Topology and component description of DriveNets Network Cloud as a sustainable solution for multipurpose networking

Revision 1.0

Author: Run Almog

DriveNets

Revision History

Revision	Date	Author	Description
1.0	18/05/2021	Run Almog	OCP Draft 01
1.1			

Table of Contents

1. Copyright License

2. Scope

This document describes the high-level connectivity of a Network cloud from a component and use case perspective.

The Network Cloud is an implementation of the OCP Distributed Disaggregated Chassis (DDC) which includes servers, switching, fabric and packet processing white boxes, interconnecting cables, and software that defines the functionality of the deployment. Every component in the network cloud is a standalone element which can be used in different topologies to fulfil a required network use-case.

The document details a specific set of components used in a network cloud deployment. However, replacing any component with another DDC standard compliant element (HW or SW) is a valid alternative.

The contribution being made here is for the concept of mounting multiple interchanging network applications over a mutual hardware abstraction layer as a standard (OCP accepted) infrastructure which by itself is an OCP contribution (DDC)

The OCP definition of DDC can be found here: https://www.opencompute.org/documents/ocp-hardware-specifications-and-use-case-description-for-j2-ddc-routing-system-pdf

3. The Various Network Cloud Components

A Network Cloud implementation can come in a formation of a standalone element which hosts all the functionality within a single box or in a cluster (Clos) topology where different standalone boxes act their separate roles. The standalone mode of operation white box is also a component in the cluster implementation.

Figure 1 shows the full topology of a Network cloud which comprises of four different purposes of the white boxes. Two types of boxes, NCP and NCF to construct the data plane in a Clos topology and two white boxes to construct the control and management planes in a star topology with the NCC hosting the base OS and various containerized functionality and the NCM acting as a dedicated control distribution plane that directly connects the NCC to all other Network Cloud components

All forwarding devices can be installed using the Open Network Install environment – ONIE (https://opencomputeproject.github.io/onie/design-spec/hw_requirements.html)

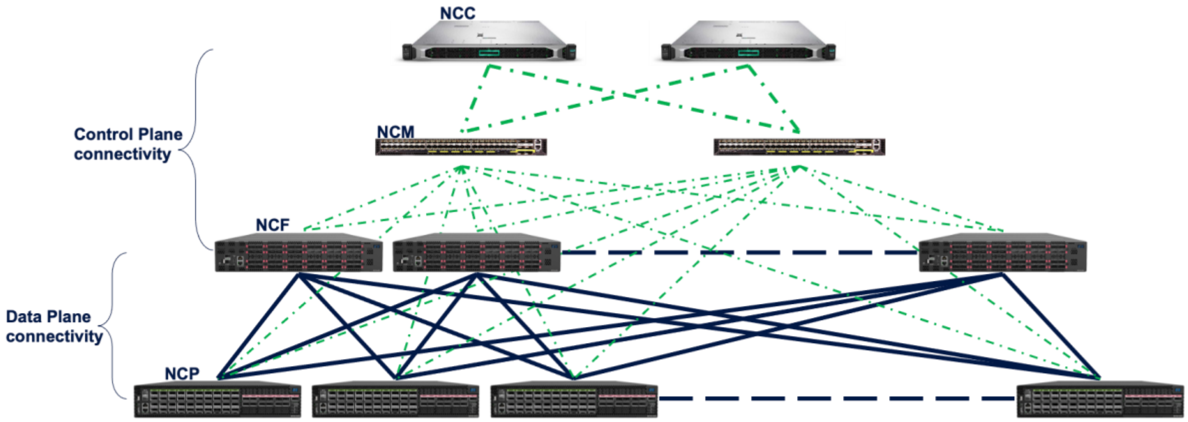

Figure 1: High-Level View of a Network Cloud platform

Cluster sizes are defined as the following:

1. Standalone: (1 NCP)
2. Small: (2 NCF, 2-4 NCP, 2 NCM, 2 NCC)
3. Medium: (7 NCF, 1-24 NCP, 2 NCM, 2 NCC)
4. Large: (13 NCF, 1-48 NCP, 4 NCM, 2 NCC)

Different cluster sizes are also available while the above is given as common examples and guideline.

