Dual-uCPE for High-Availability Retailer Services with Fault Tolerance and Load Balancing

Ying-Dar Lin, Kuan-Yu Ho, Widhi Yahya, Chi-Yu Li, Yuan-Cheng Lai, and Jeans H. Tseng

Abstract

Network function virtualization (NFV) enables the emergence of universal customer premises equipment (uCPE), which hosts various networking and computing services. uCPE supports a software-defined wide area network (SD-WAN) that provides networking services to branch or remote offices, so it has the potential to host more services than CPE and virtual CPE. However, any failure of the uCPE, which hosts all the services, used in a retailer store may result in significant losses due to unrecorded transactions or the inability to conduct transactions. For providing high-availability (HA) of computing and networking services with minimal downtime in the event of failures, this study proposes a dual-uCPE system for retailer services. Operating in a master-slave mode, the dual-uCPE hosts three retailer applications: point of sales (POS), network video recorder (NVR), and Wi-Fi networking. Three-tier HA networking, which includes dual-WAN connections with traffic balancing, virtual routers on dual-uCPE servers, and a Wi-Fi mesh network, is proposed. The HA performance of the dual-uCPE system is evaluated by measuring the failover and failback times under various failure scenarios. The POS and NVR failover (failback) times are found to be 9.12 (7.1) and 8.76 (6.92) seconds, respectively. Moreover, the failover (failback) times for the HA networking service are both 0 seconds following WAN connection failures, 4.42 (0.13) seconds after uCPE server failures, and 40.14 (30.02) seconds after access point (AP) failures. In other words, for all of the failure scenarios, the average downtime of the proposed dual-uCPE is in the order of seconds, rather than tens of minutes as for manual recovery.

Introduction

Internet service providers (ISPs) typically install customer-specific hardware at the customer's premises to provide network services such as set-top box, VPN and WAN connections [1]. However, before configuring the customer premises equipment (CPE), the ISP technician must first analyze the customer's hardware and software requirements. Furthermore, the ISP must dispatch a technician onsite to physically deploy the equipments. Such activities are time-consuming and costly, particularly when the customer subsequently decides to update the network services, necessitating the reconfiguration and deployment of one or more closed-proprietary devices. Since CPE takes the form of closed-proprietary hardware, various hardware must be installed by providers for different services or functions. Thus providing high-availability (HA) for all CPEs may double all equipment.

Virtualized CPE (vCPE) overcomes the mentioned CPE drawbacks by applying the concept of network functions virtualization (NFV) to transform the network services from dedicated hardware-based equipment to software-based functions [2]. Notably, through the virtualization of network functions, multiple functionalities can be executed in a single container or virtual machine (VM) on a commercial off-the-shelf (COTS) x86 server and managed by a virtualized infrastructure manager (VIM) [3].

Ying-Dar Lin, Kuan-Yu Ho, and Chi-Yu Li are with National Yang Ming Chiao Tung University, Taiwan.

Widhi Yahya (corresponding author) is with Universitas Brawijaya, Indonesia and with National Yang Ming Chiao Tung University, Taiwan.

Yuan-Cheng Lai is with National Taiwan University of Science and Technology, Taiwan.

Jeans H. Tseng is with Lanner Inc.

Digital Object Identifier: 10.1109/IOTM.001.2200234

vCPE running on a COTS x86 server could host computing, storage, and networking for enterprise edge computing, known as universal CPE (uCPE). Placing uCPE in some branch offices spread across geographical areas and managed by a centralized VIM enables software-defined wide area networking (SD-WAN) [4]. uCPE with SD-WAN simplifies edge networking management to provide optimal computing and networking performance.

This paper discusses uCPE development for retailers with three crucial retail services [5] such as network video recorder (NVR), point of sale (POS), and Wi-Fi service. This retailer's uCPE provides computing, networking, and storage services on the same hardware. As hosting more services than vCPE and CPE, uCPE incurs an increased risk of hardware failure due to its heavy load. However, retailers are typically required to operate 24/7, and any failure of the uCPE server may thus result in significant losses due to unrecorded transactions or the inability to conduct transactions. Consequently, HA is necessary to ensure minimal service disruption or downtime, as stipulated in the Service Level Agreement (SLA).

HA is the ability of systems to minimize service downtime by providing redundant system entities, including hardware and software. This redundant entities operates in one of four different modes: cold-standby, warm-standby, hot-standby, and active-active. The first two modes are incompatible with real-time systems since they turn off one of the redundant entities during normal operation. By contrast, in the hot-standby mode, the spare entities remains in an idle/sleep state until required, while in the active-active mode, all of the entities remains active and performs processing all of the time. However, because active-active mode initially employs all active servers to serve users, the failure of one of the servers leads to a significant deterioration of the overall system performance [6]. Thus, in

