## White Paper

# intel.

# Micro-PoP – A Converged vBNG and OLT Solution

#### Radically simplifies Point-of-Presence deployments for PON-based access.

#### Authors

Paul Mannion Intel Corporation

Brendan Ryan Intel Corporation

Jasvinder Singh Intel Corporation

Dr. Ricardo Santos BISDN Dr. Hagen Woesner BISDN Paul Runcy Tibit Communications

> Ricardo Pianta Venko Networks

#### **Table of Contents**

Abstract 1
1. Introduction 1
1.1 The Role of Software 2
1.2 Micro-Pop Use Cases 2
2. Micro-PoP Solution Overview 3
2.1 Hardware Components 3
2.2 Software Architecture 3
2.2.1 Network Interface 3
2.2.2 Date Plane 4
2.2.3 Control Plane 4
3. Performance Analysis 5
3.1 Performance Analysis Env 5
3.2 Performance Results 6
4. The Micro-PoP 8
References9

#### Abstract

This paper describes an innovative approach to combining vBNG and OLT network functions into a single consolidated software-based network element called a Micro-PoP to greatly simplify PON-based broadband access and improve Internet connectivity in such places as industrial parks and small rural areas. The paper explains this new concept and provides details of the underlying hardware components and software architecture of an actual reference design. The results of a performance analysis for this reference design in a simulated environment are provided to demonstrate the feasibility of the solution to meet the requirements of a real deployment.

#### **1. Introduction**

A Point-of-Presence (PoP) for Fiber-to-the-x (FTTx) technologies is usually composed of a set of devices that execute the different functions required to provide Internet access to end users. It involves a few basic functions:

- Authentication/Authorization/Accounting (AAA): The identification of customers and authorization to access a specific service.
- Address Assignment: The AAA framework collaborates with RADIUS servers to assign addresses from the specified pools to the subscribers.
- **Bandwidth/QoS Management:** The function that ensures that each customer receives broadband access services according to their contract.
- **Optical Line Termination (OLT):** The physical endpoint of the passive optical network, which converts the electrical signal to optical signals used in the passive optical network (PON).
- Aggregation: The aggregation of several OLTs into a single broadband network gateway (BNG).

Typically, these functions are executed by specific network elements, which may be distributed across different sites. Setting up a PoP is not trivial and involves different devices, such as OLTs, switches, routers, and servers. PoPs may also vary a great deal in terms of size, serving several hundred to several thousand customers.

In a traditional approach, a large and scalable BNG is located in the ISP central office, with OLTs being geographically distributed.

In circumstances where the concentration of subscribers is small or the local infrastructure is restricted, a Micro-PoP may be more suitable. A Micro-PoP represents a complete broadband access service in a single box, with all the required components, supporting up to 256 subscribers from a shared 10 Gbps link.

2. White Paper | Micro-PoP - A Converged vBNG and OLT Solution



he icons shown here consititute trademarks and copyrights of Telecom Infra Project, Inc. in the United States and other countries Used here under license. Unauthorized use strictly prohibited. All rights reserved

#### Figure 1: Extended Reach to Remote Areas

Alternatively, a 5G modem can be incorporated as a backup or the main broadband uplink.

Micro-PoPs can be appropriate in locations such as:

- Rural areas with small villages
- Industrial parks
- Rooftops and condominiums
- Ports

From a business perspective, Micro-PoPs would allow small entrepreneurs to easily deploy broadband access services with a very small investment.

From an end-user perspective, dedicated 10 Gbps bandwidth would be beneficial for small areas, while also offering extended reach for remote or rural areas.

#### 1.1 The Role of Software

A software-centric approach has enabled a spectrum of deployment models and evolving technologies to be adopted, providing multiple benefits to network service providers such as greater efficiency of space and power, easier maintenance and upgrades, and more granular scalability.

Performance for such a software-centric approach is greatly boosted by technologies such as the Data Plane Development Kit (DPDK), which provides highly optimized packet processing in software that is tightly coupled to the Intel processor CPU architecture. Intel processors provide native instructions and features that specifically accelerate data plane packet processing for access network functions such as the virtual Broadband Network Gateway (vBNG).