The following table lists current* components in the Network Cloud:

Network Cloud Name	DDC Name	Component Specification
NCP	*DCP100/ DCP400*	*{2-RU, EIA 19", Max Depth 30"} with 40x100G QSFP28 (for ports) and 13x400G QSFP-DD (for fabric) pluggables / 10x400G QSFP-DD (for ports).*
NCF	*DCF48 / DCF48*	*{2-RU, EIA 19", Max Depth 30"} with 48/24x400G QSFP-DD (for fabric) pluggables.*
NCM	*DCM*	*{1-RU, EIA 19", Max Depth 30"} with 48x10G/1G SFP and 6x100G QSFP28 Pluggables.*
NCC	*DCC*	*{1-2RU, EIA 19", Max Depth 30"} TP76200/TP450 Level 3 Compliant COTS server with NICs to support required number of 100G interfaces.*

* additional white-box elements are added by the ODM as technology evolves

NCP – Network Cloud Packet Forwarder

The NCP is also operational as a standalone implementation of the Network Cloud and run the full operating system and a containerized network function by itself. In such case, the fabric interfaces of the device remain idle and are not utilized.

The NCP is carrying the following components: Jericho2 (88690) NPU with 8G of HBM for deep packet buffering from Broadcom, the external OP2 BCM 16K processor to support enhanced route scale and statistics which can be used for different purposes depending upon the use case, and a x86 (D-1548 with 8-core @ 2.0GHz. 2x32G DRAM, 1x128G SSD) Intel processor

the sketched architecture of the NCP (DCP) devices as defined by OCP spec is brought here:

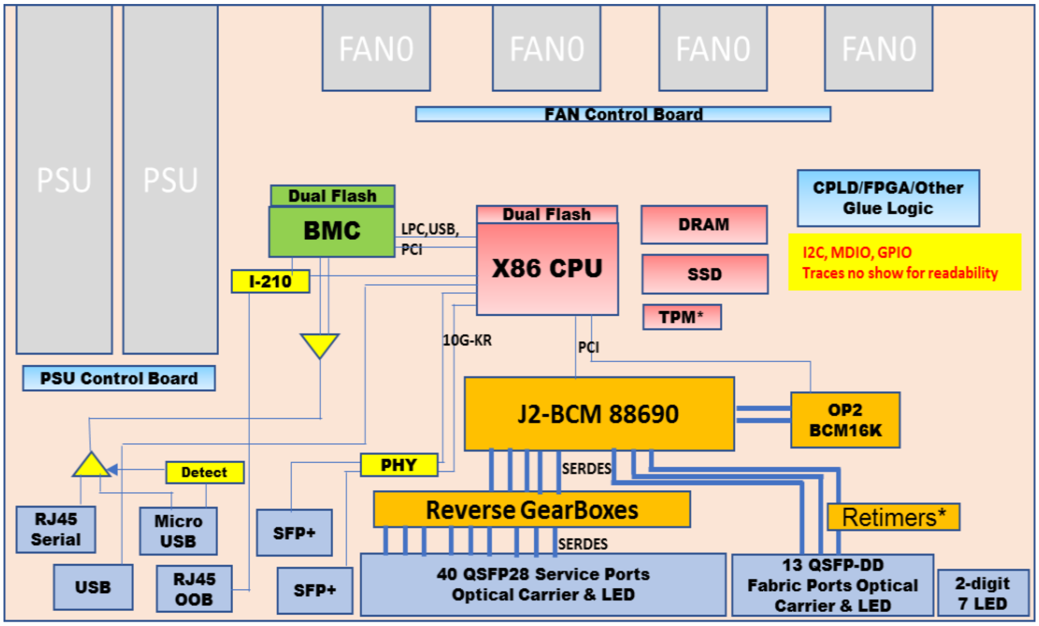

Figure 2: Architecture and Components of the NCP White-Box

NCP 100GbE interfaces support 4x10G or 4x25G QSFP break-out optics to provide connectivity to legacy equipment for various field of use.

The following NCP devices are available under the Network Cloud implementation:

S9700-53DX (UFISpace), 40x100G user ports + 13x400G fabric ports (OCP Accepted)

S9700-23D (UFISpace), 10x400G user ports + 13x400G fabric ports (OCP Accepted)

S9700-76D with TCAM (UFISpace), 32x400G user ports + 40x400G fabric ports (OCP submission pending)

AGCXD40S (Delta Networks), 40x100G user ports + 13x400G fabric ports (OCP submission pending)

NCF – Network Cloud Fabric

The NCF is only usable as the fabric components of a cluster. Its interfaces are all 400G fabric ports and it is based on the Broadcom Ramon (BCM88790) ASIC. The NCF devices can be repurposed for constructing a network cloud cluster of any required size and regardless of if the user ports in use are 400GbE, 100GbE, 25GbE, or 10GbE according to the OCP being used.