<table>
<tr><th rowspan="2">Paper</th><th rowspan="2">Target device</th><th colspan="2">Problem</th><th colspan="2">Services</th><th rowspan="2">Multi-ISP support</th><th rowspan="2">NF location</th><th rowspan="2">Objective</th><th rowspan="2">Solutions</th><th rowspan="2">Virtualization</th></tr>
<tr><th>High availability</th><th>Load balancing</th><th>Networking</th><th>Computing</th></tr>
<tr><td>[7]</td><td rowspan="8">vCPE</td><td rowspan="9">X</td><td rowspan="6">X</td><td rowspan="10">V</td><td rowspan="9">X</td><td>V</td><td>CS (residential)</td><td>Full virtualized home gateway</td><td>NFV-based vCPE</td><td>JVM</td></tr>
<tr><td>[8]</td><td>X</td><td>CS (residential)</td><td>Develop vRGW and vSTB</td><td>NFV-based vCPE</td><td>VM</td></tr>
<tr><td>[9]</td><td>X</td><td>CS (residential), cloud</td><td>Minimize network cost</td><td>Heuristic VNF placement algorithm</td><td>N/A</td></tr>
<tr><td>[10]</td><td>X</td><td>CS (residential)</td><td>PoC vRGW</td><td>SDN and NFV-based vCPE</td><td>VM</td></tr>
<tr><td>[11]</td><td>X</td><td>Cloud, edge</td><td>SDN switch based vCPE</td><td>Multiflow table</td><td>VM</td></tr>
<tr><td>[14]</td><td>X</td><td>CS, cloud/ edge</td><td>Minimize mgmt. complexity</td><td>NFV-based vCPE</td><td>VM</td></tr>
<tr><td>[12]</td><td rowspan="2">X</td><td>X</td><td>Cloud</td><td>Minimize manag. complexity</td><td>SDN-based vCPE</td><td>N/A</td></tr>
<tr><td>[13]</td><td>X</td><td>Edge</td><td>Dynamically instantiated CPE</td><td>vCPE with SDN and NFV</td><td>Container</td></tr>
<tr><td>[15]</td><td rowspan="2">uCPE</td><td>X</td><td>V</td><td>CS (IoT field)</td><td>IoT networking/ interoperability</td><td>uCPE-based IoT gateway</td><td>N/A</td></tr>
<tr><td>Ours</td><td>V</td><td>V</td><td>V</td><td>V</td><td>CS (retail stores)</td><td>Minimize down-time and maxi-mize througput</td><td>Dual-uCPE with dual-WAN</td><td>Container</td></tr>
</table>

TABLE 1. Comparison on CPE developments.

this study, the hot-standby mode is adopted as a more cost-effective and practical solution for HA uCPE.

The development of vCPE systems has attracted considerable attention in the literature [7–14]. These studies have focused primarily on the use of vCPE for the provision of triple-play services (voice, video, and data) in residential networks. The studies in [7, 8, 10, 14 focused on the application of NFV technology to realize CPE, while that in [9] attempted to optimize the virtual network function (VNF) placement for edge-cloud sites. The authors in [11–13] developed vCPE systems by leveraging software-defined networking (SDN) technology. Zhou *et al.* [15] developed a uCPE to address interoperability problems in IoT systems. However, none of these studies considered the need for HA CPE since disruption to the considered residential CPE devices affected only the family members concerned. By contrast, the present study considers the case of HA uCPE for a retail store, wherein the uCPE hosts computing and networking services to support the business processes, and any disruptions may result in serious business harm. Since computing and networking services have different requirements and characteristics, developing a HA uCPE is more challenging than developing a traditional server with only computing services.

This study proposes a dual-uCPE architecture. To the best of our knowledge, the proposed dual-uCPE architecture is the first CPE architecture to provide HA for both computing and networking services while requiring minimal hardware by hosting both computing and networking services on the same hardware. To achieve the former, the two servers in the dual-uCPE system are configured with the virtual router redundancy protocol (VRRP) in the hot-standby mode, which entails the use of one of the servers as a master and the other as a slave. Each server hosts the required POS, NVR, and Wi-Fi services and generates files and data locally. Consequently, a storage synchronization problem may occur. In the proposed system, this problem is addressed through the use of database master-master replication and a network file system. For providing the HA of networking services, the dual-uCPE adopts a redundancy approach and employs two WAN links and two gateway links per uCPE and implements the Wi-Fi network using four access points (APs) connected in a mesh topology. In the two WAN links, a weight-based routing scheme is employed to increase the throughput from the LAN to the WAN. Meanwhile, the two gateway links employ multiple VRPP instances of Keepalived opensource software (https://www.keepalived.org/) with DHCP conditional pool to assign two virtual IP addresses as network gateways to user equipments (UEs). Moreover, the four Wi-Fi APs are managed by a prplMesh (https://prplfoundation.org/prplmesh/) controller to construct a wireless mesh network with enhanced robustness againts AP failure.

The rest of this work is organized as follows: Section II presents related work for vCPE developments; Section III describes dual-uCPE architecture includes problem statements and solutions; Section IV details the system implementation; The experiment results are presented in Section V; Section VI concludes this work.

Related Work

Table 1 summarizes the previous work on vCPE and uCPE development. Herbaut *et al.* [7] proposed a Surrogate VNF approach in which a bundle of OSGi services were used as a regular home gateway alongside VNFs. The proposed system thus facilitated migration from modular home gateway to fully NFV-based home gateway. Bronstein and Shraga [8] investigated the implementation and security issues associated with vCPE deployment. Suksomboon *et al.* [9] examined the worthiness of shifting CPE functionalities to a data center using a vCPE cost optimization method based on a connected graph approach, in which the vCPE functions served as vertices and the communication costs between them served as edges. Proença *et al.* [10] presented a proof-of-concept prototype for a NFV-based residential gateway (vRGW) based on SDN and NVF functions. Huang *et al.* [11] proposed a vCPE framework for enabling the deployment of NFVs at the edge, in which the VNFs were implemented using multiple flow tables in an SDN switch. Zhou *et al.* [15] proposed a uCPE for an IoT gateway that supported ZigBee, Bluetooth Low Energy, and Wi-Fi protocols and connected to multi-cloud providers. Ericsson [12]