The flexibility of a software-centric platform also provides significant benefits to network service providers of consolidated software solutions with co-located network functions, such as BNG and OLT-PON, as described in this paper.

#### 1.2 Micro-PoP Use Cases

In this section we take a closer look at potential use cases for Micro-PoPs.

The proposed solution requires a minimal footprint that would allow connectivity to remote areas with a small



The icons shown here constitute trademarks and copyrights of Telecom Infra Project, Inc. in the United States and other countries. Used here under license. Unauthorized use strictly prohibited All rights reserved

### **Figure 2:** Dedicated Fiber Broadband Access for Small Industrial Areas

number of subscribers where a traditional deployment would not be suitable. For example, in a residential housing complex or a small village, as shown in Figure 1, a Micro-PoP could derive access from the centralized carrier provider and distribute it across a remote area. Technology such as 10 Gigabit Symmetrical Passive Optical Networking (XGS-PON) is capable of a range of 20 km between the most distantly spaced end users, enabling a wide geographical coverage area.

Support for symmetrical connectivity makes XGS-PON solutions especially suitable for corporate applications such as LAN-to-LAN connectivity. The Micro-PoP is thus very appropriate for areas with high demand for bandwidth, such as industrial parks, where there is a concentration of companies in a specific geographic location.

By including wireless access technology such as 4G or 5G in the Micro-PoP, this solution can be of additional benefit in remote areas where the availability of fiber is restricted. The availability of different access technologies may also enable high-availability access strategies based on SD-WAN technologies, as shown below.

In a typical scenario, all data traffic should break out into the Internet as close as possible to the Micro-PoP. This is not necessarily needed for control traffic, which can be centralized in order to improve the manageability of largerscale deployments.

As shown in Figure 3, the control plane of the Micro-PoP can be centralized and run in, for example, a public cloud. In this scenario, a local Ethernet VPN (EVPN) client creates an overlay LAN that makes use of multiple uplinks (typically, a high-bandwidth fiber uplink backed up by a 4G or 5G mobile uplink). By spanning a highly redundant EVPN across all uplinks, Micro-PoPs have an improved reach, even in the case of link failures. At the same time, the local breakout into the Internet assures that traffic costs for the cloud installment remain competitive.

In combination with the SD-WAN client running on the Micro-PoP, the Control and User Plane Separation (CUPS) principle is exploited to reduce the operational cost of remotely managing hundreds of small devices.

The centralization of control also leads to better support for user mobility. New architectural approaches, which point towards the integration of the fixed access network into the 5G core network, are simplified by using tunnel overlays, as depicted in the figure below.

Note, however, that the control plane of the Micro-PoP can remain physically co-located with the data plane, as will likely be the case for smaller installations with only a few Micro-PoPs.



Figure 3: Micro-PoP with Smart Routing/SD-WAN

#### 2. Micro-PoP Solution Overview

The following sections describe the key hardware components and software architecture of an actual Micro-Pop solution design.

#### 2.1 Hardware Components

The hardware components of the proposed solution are as follows:

- Silicom uCPE 1U Half Blade Form-Factor
- Internal Power Supply with Optional External Power Supply for Redundancy
- Intel Atom<sup>®</sup> C3000 Series Processor or Intel<sup>®</sup> Xeon<sup>®</sup> D Processor
- 2 x Tibit MicroPlug OLT SFP
- 4 x 10GbE Intel<sup>®</sup> Ethernet Server Adapter X710 SFP+
- 1GbE Admin Ports
- Dual 4G/5G Sim Card

A Micro-PoP based on the Intel Atom processor may be deployed on the Silicom Cordoba or Madrid Universal Customer Premise Equipment (uCPE) platforms shown below. The Micro-PoP may also be deployed on uCPE platforms based on Silicom and the Intel® Xeon® D processor, such as Barcelona or Bilbao. See Figure 4 and Reference 1 for further details on Silicom uCPE platforms.



