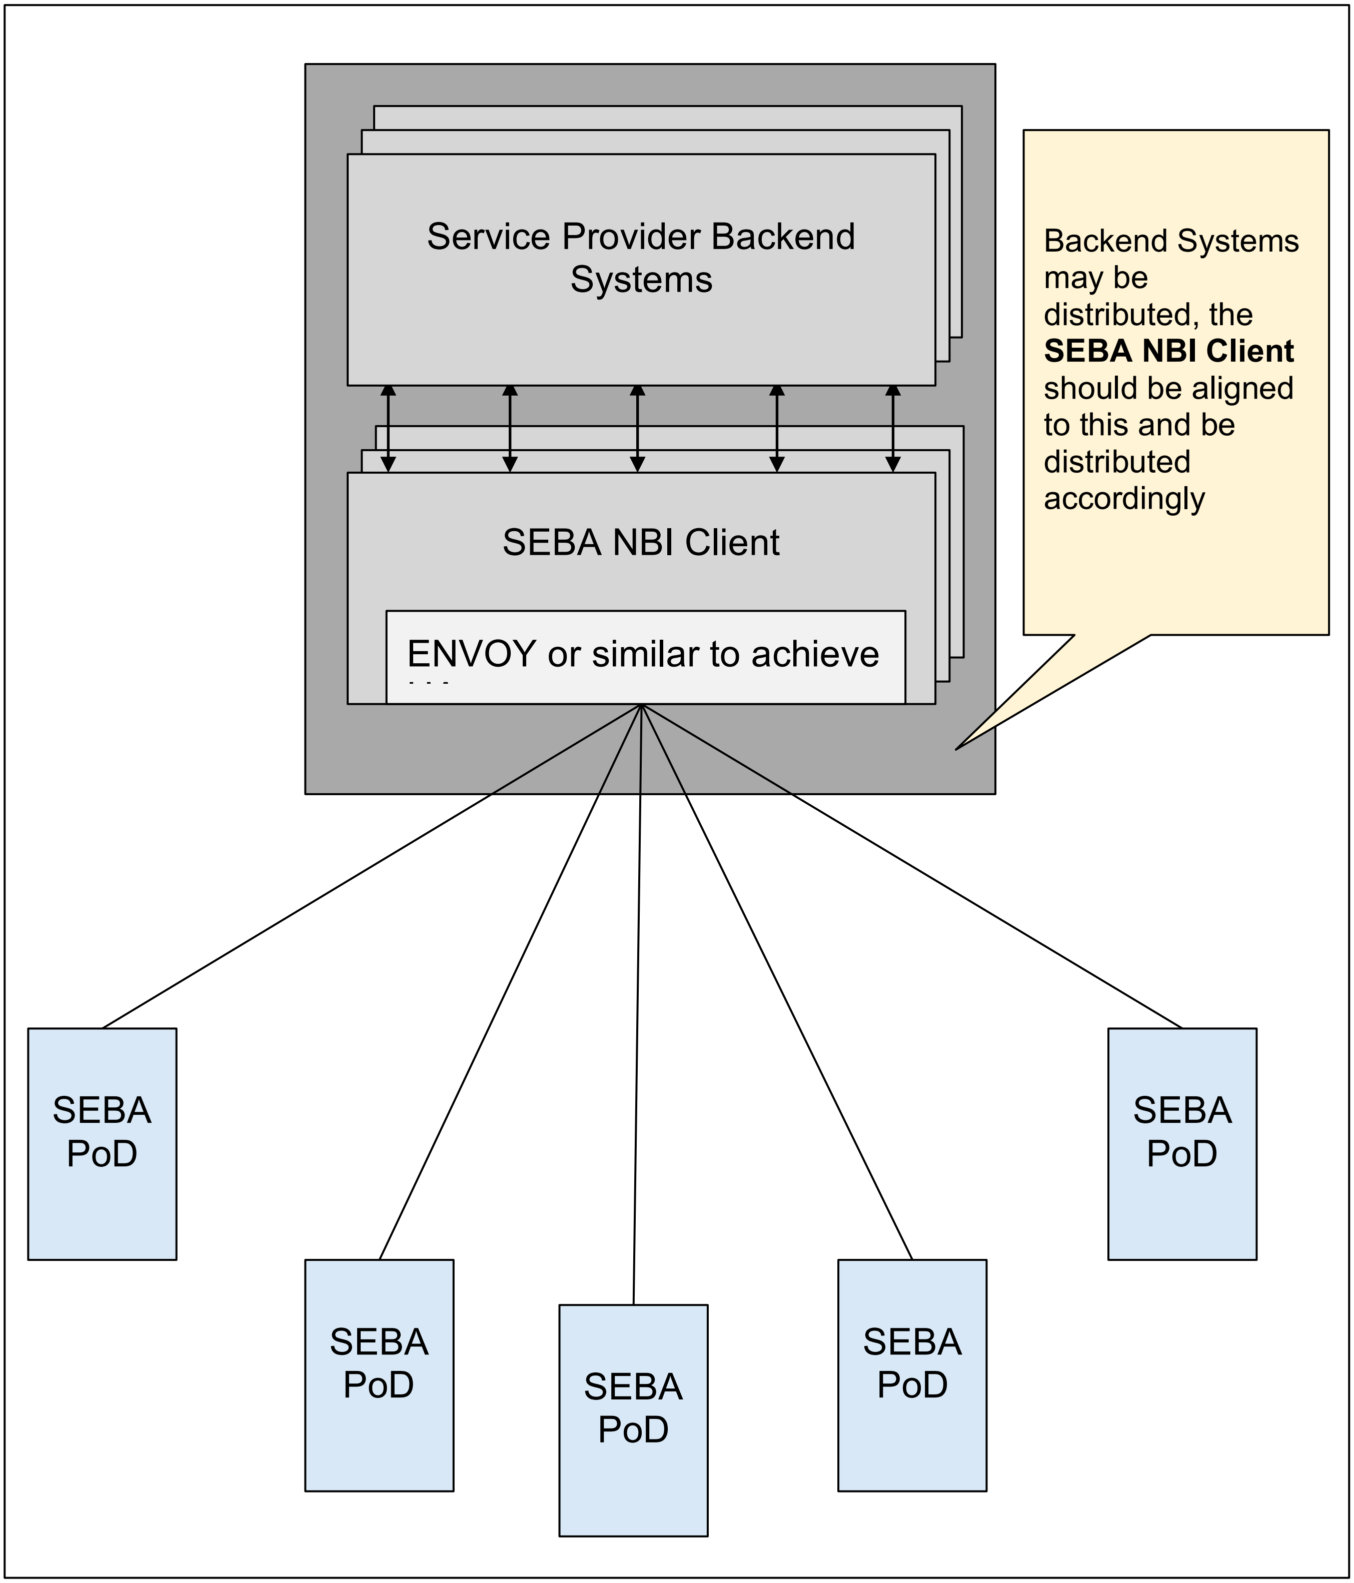
Service Provider Backend Systems
SEBA NBI Client
ENVOY or similar to achieve
Backend Systems may be distributed, the **SEBA NBI Client** should be aligned to this and be distributed accordingly
SEBA PoD
SEBA PoD
SEBA PoD
SEBA PoD
SEBA PoD