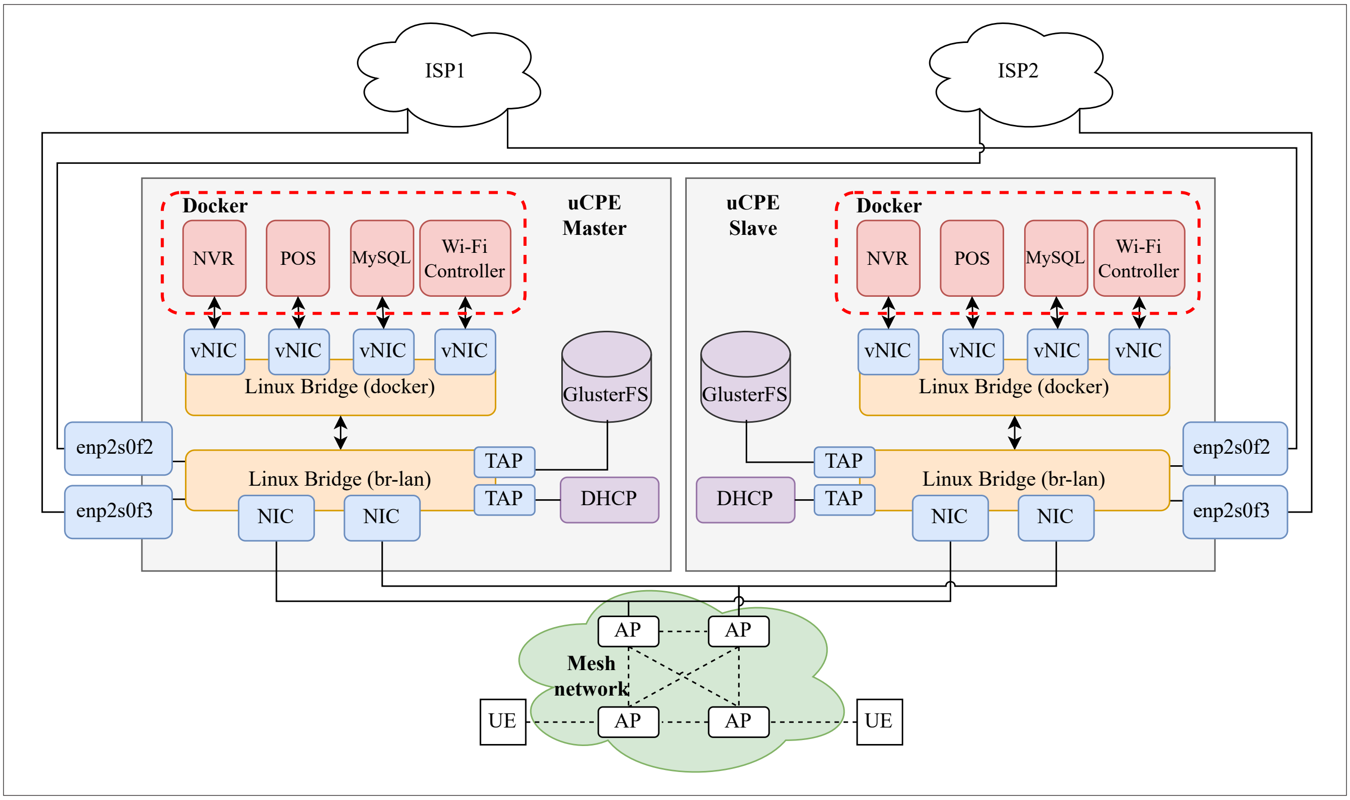

FIGURE 1. Design of dual-uCPE architecture.

proposed an SDN-based vCPE that enabled service providers to differentiate their services through the rapid launch of innovative and personalized services. Meneses *et al.* [13] considered the problem of dynamic SDN-based vCPE orchestration at the edge level and utilized VNF migration to achieve cluster-level load balancing. Proença *et al.* [14] presented a prototype vCPE environment with two service models for residential and business scenarios, respectively. In the residential model, the VNFs were all deployed on the operator's infrastructure. By contrast, in the business scenario, the VNFs were deployed both at the customer and at edge/cloud sites.

However, none of the prior works in Table 1 takes the HA requirement of the CPE or the deployment of computational services on the CPE into account. In most cases, their primary objective is simply to provide triple-play services to residential customers. Furthermore, due to the limited resources available to the CPE, some of them deploy the VNFs at the edge and/or the cloud [9, 11–13]. However, maintaining the VNFs on the customer premises can improve their protection against security threats [15]. Accordingly, Ericsson [12] considered load balancing customer the traffic to the cloud-based vCPE, while Meneses *et al.* [13] performed load balancing in optimizing the placement of the VNFs at the edge. None of them considered load balancing traffic across two or more WAN links, which would also improve customer internet connection availability by utilizing multiple ISPs.

Dual-uCPE Architecture

Figure 1 shows the architecture of the dual-uCPE proposed in the present study for providing a retailer store with three primary services: POS, NVR, and Wi-Fi. The uCPE merges both computing and networking services, and thus the problem statement comprises two elements, namely providing HA computing services and providing HA networking services.