Silicom Madrid

**Figure 4:** Silicom Cordoba and Madrid, Intel Atom<sup>®</sup> Processor-Based uCPE

A key innovation in this solution is the insertion of a pluggable OLT, the MicroPlug OLT (shown in Figure 5), into two ports of a general-purpose compute board, an Intel® Ethernet 700 Series Network Adapter 4x10G NIC.



#### Figure 5: Pluggable OLT

The Tibit MicroPlug OLT is a small, low-power, low-cost pluggable device that allows a high-performance Ethernet port to become an XGS-PON port. It provides great flexibility at the edge, enabling Ethernet switches, routers, or almost any device with an Ethernet port to become an FTTx device. By definition, it is a software-defined networking (SDN) solution that can be managed from anywhere in the service provider infrastructure, even the cloud. See Reference 2 for further details on the Tibit MicroPlug OLT.

#### 2.2 Software Architecture

A reference software architecture for a Micro-PoP solution is shown below.



#### Figure 6: Micro-PoP Software Architecture

#### 2.2.1 Network Interface

The simplest Micro-PoP deployment uses two interfaces in an X710 NIC. One interface (Port 0) is connected to a Tibit pluggable OLT module, and the second (Port 1) is connected to the uplink next hop, as shown in Figure 6. The NIC can also benefit from Dynamic Device Personalization (DDP) and, alongside the instantiation of virtual functions, traffic can be intelligently steered in the CPE data plane. More specifically, in this Micro-PoP vBNG setup, a Point-to-Point Protocol over Ethernet (PPPoE) DDP filter is used, which enables PPPoE session control frames to be separated from data plane traffic at Port 0. The uplink port (Port 1) does not have control plane forwarding requirements. Therefore, virtual functions that are created on top of it are used solely for forwarding data plane traffic. The simplest Micro-PoP vBNG deployment can be done with a single data plane virtual function on the uplink port.

However, the OLT Port needs two management frames to redirect two types of control plane traffic:

- **Tibit's API (TAPI) Management Frames:** Management traffic exchanged between Tibit's PON Controller and the pluggable OLT module(s).
- vBNG PPPoE Session Management Frames: PPPoE session control traffic (e.g. PADI/O/R/S), exchanged between the vBNG PPPoE server and the PPPoE clients.

This solution therefore requires a minimum of three virtual functions, two for the control plane and one for the data plane. However, this architecture can scale the number of vBNG instances by creating additional data plane virtual functions per port and assigning them to other Micro-PoP vBNG container applications.

#### 2.2.2 Data Plane

The Micro-PoP vBNG data plane is based on DPDK's IP Pipeline application. The application's design allows for the creation of modular pipeline blocks that operate with matchaction tables. Each pipeline block can also be assigned to a dedicated CPU core.

The Micro-PoP vBNG IP Pipeline application features two series of pipeline blocks, the downlink and uplink pipelines. Each pipeline processes packets received on an input port, which are then sent through a separate output port. For example, traffic from the downlink pipeline enters through a data plane virtual function created on the uplink port and exits through a data plane virtual function created on the OLT port. Each pipeline block in both the uplink and downlink pipelines is assigned to a separate CPU core.

#### **Downlink Pipeline Overview**

**ACL:** This firewall forwarding table, with 5-tuple lookup, can be programmed with static rules that are common to all subscribers or with rules specific to each subscriber.

**Traffic Management:** All downstream packets pass through a 5-level (port, subport, pipe, traffic class, and queues) hierarchical QoS (HQoS) scheduler that assigns guaranteed bandwidth to different subscriber services. For each subscriber, the HQoS scheduler assigns multiple queues which are mapped to traffic classes (up to 13) to host the traffic for different services.



#### **Uplink Pipeline Overview**

**ACL:** As with the downlink pipeline, this firewall pipeline block forwarding table has 5-tuple lookup and is programmed with a set of rules that are applied to all subscribers or individually on a per-subscriber basis.

**Flow Classification:** At this stage, subscriber traffic is classified using access packet header fields to identify the session. Subsequently, traffic metering action (e.g., srTCM or trTCM algorithms) is performed on the subscriber session traffic to determine its compliance with a service contract. The subscriber flows that conform to a specified rate are sent to the next stage of the pipeline while packets violating the contract are dropped.