Supported NCF devices include the S9705-48D (UFISpace – OCP Accepted) and the AGCC048 (Delta Networks - OCP submission pending)

The FRU (Fans and power) are common to both the data plane white boxes (NCP and NCF) and can be grouped to a mutual maintenance inventory

NCM – Network Cloud Management

The NCM is a carrier access/aggregation switch carrying 48xSFP+ ports and 6x100G ports in a 1RU form factor. This device can also act as a data-center top of rack switch and is installed using standard ONIE

Available NCM devices include the Acton/Edgcore AS5916-54XL (OCP Accepted)

NCC – Network Cloud Controller

The NCC is a standard COTS server running native Linux operating system and can be used as a standard Linux server for any other purpose besides NCC.

Available NCC hardware include the HP Proliant DL380P Gen10 (868703-B21) and Dell PowerEdge R640 (High end) for large clusters and HP Proliant DL360P Gen10 (86959-B21) and Dell PowerEdge R640 (low end) for Smaller clusters.

4. The Various Network Cloud Use Cases

The original definition of the DDC was set to replace a P and a PE router in high capacity and stringent network availability requirements similar to those met by the high-end chassis-based router implementations. The Network Cloud implementation extends this requirement and is established as a cloud native platform which runs containerized network functions on the data plane white boxes as well as the control devices on a cluster or a standalone deployment.

Such functions include a variety of routing functions suitable for specific use-cases such as core functions (P/PE as defined by AT&T), internet peering, VPN, aggregation, etc. as well as non-routing network functions such as firewall (FW), distributed denial of service (DDoS), deep packet inspection (DPI), Load Balancing (LB), 5G core (5GC), broadband Network gateway (BNG) and in fact, any type of disaggregated network function can be mounted onto the Network Cloud infrastructure and run as a containerized service within the network itself.

With the size of the infrastructure defined by the number of boxes in the deployed topology and the functionality defined separately by whichever containers are mounted and run within the NOS, the use cases which can be fulfilled using a network cloud are a superset of routers and dedicated network appliances as well as being a readied platform to mount any future service which will be implemented in a modern cloud native manner

5. The Sustainability Nature of The Network Cloud

Sustainability solutions are enabled when equipment can service multiple different purposes which are controlled via software. Such implementations are very common to anything cloud related as it is an exact example of generalized hardware and specific software-controlled usage of that hardware.

Network Cloud does exactly that to networking. A Disaggregated and Distributed network element is constructed of generalized hardware elements, well defined by an open standard and such that operate whichever application is dictated by a separate software that can operate a cluster of such devices.

Taking a Network Cloud cluster implementing a core router and “breaking” it into 2 separate routers running aggregation and Provider edge functions in two separate location as well as mounting a DDOS application on one and a DPI function on the other is a valid option. Simply reconnect the devices in the new desired scale and mount them with the relevant software piece that defines the required functionality.

This prolongs the life span of any of the Network Cloud white-boxes even within the same network, within the same economy and within the same geography.

Further variants of reuse examples for these white boxes can be when the network evolves into new speeds with the fact that the same fabric devices can serve the new incoming technology (no limitation to the number of NCF used in a cluster as you would find in a chassis).

Another variant of a sustainable solution can be found when what was once a flashy high-end deployment becomes outdated for its deployment environment (such as the example of HPC clusters). Decommissioning a Network Cloud cluster and rebuilding it in a developing country where it can proudly serve for several more years becomes a possibility due to the decoupling of the vertically integrated traditional router implementation and replacing it with a software/hardware or, more importantly, a function/capacity separation that enables flexibility in repurposing the existing hardware to serve evolving use cases.

The architecture of the network cloud is the enabler for this separation to take place. Creating the infrastructure for Control and User Plane Separation (CUPS) dictates that the hardware should be built with a level of simplicity so that it can be easily sourced and easily addressable by various applications. Software should be built as cloud native to seize this advantage and agnostic to the hardware so that it can scale up and down on demand as well as be mounted onto various data plane options.