Design Issues

High-Availablity Computing Services: Failure to access retailer apps like POS or NVR results in serious business harm because these apps provide transaction and security services, respectively. Given the deployment of dual-uCPE servers, one server should be chosen as the master to handle the arrival tasks, while the other should be nominated as a backup, which resides by default in an idle state, but replaces the master uCPE in the event that it fails (referred to as failover). The backup uCPE is deliberately not powered off in order to minimize the failover time. When the master server is subsequently recovered from the failure state, the backup uCPE returns the master role and reverts to an idle state (referred to as failback). Notably, the failback and failover processes must be transparent to the users. That is, they should not be required to modify the IP address destination for the specified services.

Each uCPE in the dual-uCPE system runs POS, NVR, and all the network functions required to support the Wi-Fi service, including the Wi-Fi controller, gateway, and DHCP server. The computing and networking services are composed of a collection of application files and databases. Any changes made to the files or databases in one uCPE must be replicated in the other. In particular, given dual-uCPE boxes with POS, NVR, and network services, the aim is to maintain identical service files and databases on both uCPE servers during all the uCPE server operations, including failover and failback. By doing so, clients attempting to access these services receive identical results regardless of which server they connect to.

High-Availability Networking Services: Retail stores are typically distributed across a wide geographical area. Consequently, a WAN or Internet connection is required to connect them to corporate headquarters. The Internet also connects the local store databases to a cloud-based central database. Thus, to ensure file and database synchronization, a HA Internet connection is required. In particular, given dual-uCPE servers, each connected to multiple ISPs, the aim is to balance the uplink and downlink traffic of the ISPs in such a way as to maximize the network throughput while maintaining high-availability.

Wi-Fi connects retailer equipment and customers to the Internet, improving shopping experiences. The Wi-Fi service is thus offered by installing APs at various locations across the retail store. That is, given dual-uCPE servers and multiple APs, the dual-uCPE servers should act as network gateways of the retail store local network, and the aim of the design problem is thus to realize a HA topology design and a dynamic gateway assignment mechanism to distribute the clients across the multiple gateways.

Solution Ideas

Computed NVR, POS, and WiFi with High Availability: The POS, NVR, and network functions supporting the Wi-Fi network are all containerized on each of the dual-uCPE servers. As shown in Fig. 1, each container has a virtual Network Interface Card (vNIC) that is bridged to a physical interface with a physical IP address. As shown in Fig. 2, at any given time, VRRP with VRID 1, which runs on both uCPE servers in a master-backup mode, assigns a virtual IP (VIP) to the physical IP address of one of the uCPE server with the highest priority (the master uCPE). All of the retailer applications are then run on this uCPE and halted on the backup uCPE.

A VIP address is assigned as the user interface for the retailer's applications and databases. Client requests addressed to that VIP must therefore be redirected to the physical IP address of the master uCPE. The master uCPE periodically sends an advertisement to the all uCPEs. If the backup uCPE does not receive three consecutive advertisements from the master uCPE, it inherits the VIP and changes its role to master uCPE. When the master uCPE is subsequently recovered, it sends an advertisement with the highest priority once again. On receiving this advertisement, the slave immediately returns the VIP to the master and resumes its role as the backup uCPE. Since the customers access the retail services over the VIP, the failover and failback processes are both transparent to them.

The uCPE computing services contain files and databases. To maintain consistency in the dual-uCPE system, all of the changes to a file or database on one uCPE server must be replicated to the other such that, if one of the servers fails, or is taken offline for maintenance, the files and databases can still be accessed on the other server. For the dual-uCPE proposed in the present study, database consistency is maintained using MySQL master-master replication, while file consistency is achieved using GlusterFS volume replication.

Load Balancing and Dynamic Gateway Assignment for Wireless Mesh Network: As shown in Fig. 3a, the dual-uCPE system utilizes two ISPs to achieve HA Internet connectivity. However, rather than employing a master-backup mode, the egress link employs load balancing, whereby Iptables are configured to operate in a multihomed mode with a round-robin scheduler. Notably, the load balancing mode not only provides HA, but also maximizes the outgoing network throughput.

To achieve HA Wi-Fi network, APs are configured as a mesh network that interconnects APs over multiple links. The dual-uCPE system uses two VIPs as mesh network gateways, and thus faces a multi-gateway issue in the local network of the retailer store. In the proposed dual-uCPE design, this issue is addressed by two mechanisms: multiple VRRP groups and multiple DHCP-server load balancing. The multiple VRRP groups, denoted as VRID 1 and VRID 2, respectively, ensure that the network gateways given to the ports of the two uCPE servers are fail-safe. Thus, if one uCPE server fails, the other server backups the network gateway, as shown in Fig. 3b. (Note that the associated mechanism is similar to that used to achieve the HA computing services.) Meanwhile, multiple DHCP-server load balancing is used to balance the number of devices connected to the dual-uCPE gateways by assigning a distinct virtual gateway to a distinct client machine, as shown in Fig. 2b. Since the IP and gateway configurations are both handled by the DHCP server, the client is unaware of the gateway being used (i.e., the load balancing process is transparent to the user).