**Routing:** This is the final stage of the uplink pipeline where the subscriber PPPoE headers are stripped and the packets are routed to the next hop network interface for transmission. If required, any core network encapsulations, such as Multiprotocol Label Switching (MPLS), are applied here.

#### 2.2.3 Control Plane

The Micro-PoP control plane integrates both OLT and vBNG management components. Both are summarized in the following sections.

#### **OLT Control Plane**

The OLT control plane ensures that the pluggable OLT modules connected to the Micro-PoP are operational and can forward the data plane traffic to and from the BNG. For example, the VLANs corresponding to the subscriber traffic that the BNG forwards to the PON must be allowed between the Optical Network Units (ONUs) and the OLT. This control plane is formed of the following three components:

**Tibit PON Manager:** The Tibit PON Manager, a web application that enables the configuration of the Tibit pluggable OLTs. It interacts with a MongoDB database by reading the available OLT state, displaying current configurations in the web UI and updating the database with any changes performed through the web UI.

**Mongo DB:** The database used to store the state and configuration data for each OLT connected to the Micro-PoP CPE.

**Tibit PON Controller:** The application responsible for the communication with Tibit pluggable OLT modules via the Tibit API. The PON Controller reads configuration data from the MongoDB and applies it to the OLT modules if it differs from the current state.

- 1GbE Admin Ports
- Dual 4G/5G Sim Card

A Micro-PoP based on the Intel Atom processor may be deployed on the Silicom Cordoba or Madrid Universal Customer Premise Equipment (uCPE) platforms shown below. The Micro-PoP may also be deployed on uCPE platforms based on Silicom and the Intel Xeon D processor, such as Barcelona or Bilbao. See Reference 1 for further details.

#### vBNG Control Plane

The vBNG control plane is responsible for the session state of each subscriber and the data plane forwarding state. Since the control plane is decoupled from the data plane, it is possible to deploy the control plane components inside the Micro-PoP CPE or remotely in a separate server or in a virtual machine. This reference vBNG control plane is composed of the following components:

- Accel-ppp: The PPPoE or IPoE server, responsible for handling the subscriber session protocol state. It has three main roles in the vBNG control plane architecture:
  - **Subscriber State:** Exchanges PPPoE session control traffic with the control plane virtual function and keeps state information for each subscriber.
  - **RADIUS Client:** Communicates with the AAA server using the RADIUS protocol. The AAA server can be located inside or outside the Micro-PoP CPE.
  - Interface to Session Control Interface: Signals the Packet Forwarding Control Protocol (PFCP) control plane when to start and terminate its sessions by publishing session events through the Redis bus (e.g., session starting, finishing).
- **PFCP CP:** The PFCP control plane endpoint, responsible for the session control interface between the vBNG control and data planes. It manages the subscriber forwarding state though the creation, modification, and deletion of sessions in the respective PFCP user plane endpoint. Apart from the original PFCP stack implementation, it also includes extensions from BBF TR-459 to handle broadband case-specific information elements in its messages.
- **PFCP UP:** The PFCP user plane endpoint programs the subscriber session forwarding rules on the vBNG data plane instance based on the information exchanged with the PFCP control plane. The default API to the IP Pipeline application uses telnet as its communication protocol, though it can also support other data plane backends, such as P4 Runtime or OpenFlow.
- **Redis:** A high-performance, low-complexity message broker, used as an interface between Accel-ppp and the PFCP control plane to exchange session events.
- **FreeRADIUS:** An optional AAA server. When deployed in the Micro-PoP CPE, it uses a database backend (e.g. MySQL), where the subscriber data is stored.

#### **3. Performance Analysis**

#### 3.1 Performance Analysis Environment

Figure 8 shows the testing environment used in an evaluation of the Micro-PoP design described in this paper. To evaluate its performance, the number of concurrent PPPoE sessions was varied between 32 and 1024 and L2 frame sizes were varied between 64 and 512 bytes. Five iterations were run of each test and the received L2 throughput and total received packet rate was measured for each iteration.