Figure 8: SEBA NBI Client Role

The SEBA NBI Client is tightly coupled with the Service Provider backend and will be Service Provider specific.

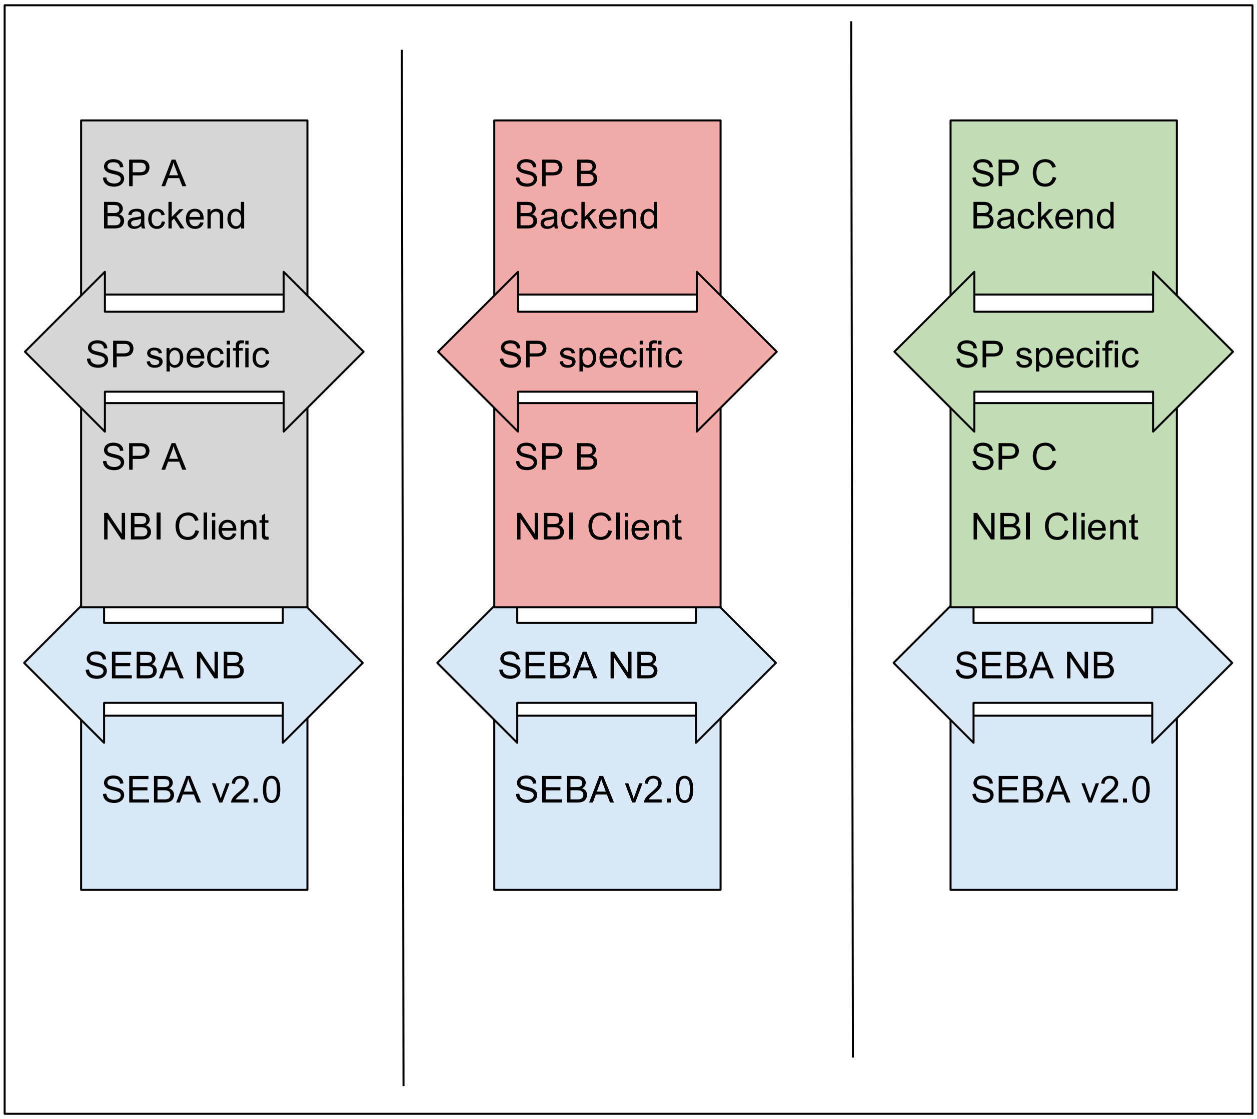

Figure 9: SP Backend, SEBA NBI Client (per SP), and SEBA NB API

SEBA should be self-contained and shall not be affected by any changes that may occur in the deployment environment.