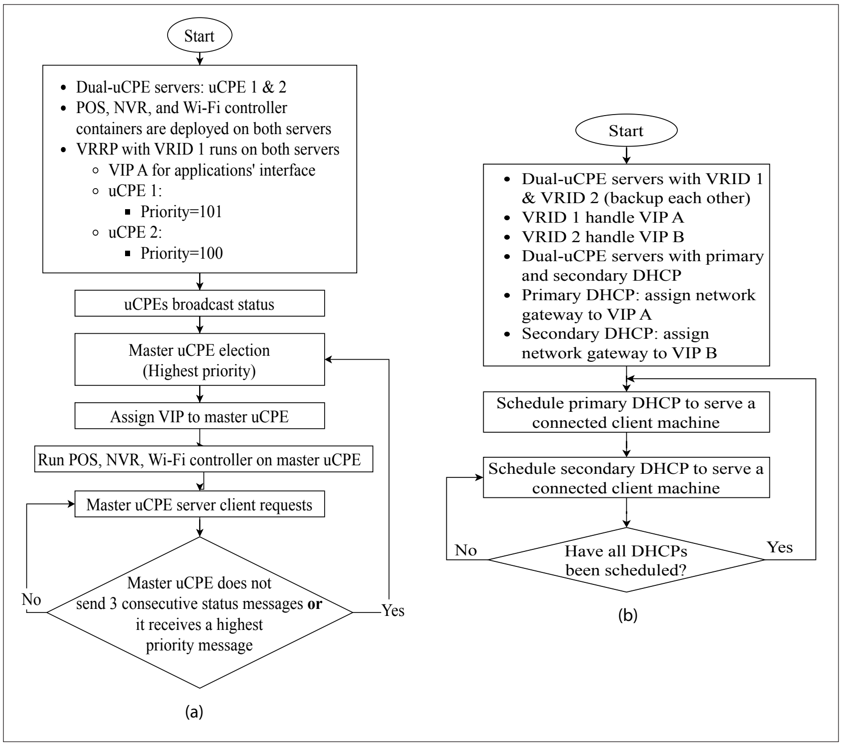

FIGURE 2. HA computing and networking: a) HA computing; b) dual-DHCP scheduling for HA networking.

Implementation

The HA dual-uCPE was built using two Lanner NCA-1515A boxes. Each uCPE was equipped with an Intel Atom C3000 CPU, 16GB RAM, and 6x GbE RJ45 ports. Moreover, each uCPE was linked to the Internet by X and Y ISPs and to a LAN via multiple links. IP address ranges 192.168.1.0/24 were used on the LAN, where the LAN comprised four APs organized in a mesh topology. Docker was used to containerize applications and networking services since it has a smaller footprint than VMs and is supported by automation deployment tools like Kubernetes (www.docker.com). This section presents the implementation details of the HA computing and networking services, respectively.

High-Availability Computing Services Implementation

Dual-uCPE servers with VRRP: This project implemented Keepalived 1.3.9 on top of Ubuntu 18.04. As shown in Fig. 3a, two VRRP instances, VI_1 and VI_2, were installed on each uCPE server. For each server, the VRRP instances were assigned to virtual router IDs (VRIDs) 1 and 2, respectively, where each VRID contained a pool of physical devices with a mutual backup capability.

Each VRRP instance was responsible for assigning a VIP to a physical IP address. In particular, VI_1 of uCPE 1, VIP 192.168.1.20, was assigned to the master physical IP 192.168.1.10, while VI_1 of uCPE 2, 192.168.1.11, was assigned to the backup physical IP. Similarly, VI_2 of uCPE 2, VIP 192.168.1.21, was assigned to the master physical IP 192.168.1.11, while VI_2 of uCPE 1, VIP 192.168.1.10, was assigned to the backup physical IP.

In other words, each VIP held two physical IPs for the master and backup, respectively.