#### Figure 8: Micro-PoP Performance Analysis Environment

Due to the different computational requirements of the uplink and downlink pipelines, the tests were conducted separately for both scenarios. Table 1 shows the number of forwarding entries installed on each pipeline instance. Note that one CPU core was used per pipeline instance, on both the uplink and downlink scenarios. One PPPoE session per subscriber was also assumed.

#### Table 1: Number of Forwarding Entries per Pipeline Instance

Table	Number of entries
Uplink/Downlink Firewall	1 (static configuration)
Uplink/Downlink Routing	1 per subscriber (32-1024)
Uplink Flow Classification	1 (static configuration)
Downlink Traffic Management	1 per subscriber (32-1024)

The system configuration details for Micro-PoP performance analysis are provided in the following tables. Data plane configurations based on Intel Atom processors and Intel Xeon D processors were evaluated.

#### Micro-PoP Data Plane Server Configuration – Intel Atom® Processor (DUT)

Item	Description
CPU	Intel Atom <sup>®</sup> processor C3958 2.00GHz 16 Cores
Platform	HawkeyeTech HK-2010
BIOS and Microcode	BIOS: 5.13, 01/23/2019 Microcode: 0x2e
Memory	16 GB 2400 MHz DDR4
NICs	Intel® Ethernet Server Adapter X710 Quad- Port 10GbE SFP+ (NOTE: 2 of 4 ports used) Intel® Ethernet Server Adapter I350 1GbE (Control plane)
Operating system	Ubuntu Server 20.04.2 LTS
Linux kernel version	5.4.0-47-generic
DPDK version	20.11
Tester and Date	BISDN, 06/07/2021

#### Micro-PoP Data Plane Server Configuration – Intel® Xeon® D Processor (DUT)

ltem	Description
CPU	3 <sup>rd</sup> Gen Intel® Xeon® D processor 10-Core Pre-Production Sample
Platform	Intel Brighton City Customer Reference Board
BIOS and Microcode	BIOS: IDVLCRB1.86B.0019. D04.2103170050, 03/17/2021 Microcode: 0xfd0001d0
Memory	8 GB 2666 MHz DDR4
NICs	Intel® Ethernet Server Adapter X710 Quad- Port 10GbE SFP+ (NOTE: 2 of 4 ports used)
Operating system	Ubuntu Server 20.04.2 LTS
Linux kernel version	5.4.0-47-generic
DPDK version	20.11
Tester and Date	BISDN, 06/16/2021

#### Micro-PoP Control Plane Server Configuration

Item	Description
CPU	Intel® Xeon® Gold 6238R CPU 2S 112 Cores 2.2GHz
Platform	Supermicro X11DPi-N(T)
BIOS and Microcode	BIOS: 3.3a 09/18/2020 Microcode: 0x5003102
Memory	256 GB 2933 MHz DDR4
NIC	Intel <sup>®</sup> Ethernet Network Adapter 25G Quad-Port E810-XXVDA4 (NOTE: Lower-speed NIC can be used)
Operating system	Ubuntu Server 20.04.2 LTS
Linux kernel version	5.8.0-48-generic

#### Traffic Generator (T-Rex) Server Configuration

Item	Description
CPU	Intel® Xeon® Gold 6238R CPU 2S 112 Cores 2.2GHz
Platform	Supermicro X11DPi-N(T)
BIOS and Microcode	BIOS: 3.3a 09/18/2020 Microcode: 0x5003102
Memory	256 GB 2933 MHz DDR4
NIC	Intel® Ethernet Network Adapter 25G Quad- Port E810-XXVDA4
Operating system	Ubuntu Server 20.04.2 LTS
Linux kernel version	5.4.0-48-generic
T-Rex version	2.89

#### Additional Equipment

Item	Description
OLT	Tibit TXM-MPOLT-05I
ONU	BFW WAC-150
Switch	Accton AS4610-30T

#### **3.2 Performance Results**

Data plane performance results are shown in the following graphs for the performance analysis environment described in the previous section.

