SOLUTION BRIEF

DOCSIS EXCELS ON SUPERMICRO SERVERS

Optimized Server Delivers Extreme Performance for Cable Operators

Supermicro Hyper-E Edge Server

TABLE OF CONTENTS

SUPERMICRO

As a global leader in high performance, high efficiency server technology and innovation, we develop and provide end-to-end green computing solutions to the data center, cloud computing, enterprise IT, big data, HPC, and embedded markets. Our Building Block Solutions® approach allows us to provide a broad range of SKUs, and enables us to build and deliver application-optimized solutions based upon your requirements.

Executive Summary

Supermicro and Intel® have jointly demonstrated an industry-leading download data throughput speed of 317 Gbps while consuming 847 Watts of power over a DOCSIS access plant workload on the Supermicro 2U Hyper E platform (SYS-220HE-FTNR). This system is powered by dual 3rd Gen Intel® Xeon® Scalable processors. This achievement underscores Supermicro's commitment to customers and partners that need to deliver industry-leading solutions at a competitive cost. In addition, Supermicro is the first-to-market partner of Intel with servers that are developed and manufactured in San Jose, CA.

The Supermicro hardware platform demonstrates the latest innovation in system design to meet power efficiency and green computing technology. The Supermicro enterprise computing, storage, and networking solutions combine flexible configuration and tool-less modular designs to give the best-in-class performance, accessibility, and serviceability.

Introduction

Traditional cable access networks are usually built around Cable Modern Termination Systems (CMTS) or Converged Cable Access Platform (CCAP), locked into the technology's hardware-bound limitations. The Intel Access Network Dataplanes project will host access network data plane reference implementations optimized for high-performance packet processing on Intel architecture and deployable on the Network Function Virtualization (NFV) platform. The Intel Dataplane function is a complete solution and mimics the DOCSIS data plane. The vCMTS tool reserves CPU cores for the control plane, scheduler, OS, and telemetry function, as shown in Figure 1. It is a simulation tool with no high availability and redundancy, maximizing the throughput on a single server. The 3rd Gen Intel® Xeon® Scalable processor is used for these tests is an Intel® Xeon® 8368 processor, which contains 38 cores and has a base clock rate of 2.4 GHz.