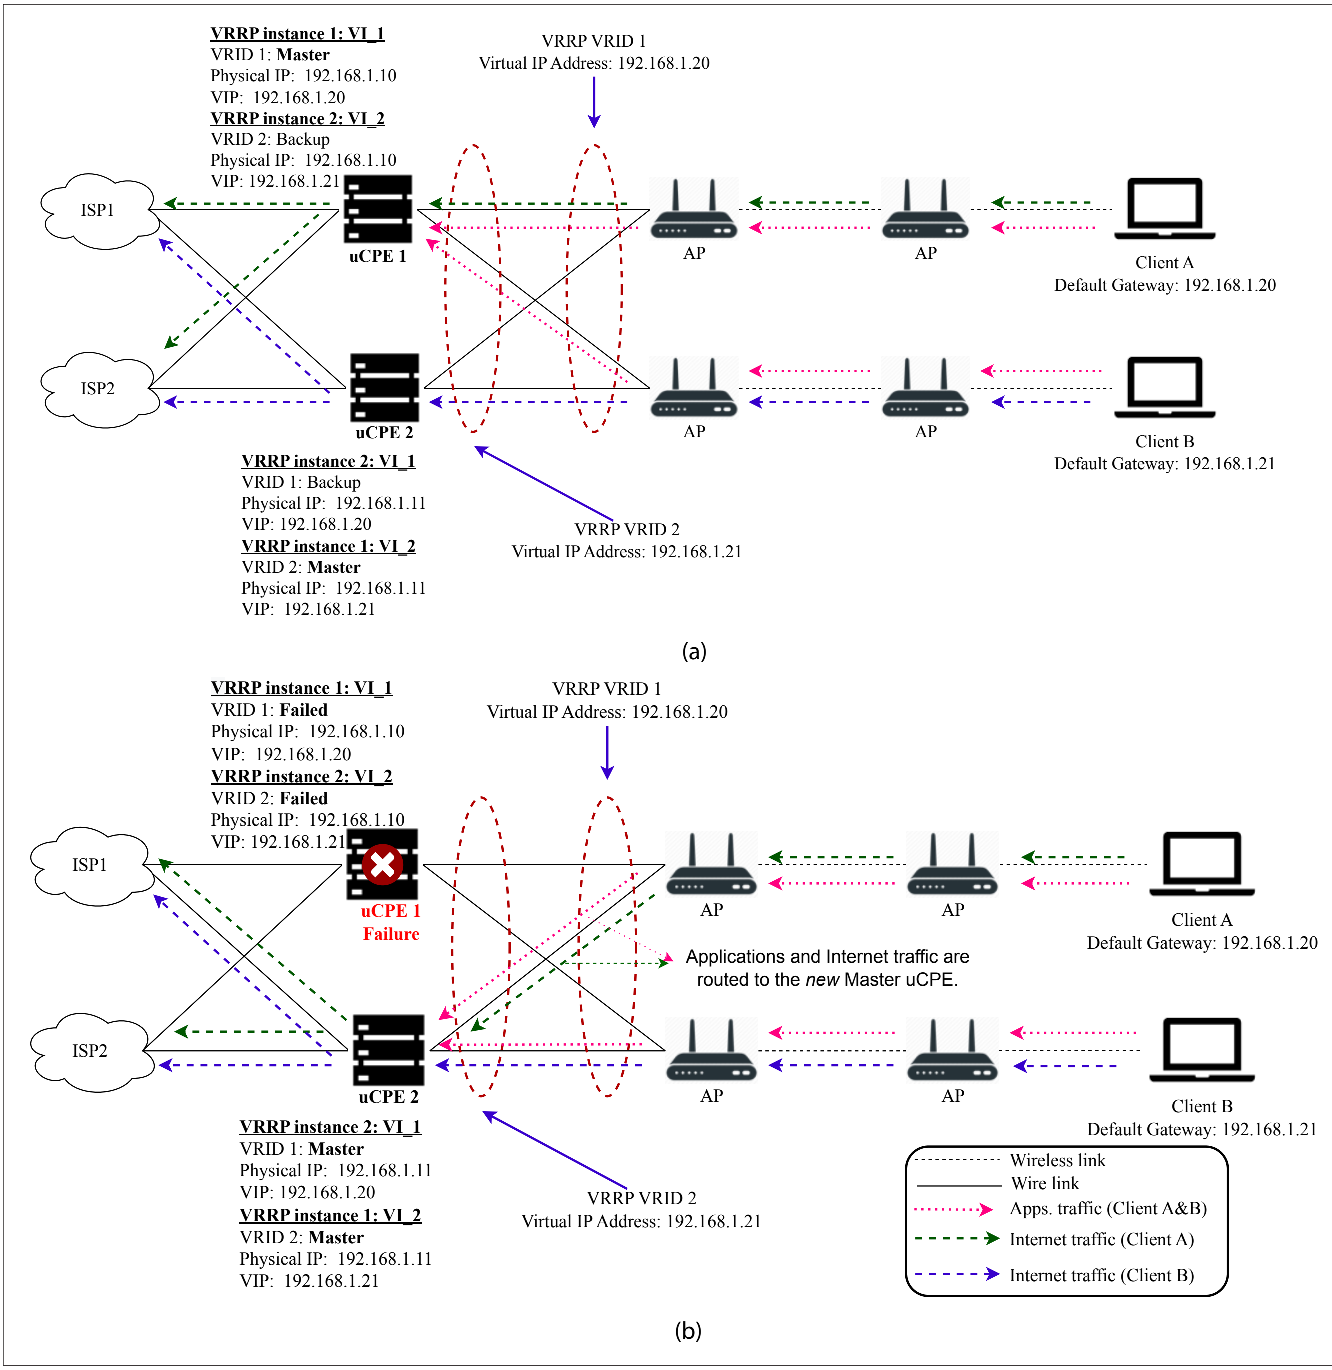

FIGURE 3. Dual-uCPE system testbed: a) normal traffic direction; b) traffic direction during single uCPE server failure.

File and Database Synchronization: The files and databases for the retailer services were stored on a volume on a Linux system and were controlled by the file system. During failover and failback, files on one of the uCPE servers may be modified. Thus, without file replication, inconsistencies in the files will occur if the uCPE recovered from the failure state subsequently serves clients. To avoid this problem, a replicated volume named "datavol" was created and mounted on both uCPE servers. Database synchronization was then implemented using the MariaDB master/master replication scheme. In particular, MariaDB (https://mariadb.org/) was run on both uCPEs, and the running application accessed the database via the VIP connected to the database system. In the case of a single uCPE failure, the application thus retained the ability to access the database since the VIP was automatically reallocated to the running uCPE server by Keepalived.

High-Availability Networking Services Implementation

As shown in Fig. 3a, the dual-uCPE system provided HA Internet connections via two ISPs. A HA LAN service was additionally provided though a Wireless Distributed System (WDS) and fault tolerant DHCP servers.

High-Availability and Load Balancing of Internet Connections: Each uCPE server connected to the LAN through two physical interfaces each of which was connected to an AP. The configured br-lan on each uCPE server interface routed traffic from the LAN to one of the defined physical interfaces (enp2s0f2 and enp2s0f3), where each interface was connected to an ISP. Outgoing traffic was balanced between the two interfaces by assigning them the same weight and using a simple round-robin scheme. In addition, network address translation (NAT) was configured at each of the two interfaces.

High-Availability and Load Balancing of LAN connections: To provide HA, a LAN testbed comprised of four APs running the prplWrt operating system was organized as a mesh topology. Two of the APs were connected via wired connections to both uCPE servers and served as the main APs providing Internet connectivity. The other two APs served as relays, connecting the primary APs and clients through a multiple wireless links.