Recapping on the various network cloud use cases which start from core routing to other routing use-cases and extended into other network functions like FW, DDoS, DPI, LB, 5GC, BNG, etc. such functions are agnostic to the data plane and are relating to the data plane via the network cloud abstraction layer. This is where the original purpose for which the network was placed for changes from a constraint to an opportunity. Adding a function on top of an existing deployment or changing functionality into something else extends the longevity of the CapEx invested into the network.

Further than this, functions that run on top of the network cloud are not necessarily the well known network functions mentioned above but could also be new functions which are developed to capitalize on the extended capabilities of the network (e.g. 5G enabled apps), create differentiation for the service provider (e.g. console-free gaming apps), analytics based functions which benefit from network generated data points, etc. these functions can be coupled to a still running “traditional” network function or be mounted on the cluster after it served it’s original purpose.

6. A Full Breakdown of A Network Cloud Bill of Material (BoM)

Ordinal	Role	Item Name	Description	Manufacturer	SKU	Quantity
1	NCP	NCP-40C (AC Power)	2RU-40x100G QSFP Service Ports + 13x400G QSFPDD Fabric Ports, [Jericho2 with HBM Deep buffer, OP2 KBP, 1 128G SSD, AST2400 Dual Flash BMC, AC PSU FRU, C20 to C19 Power cord].	UfiSpace	S9700-53DX-JB4 AC	24
2	NCF	NCF-48CD (AC Power)	48x400G QSFPDD 2RU Fabric Element based on [2xRamon 1x 128G SSD, AST2400 Dual Flash BMC, AC PSU FRU, C20 to C19 Power cord]	UfiSpace	S9705-48D-480 AC	6
3	NCC	NCC DL360 (AC Power)	HPE DL360 Gen 10	HPE	868703-B21-DL360	2
4	NCM	NCM (AC Power)	48-Port 10G SFP+ with 6x100G QSFP28 uplinks, ONIE software installer, Broadcom Qumran 800Gbps with 6GB packet buffer, ntel Xeon® Processor D-1518 CPU	Edge-Core	5916-54XL-OT-AC-F	2
5	Fabric-Cables	Fabric 400G Ports (NCP<->NCF)	400G AOC QSFP-DD	Innolight	C-DQ8FNM005-N00	144
5	Fabric-Cables	Fabric 400G Ports (NCP<->NCF)	HiWire CLOS AEC, 2.5m, 400G	Credo	CAC425321D1D-B0-HW	144
6	Mgmt-Cables	Cable RJ45 CAT6/7 Orange for console Spare	Cable RJ45 CAT6/7 for console		DN-RJ030O	30
7	Mgmt-Cables	Cable RJ45 Blue Management Servers + iLO	Cable RJ45 Blue Management Servers + iLO		DN-RJ030B	4
8	Mgmt-Cables	Cable RJ45 CAT6/7 (NCP/NCF <-> NCM)	Cable RJ45 CAT6/7		DN-RJ030G	30
9	Mgmt-Cables	Cable RJ45 CAT6/7 0.5m (NCM <-> NCM 17)	Cable RJ45 CAT6/7		DN-RJ005G	2
10	Mgmt-Cables	Fiber cable SM (NCC/NCM <-> NCM)	Fiber cable LC-LC SM		DN-FC-LC-SM030	6
11	Mgmt-Cables	Fiber cable MM (NCP/F <-> NCM)	Fiber cable LC-LC MM		DN-FC-LC-MM030	60
12	Mgmt-XCVRs	Finisar 10G SR SFP+ 200m for ctrl	OOB NCP/NCF<->NCM	Finisar	FTLX8572D3BCL	120
13	Mgmt-XCVRs	Colorchip 100G CWDM4 QSFP-28 2km, ctrl & NIF	OOB NCC/NCM<->NCM	Colorchip	C100-Q020CWDM402D	12
14	Mgmt-XCVRs	Avago Copper SFP 1G	100m, for ipmi	Avago	ABCU-5710RZ	32

7. Conclusions and future work

CUPS architecture grants the ability to source and consume data plane independently from the service it runs. This allows the function and the hardware to age separately and alter the longevity of the investment from the shortest element to the longest element in the BoM.

DriveNets puts this writing as a concept paper and not a specific solution recommendation. Preferably, this approach will be further extended by vendors to create functions which can be mounted onto a OCP-DDC and solution providers to propose specific applications based on this circular economy approach and make these available on the OCP marketplace.