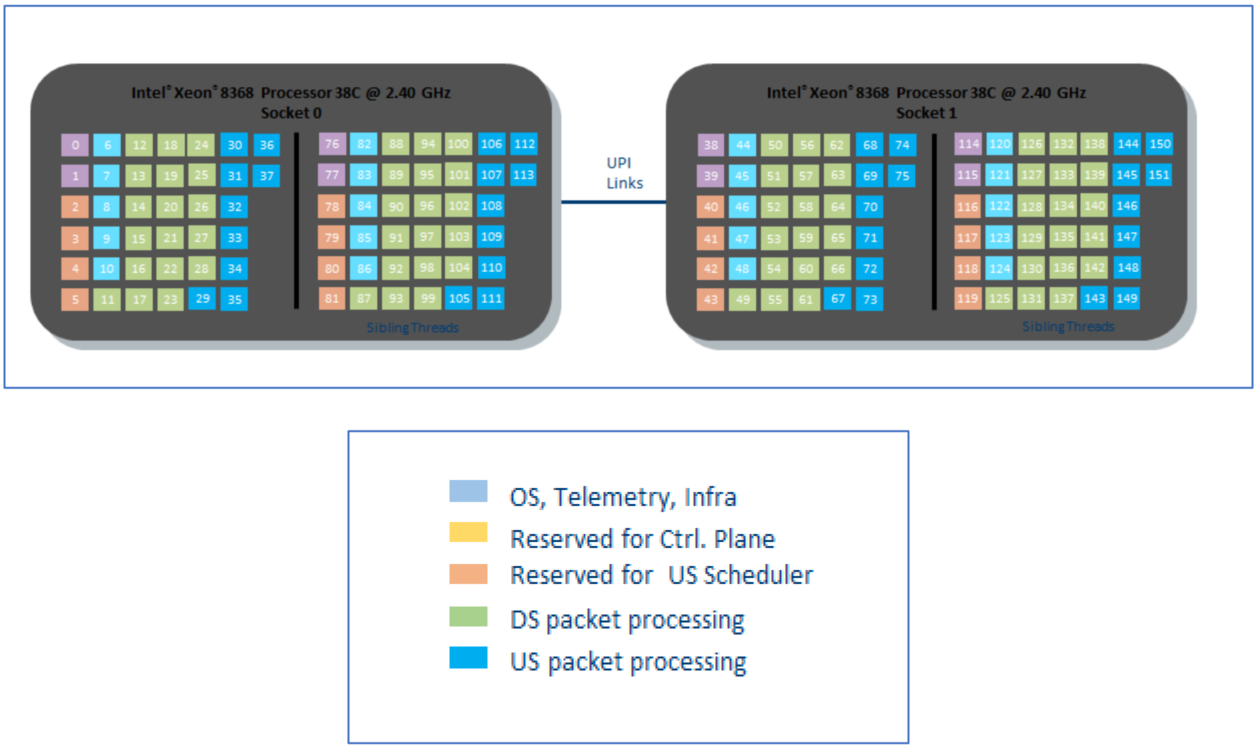

Figure 1 - Reserving Cores for Different Functions

The system used for this development consists of a dual socket Supermicro HYPER-E server SYS-220HE-FTNR. Each of the 38 cores per socket is assigned to execute specific tasks. By assigning specific tasks on a per core basis, the entire system is used, reducing both CAPEX and OPEX costs. The following list describes the number of cores that are dedicated to various tasks.

- 2 cores per socket for OS, telemetry, and infrastructure service
- 4 cores reserved for DOCSIS control plane traffic
- 5 cores reserved for US Scheduler
- 27 cores for DP traffic to run 18 Service Groups
- 18 cores for Downstream Service Group
- 9 cores for Upstream Service Group

© 2021 Copyright Super Micro Computer, Inc. All rights reserved

Service group (SG) configuration:
• 6OFDM channel configuration
• AES, SW only
• IMIX2 traffic simulation

Supermicro Hyper-E Server

The Supermicro Hyper-E is a compact, high-performance, highly configurable edge server designed for a wide range of telco installations. The Hyper-E supports dual 3rd Gen Intel® Xeon® Scalable processors and can be configured with multiple FPGA GPU cards for Edge AI/ML applications such as predictive analytics, helping companies quickly tackle operational issues with predictive tools identify future problems before they become service bottlenecks. In addition, it equips telecom-standard features, including optional NEBS Level 3 compliance, AC and DC power version, and a short-depth front I/O chassis for deployment in central offices and micro data centers.

Figure 2 - Supermicro Hyper-E Server

The Hyper-E (SYS-220HE-FTNR) is an excellent system for vCMTS for the following reasons:
• 2U short-depth design (574mm for 600mm rack)
• Dual socket support for 3rd Gen Intel® Xeon® Scalable processors up to 270W
• Front-access I/O & tool-less design for best serviceability
• Redundant DC or AC power supplies to support high-availability operation
• Up to 8 x PCI-E 4.0 slots, which can support the 2x Intel 200G and the 4x Intel QAT cards used in the vCMTS implementation
• 2x slim AIOM (OCP 3.0 compatible) slots for additional networking connections

The components of the vCMTS solution include:

© 2021 Copyright Super Micro Computer, Inc. All rights reserved

Supermicro Product ID	Description	Qty
SYS-220HE-FTNR	2U Hyper-E server with redundant 2000W AC PSU	2
P4X-ICX8368-SRKH8	Intel® Xeon® 8368 2P 38C/76T 2.4G 57M 11.2GT 270W 4189 D2	4
MEM-DR416L-HL02-ER32	16GB DDR4-3200 ECC Memory	32
HDS-I2T0-SSDSC2KG019T8	Intel D3-S4610 1.92T SATA 6Gb/s 3D TLC 2.5" 7mm 3DWPD Rev.3	2
N/A	Intel 200G Chapman Beach NIC	4
N/A	(optional) Intel 8970 QAT	4
N/A	QSFP28 to QSFP28 cable - Intel PN: DQF8503-4E05-IN	4

The reference architecture for the Dataplane and NFV stack for the Cable Access Network is shown below:

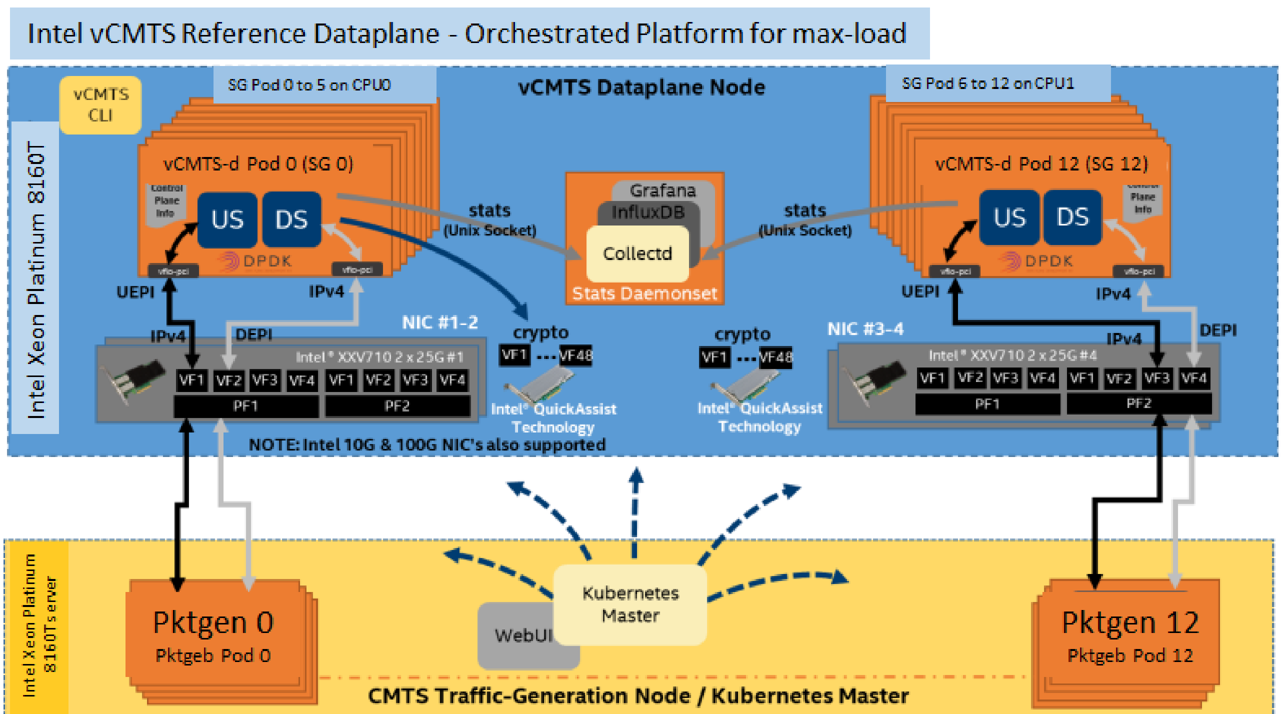

Figure 3 = Reference Architecture

© 2021 Copyright Super Micro Computer, Inc. All rights reserved

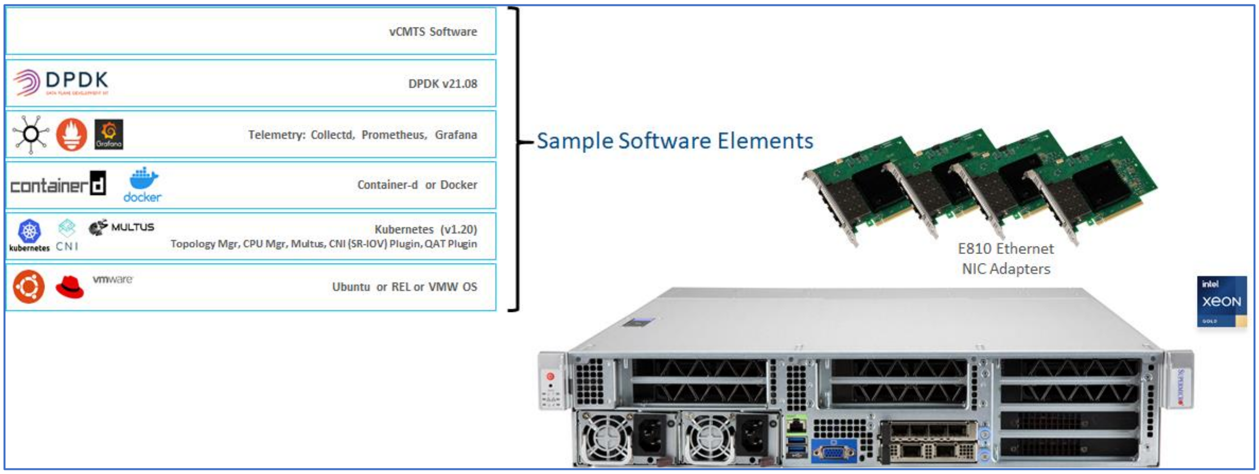

Figure 4 - Software Layers

On the vCMTS data plane node, multiple Docker containers host DPDK-based DOCSIS MAC upstream and downstream data plane processing for individual cable service groups, allowing them to be instantiated and scaled independently. On the CMTS traffic-generation node, Docker containers host DPDK Pktgen-based traffic generation instances, which simulate upstream and downstream traffic into corresponding vCMTS data plane instances. Under Kubernetes, each vCMTS data plane POD represents a service group with separate containers for downstream DOCSIS MAC data plane processing. The DOCSIS control plane is simulated through a JSON file containing subscriber cable-modem information. DOCSIS upstream scheduling is simulated through the generation of UEPI-encapsulated cable-modem DOCSIS stream segments using PCAP files. Telemetry functions run in Docker containers as a Daemonset under Kubernetes. A comprehensive set of vCMTS data plane statistics and platform KPIs are gathered by the open-source Collectd daemon. A Grafana dashboard is provided to visualize these metrics via InfluxDB. The reference platform also contains a power manager, which reduces the power consumption of the CPU cores during quiet network periods. The entire system is orchestrated by Kubernetes, which is an open source container manager and orchestrator. It automates the deployment, scaling, and operational functions associated with application containers. Intel plugins for Kubernetes are used for infrastructure management functions such as CPU core management and assignment of SR-IOV interfaces for NICs.