High-availability was obtained by configuring the two primary APs connected to the dual-uCPE as network gateways. Each of the dual-uCPE servers ran the DHCP server (https://www.isc.org/dhcp/) in the master-slave mode in order to perform dynamic IP configuration during failover and failback. Moreover, both DHCP servers peered with each other and configured the client IP address in the range 192.168.1.100-192.168.1.254. To balance the client-to-Internet traffic, the master and slave DHCP servers assigned distinct IP gateways, 192.168.1.20 and 192.168.1.21, to the client machines.

Experimental Results

This section starts with performance evaluation parameter settings and brief evaluation scenarios for dual-uCPE architectures. Experimental data are discussed last.

Parameter Settings

The parameter settings are shown in Table 2. The Dual-uCPE servers, hosting POS (https://opensourcepos.org/) and NVR (https://shinobi.video/), were connected to a local mesh network and two ISPs. Clients are connected to the mesh network and and generated requests to both the computing services and public websites on the Internet.

Category	Parameter	Value
Dual-uCPE	Master uCPE	1
	Slave uCPE	1
Mesh network	Access point	4
Dual-uCPE hardware	Master uCPE	NCA-1515A
	Slave uCPE	NCA-1515A
Mesh network hardware	Access point	GL-B1350
	Wireless link bandwidth	300 Mb/s
	Wire link bandwidth	930 Mb/s
Internet access	ISP 1 bandwidth	40 Mb/s
	ISP 2 bandwidth	40 Mb/s
Experiment	Time duration	2 minutes
	Generated request	15 req/s

TABLE 2. Parameter settings.

Evaluation Scenarios

Three evaluation scenarios were considered, where one scenario addressed the HA computing service issue, while the other two addressed the HA networking issue.

Computing Services Failover and Failback: For two minutes, clients were connected to an AP and generated burst requests using Jmeter to the NVR and POS applications. One of the master uCPE servers was then shut down and the failover time was determined by calculating the time interval between successful requests. The failure recovery mechanism was evaluated by restoring the down uCPE to the system, prompting the client requests to be automatically forwarded to the master uCPE. The time interval between successful requests was then taken as the failback time. The consistency of the files and databases was also investigated during the failover and failback scenarios.

Networking Failover and Failback: To evaluate the HA networking performance of the dual-uCPE system, three networking failure scenarios were considered. In the first scenario, networking failure was applied to one of the ISP links. Due to the load balancing configuration of the proposed dual-ISP arrangement, the system performance is degraded if one of the uCPE fails. Accordingly, in the first failure scenario, the performance with and without load balancing was compared. In the second scenario, failure was applied to one of the uCPE servers. The dual-uCPE hosts virtual routers to route traffic from the LAN to the Internet, and vice versa. Thus, if one of the uCPE fails, the Internet connection is disrupted. In the third scenario, the networking failure was applied to one of the mesh APs. Burst traffic was generated from one of the clients to a public website on the Internet. The failover and failback times were then calculated based on the interval between successful requests during the failure events.

Network Load Balancing: The dual-uCPE has a dual-WAN and dual-gateway structure to provide not only HA, but also an improved network throughput. In the third evaluation scenario, network throughput comparisons were made between four systems:

1. Dual-uCPE with dual-WAN
2. Dual-uCPE with single WAN
3. Single uCPE with dual-WAN
4. Single uCPE with single WAN

In the comparison process, traffic was generated from UEs attached to the APs to two servers on the Internet using iPerf3.

Results

Computing Services Failback and Failover: A batch of 15 requests per second was sent to the POS and NVR programs, and to a public website on the Internet, for two minutes in order to test the virtual router. During the experiment, the master uCPE was turned off and then back on to trigger the failover and failback mechanisms, respectively. To handle the failure, the VRRP assigned all the VIPs to the slave, including that used as the network gateway. As shown in Fig. 4a, NVR spent 8.76 and 6.92 seconds for failover and failback, respectively, while POS spent 9.12 and 7.1 seconds. The obtained applications' failover and failback times include network failover and failback times of 4.42 and 0.13 seconds, respectively, which is the time utilized for VIP (Apps interface) reassignment. Excluding network fail-recovery time, the failback times of both applications were greater than the failover times since, when the failed server was being recovered, client requests were still routed to the slave uCPE, where they were subsequently dropped due to ongoing reassignment of the VIPs to the master uCPE. The database synchronization was successfully tested by adding a user prior to a fail conditions and then revoking the new user's login after this fail.

POS and NVR benchmarking was performed by sending burst requests to the application websites. As illustrated in Fig. 4c, NVR managed 286 requests/s. However, POS was only able to handle 36 requests/s due to its need to query a greater number of database tables.

Networking Failover and Failback: Burst requests were sent to a public webpage on the Internet. During the request generation process, three network failure scenarios were instigated involving failures on the outgoing Internet link, uCPE server, and one of the mesh APs. Figure 4b shows the failover and failback times for each scenario. The outgoing Internet link failure scenario did not result in a sending packet timeout since the rerouting process in the uCPE server is extremely fast. However, uCPE server failure scenario required VIP reassignment, and incurred failover and failback times of 4.42 and 0.13 seconds, respectively. The network failover time is higher than the failback time because in failover, the slave uCPE takes over the master role (VIP) only when it does not receive three consecutive advertisements. By contrast, the failback mechanism is initiated as soon as the slave receives the first advertisement of the recovered master with a higher router priority than its own. In AP failure scenario, when one of the mesh APs attached to