Note that because Forward Error Correction (FEC) mode is enabled in the Tibit MicroPlug OLT, the maximum effective traffic rate for Ethernet frames is approximately 8.5Gbps due to overhead of FEC, XG-PON Transmission Convergence (XGTC) and XG-PON Encapsulation Method (XGEM).



**Figure 9:** Micro-PoP Downstream L2 Throughput – Intel Atom<sup>®</sup> Processor-Based Data Plane



Figure 12: Micro-PoP Upstream Packet Rate – Intel Atom® Processor-Based Data Plane



Figure 10: Micro-PoP Downstream Packet Rate – Intel Atom® Processor-Based Data Plane



**Figure 13:** Micro-PoP Downstream L2 Throughput – Intel<sup>®</sup> Xeon<sup>®</sup> D Processor-Based Data Plane



**Figure 11:** Micro-PoP Upstream L2 Throughput – Intel Atom® Processor-Based Data Plane



**Figure 14:** Micro-PoP Downstream Packet Rate – Intel® Xeon® D Processor-Based Data Plane

#### 8. White Paper | Micro-PoP - A Converged vBNG and OLT Solution



Figure 15: Micro-PoP Upstream L2 Throughput – Intel® Xeon® D Processor-Based Data Plane

In a practical setting, the total 10 Gigabits-per-second would be asymmetrically shared between upstream and downstream traffic in a 10/90 or 20/80 proportion. Also, typically, upstream frames are on average smaller than downstream frames. This is mostly due to the nature of residential Internet access, which is dominated by video streaming. For this paper, both traffic directions have been measured separately from one another.

In the uplink scenario, the Atom-based Micro-PoP can achieve a maximum of approximately 7 Mpps. With such packet processing limits, it is possible to achieve that maximum throughput for all the uplink sessions with 192-byte frames. With the Micro-PoP based on the Intel Xeon D processor we observed line-rate throughput for any packet size. When sending the smallest frame size of 64 bytes, the Intel Xeon D processor Micro-PoP handled more than 12 Mpps without any performance bottlenecks.

Due to the higher computational requirements for the HQoS pipeline stage, the number of packets processed (and therefore the measured throughput) in the downlink pipeline was lower than the uplink pipeline measurements with the Intel Atom processor-based Micro-PoP. Nonetheless, we observed that the maximum downlink throughput was achieved starting from 256-byte frames in all of the sessions tested in the Intel Atom processor-based Micro-PoP. Similarly, as in the upstream measurements, the Intel Xeon D processor-based Micro-PoP reached line-rate throughput with all packet sizes.



**Figure 16:** Micro-PoP Upstream Packet Rate – Intel<sup>®</sup> Xeon<sup>®</sup> D Processor-Based Data Plane

#### 4. The Micro-PoP - A Unified vBNG and OLT Solution that Simplifies PON-Based Access

The disaggregation trend in the telecommunications industry has created a great deal of room for innovation. When the different pieces are separated, enormous possibilities are created by combining those pieces in a different fashion and taking advantage of powerful COTS hardware. A software-based solution on a general-purpose platform with Intel processor CPUs offers the flexibility to enable deployment at the pace of innovation.

The increased availability of fiber connections offers the possibility of high-speed, broadband access to areas where the concentration of customers is not sufficient for a large deployment. Self-contained solutions such as the Micro-PoP, which demand fewer resources, become a viable alternative.

Tibit's pluggable OLT, installed in a powerful and versatile Intel processor CPU-based network function platform, is capable of running all the necessary software components in the same box, creating an opportunity for symmetric broadband services that can be built with minimal, costeffective infrastructure.









#### References

- 1. Silicom UCPE Platforms: https://www.silicom-usa.com/4g-5g-solutions/5g\_fwa https://www.silicom-usa.com/pr/edge-networking-solutions/edge-cpes
- 2. Tibit Microplug (OEM) https://tibitcom.com

Performance varies by use, configuration and other factors. Learn more at www.Intel.com/PerformanceIndex.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Your costs and results may vary.

Intel technologies may require enabled hardware, software or service activation.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others. 0521/DL/ICKK/PDF 347021-001 EN