© 2021 Copyright Super Micro Computer, Inc. All rights reserved

Results

Using the Supermicro Hyper-E Servers (as described above), the I/O speed was measured at 317.2Gbps on a single server, and the system uses 849Watt of power consumption. The figure below compares Upstream (US) and Downstream (DS) performance and resulting performance with 1% and 2% line rates.

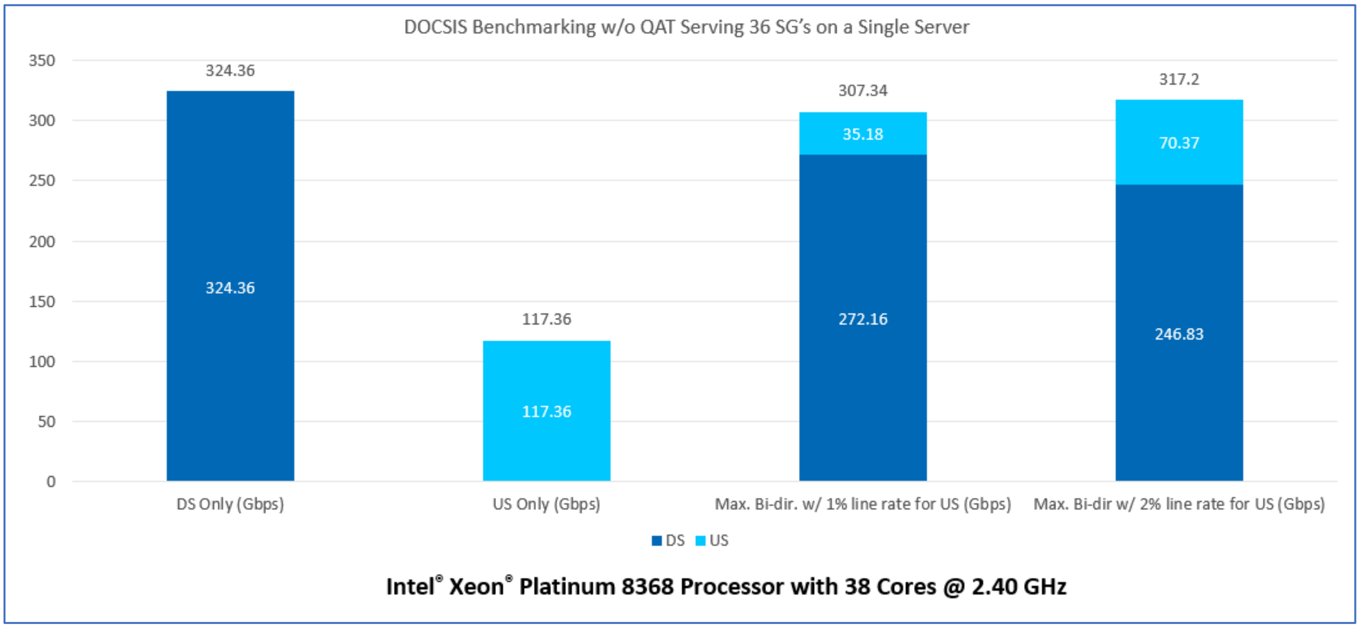

Figure 5 - Performance of DS and US

	Legacy CMTS	Virtual CMTS
		
Rack Space	14 RU	2 RU
Density	6 SG's per RU	18 SG's per RU
Throughput	1 Gbps per SG	~ 8 Gbps per SG
Power	~109w per SG	~31w per SG
HW Cost	~ $200k	Up to 5x Less
	CMTS Chassis + Line Cards	

Figure 6 - Comparing Legacy CMTS with a Virtual CMTS

© 2021 Copyright Super Micro Computer, Inc. All rights reserved

Benefits of a Virtual CMTS:
• Cost effective: Save cost from legacy CMTS deployment
• Reduce footprint: Reduce the footprint by 7x
• Power efficiency: Save power by 3.5X within the same service group
• Best performance: achieved 8Gbps per service group
• Higher Service Group density: Reach 3X more Service Groups per RU space

Conclusion

The Intel vCMTS tool demo aims to help service providers to understand the potential of a virtualized architecture and simulate the result. Supermicro offers a compact and optimized hardware platform that exceeds the Intel vCMTS reference design guide's requirement and delivers world-class performance and efficiency for Multiple System Operators (MSO) CMTS application workloads.

© 2021 Copyright Super Micro Computer, Inc. All rights reserved