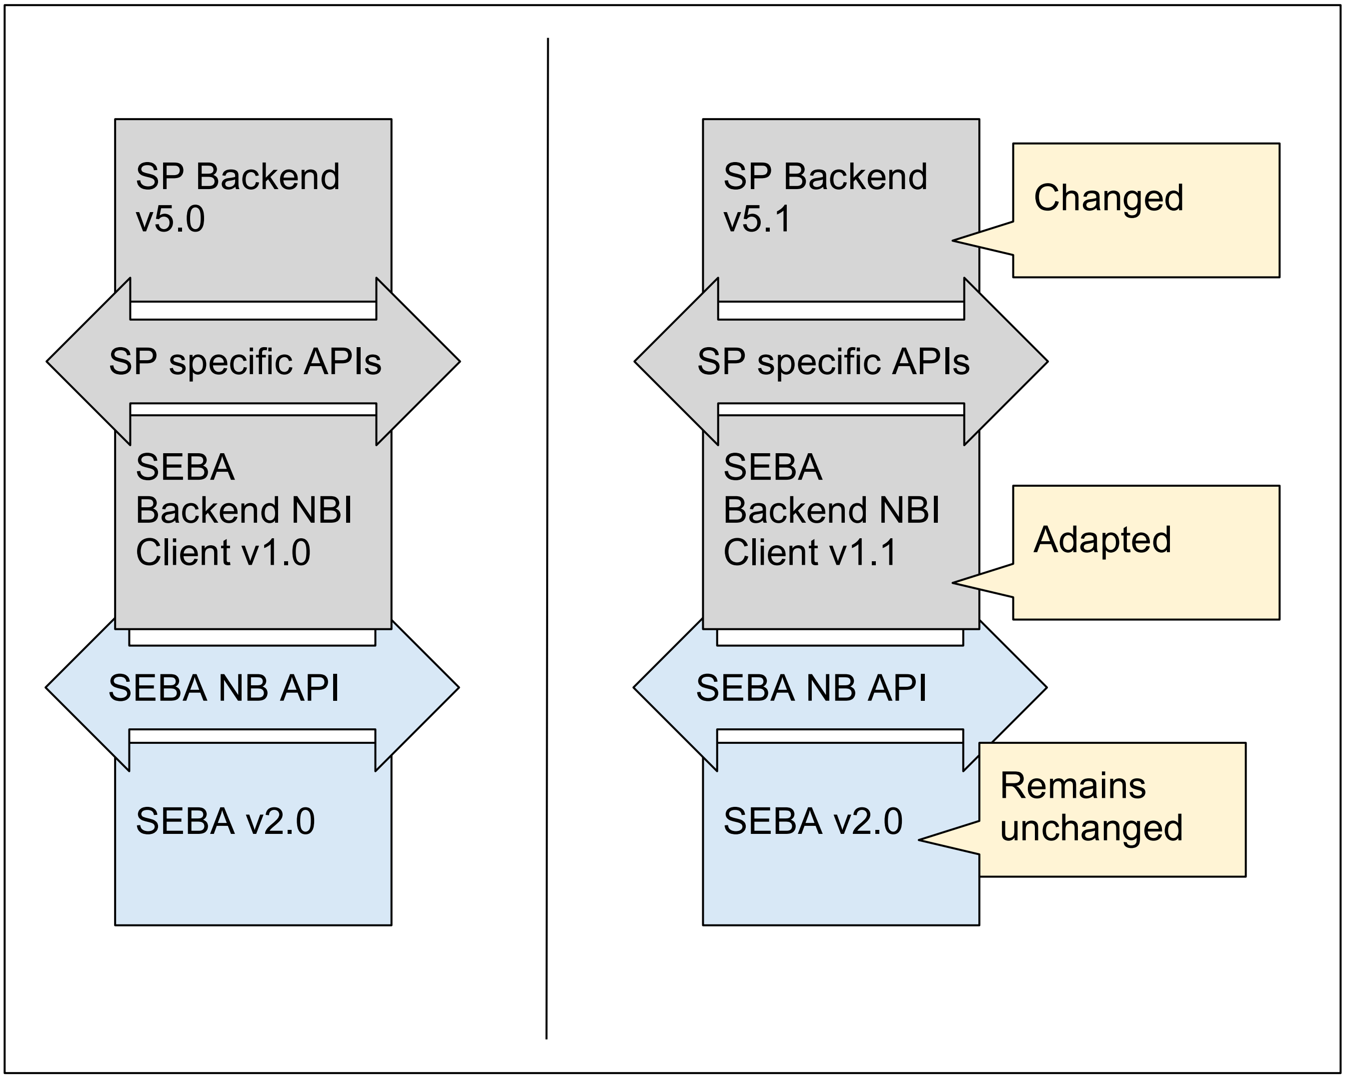

Figure 10: Example - No change to SEBA NB API

Similarly, version upgrades of SEBA shall not affect the backend systems.

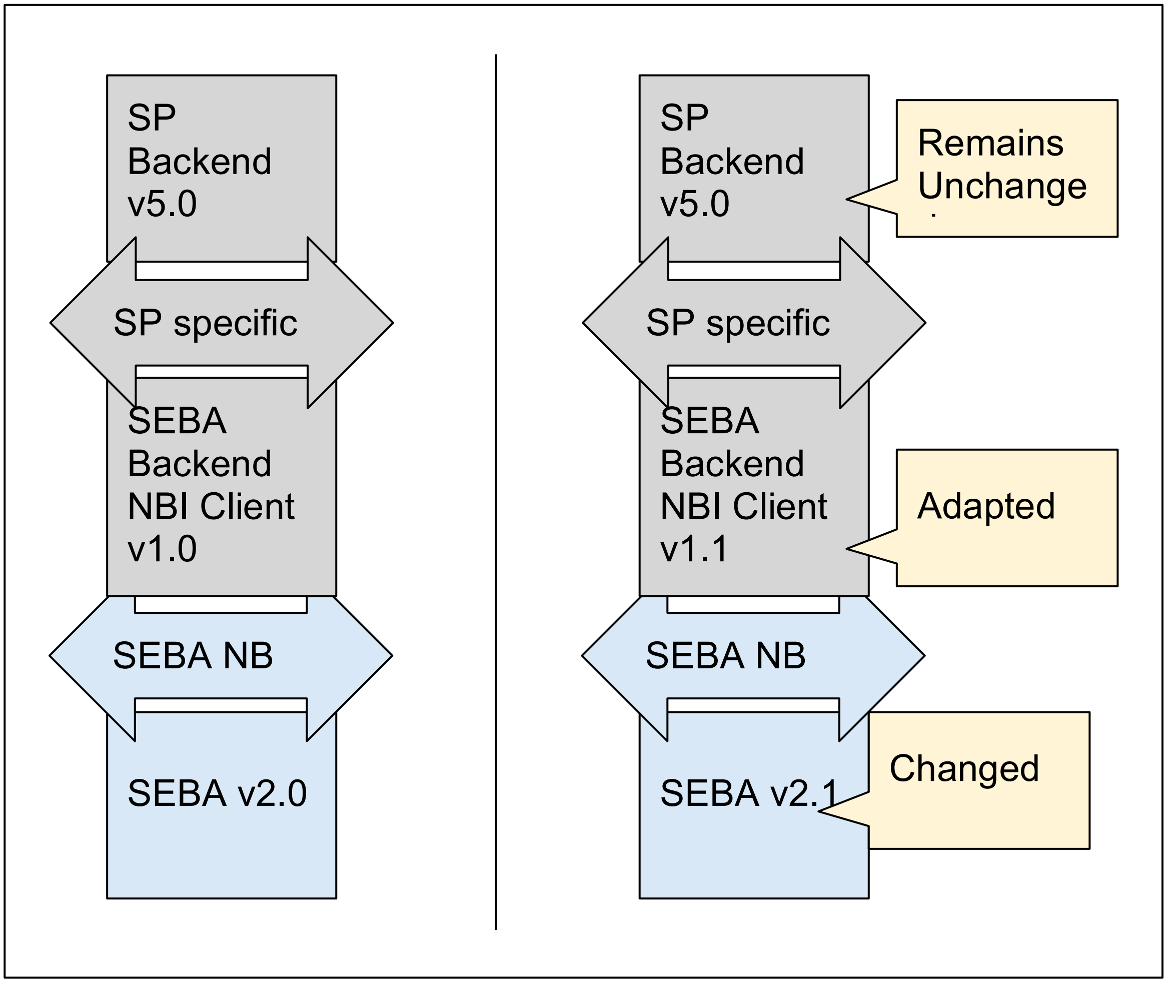

Figure 11: Example - No change to SP Backend

The adapter shall invoke relevant methods on the SEBA NB API and implements a set of callback API's for SEBA to call. The remote invocation shall be made through gRPC.

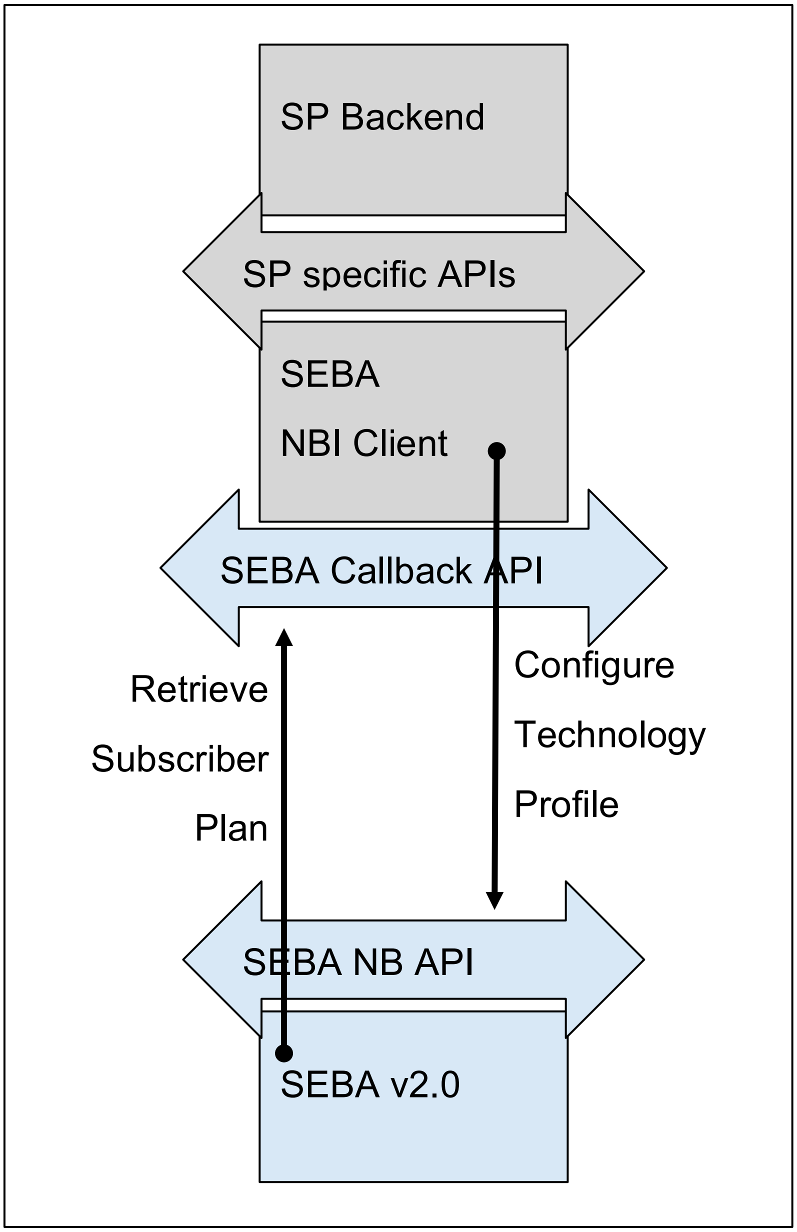

Figure 12: SEBA Callback API

The SEBA NB API will provide the following functionalities:

POD Management

- Provide inventory information for common hardware components
- Download and manage software upgrades for ONOS, ONOS Apps, and VOLTHA components
- Monitor common hardware resources
- Provide detail views on CPU utilization, component states, POD status, Container status
- Status Reporting
- Alarm Management
- Performance Monitoring
- Profile Management
- Define and manage technology profiles required for various hardware adapters
- Service Provisioning

OLT Management

- Assign CLLI associated to specific hardware inventory via serial number (or other unique identifier)
- Retrieve OLT hardware inventory information
- Manage OLT software and upgrades
- Reset OLT hardware
- Manage associated OLT database configurations
- Delete OLT hardware
- Run available OLT diagnostics and retrieve results
- Retrieve inventory information for SFP devices plugged into OLT ports

ONT Management

- Assign ONT to specific OLT port and assigned ONT number via serial number (or other unique identifier)
- Map upstream ONT identifications (OLT CLLI ONT port) to dynamic VOLTHA assignments
- Retrieve ONT hardware inventory information
- Manage ONT software and upgrades
- Reset ONT hardware
- Manage associated ONT database configurations
- Delete ONT hardware
- Run available ONT diagnostics and retrieve results

- Retrieve inventory information for SFP device plugged into the ONT
- Disable the ONT
- Manage the ONT UNI port
- Reset ONT UNI
- Disable the ONT UNI

Status Reporting

- Alarm Management
- Performance Monitoring
- Profile Management
- Define and manage technology profiles required for various hardware adapters
- Service Provisioning

3.1.14 Design in Motion - Use Cases (SEBA POD for PON Technology)

3.1.14.1 Day 0

SEBA HW Installation

- SEBA Rack Installation
- SEBA vOLT Installation & Fiber wiring
- SEBA Compute Nodes Installation

SEBA platform Installation

- Container based
- Uses precompiled container images
- Able to install with a single command line step

SEBA POD Configuration

- Done via human readable site.config (file)
- Applies POD id and basic configurations
- Configures the external management path
- Configures management networking within the POD, between the Compute, ASG, vOLT

SEBA POD Activation

- The SEBA platform runs a sanity test suite and becomes active if it passes
- Calls a POD_Activated(POD_ID) on the callback endpoint.
- If the sanity test fails, use API to raise an error on the callback endpoint

3.1.14.2 Steady State

The use cases in this section may already be covered by management use cases in other sections above. This section could help to be more specific about POD lifecycle management use cases.

Add OLT

- Add container(s) as needed, download software, apply configuration

This section requires expansion with more use cases from SEBA implementations.

3.1.14.3 Fault Detection and Recovery

There can be local or global control loops for fault detection and recovery.

This section should provide uses cases for these control loops.

- Local control loops within SEBA
- Global control loops (outside of SEBA within the Carrier Automation Platform?)

3.1.15 Tooling

Tooling of software is oriented to the effectiveness and efficiency of operations.

Operators have experience with the development of operator portals, logging and search mechanisms, correlation and analytics functions to improve the effectiveness of operations teams to troubleshoot and correct issues proactively or reactively for customers.

Operators should collaborate on Tooling features in SEBA that provide the most value in the exemplar platform.

3.2 SUPPORTING ACTIVITIES

3.2.1 Operational Plan

3.2.1.1 Physical Environment

The physical environment for a SEBA POD is likely to be mostly in a Central Office, but operators may pursue options to deploy a SEBA POD or its elements in a Data Center.

3.2.1.2 Physical Requirements

The physical requirements for a SEBA POD cover multiple areas such as space, rack placement, power, operating temperature ranges, cooling and heat dissipation.

Standards for a Central Office derive from telco industry standards and are referenced in Open Compute Project (OCP) definitions for Telco. Operators may also provide more specifics about their Central Office environments.

Standards for a Data Center derive from the data center industry, and support a wider range of open devices, as the standards may be less constraining for some requirements such as the operating temperature range.

3.2.2 Ecosystem Component Assessment

3.2.2.1 Open Source Software

Open source software should follow the guidelines of ONF as to the open software licenses that ONF projects can use to incorporate open source.

Operators and members contributing code to the ONF open source are responsible to conform to the guidelines of their companies or organizations to contribute code to ONF.

3.2.2.2 Open Hardware

The classification of open hardware will follow the definitions within the OCP, such as the "OCP Accepted" and "OCP Inspired" trademark definitions.

3.2.2.3 Functional Decomposition Supplier Consistency

Supplier consistency in the functional decomposition of the SEBA exemplar platform and into specific implementations is desired and encouraged in order to align with suppliers an developers, while not limiting innovation to improve cost, performance or reliability.

The operators, suppliers and developers should proactively collaborate to communicate about implementations and product roadmaps and evolving open software technologies, and to propose the controlled evolution of solutions.

3.2.3 Operator Specific Addenda (system impacts, etc.)

Operators should provide specific addenda here that require special attention to the SEBA project, as derived out of their experience, requirements, or consequences of their SEBA implementations.

3.2.4 Key Outstanding Questions

This section will capture any current outstanding questions from the provider, supplier or member groups for ONF and SEBA.

There are currently no open questions.

End of

SDN ENABLED BROADBAND ACCESS (SEBA)

Reference Design

Write to rdspec@opennetworking.org with comments or questions.

Open Networking Foundation