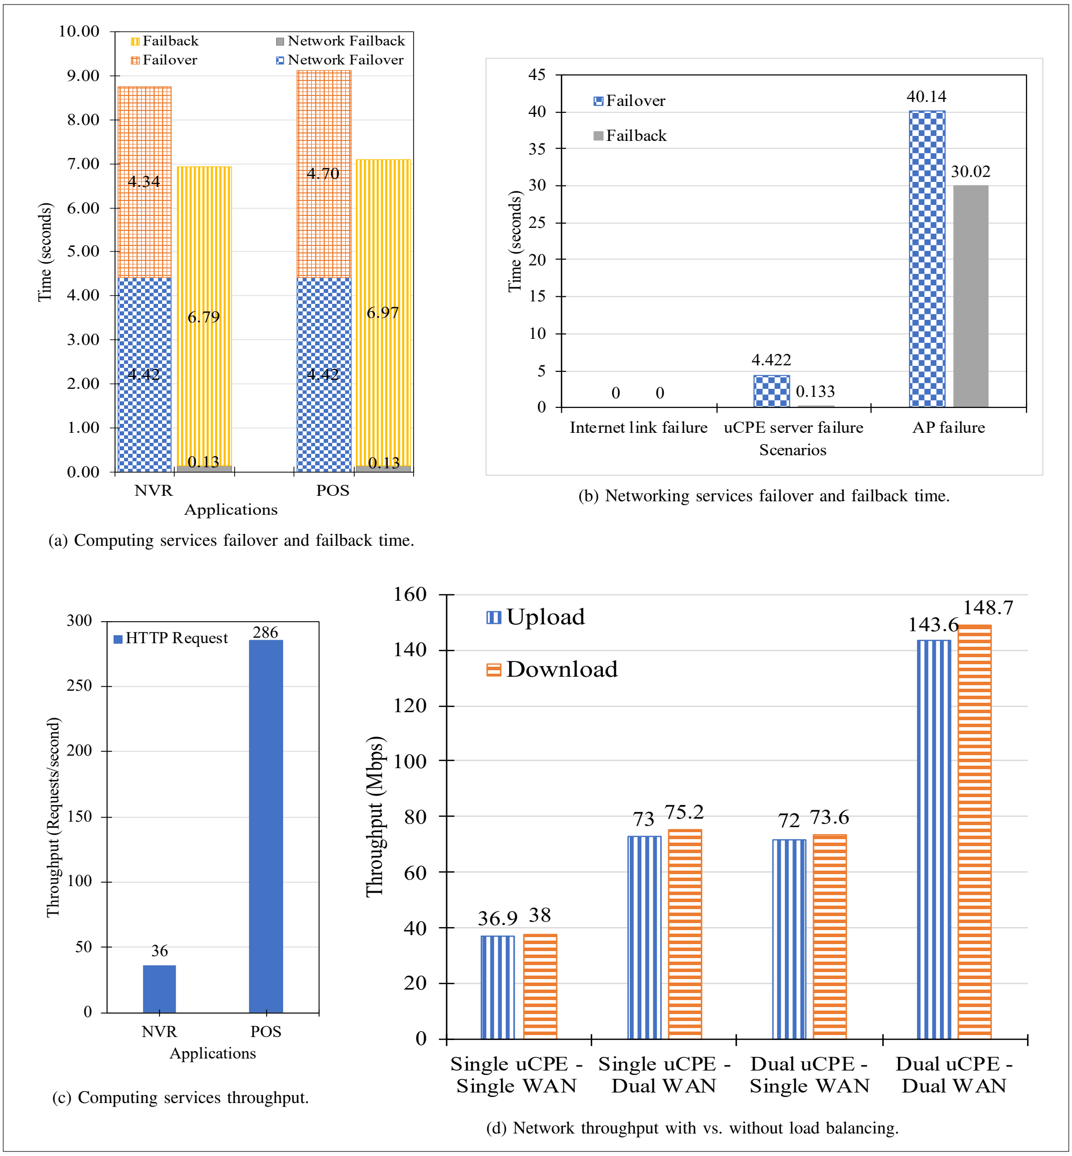

(a) Computing services failover and failback time.

(b) Networking services failover and failback time.

(c) Computing services throughput.

(d) Network throughput with vs. without load balancing.

FIGURE 4. Dual-uCPE system performances: a) networking services failover and failback time; b) computing services throughput; c) network throughput with vs. without load balancing.

a gateway fails, the mesh network must discover and reconnect the WDS to the backup AP and then reconstruct the network. Consequently, the failover and failback mechanisms have relatively longer times of 40.14 and 30.02 seconds, respectively.

Network Load Balancing: HA networking was achieved by applying the load balancing mode in the dual-WAN links and dual-LAN gateway links. Fig. 4d compares the throughputs of the four considered uCPE systems. As shown, the proposed dual-uCPE system with dual-WAN achieves the highest upload and download throughputs of 143.6 and 148.7 Mb/s, respectively, due to its use of four WAN links with two links per uCPE. The dual-uCPE system with a single WAN represents the case where one of the WANs links in the designed dual-uCPE system fails. As shown, the system thus achieves lower upload and download throughputs of 72 and 73.6 Mb/s, respectively. For the single uCPE system with dual-WAN, the upload and download throughputs are 73 and 75.2 Mb/s, respectively. In other words, the throughput performance is similar to that of the dual-uCPE with a single WAN since the number of WAN links is equal to two in both cases. Finally, the single uCPE with single WAN system has the lowest throughput of the four systems with upload and download throughputs of 36.9 and 38 Mb/s, respectively, due to its use of a single WAN link.

Conclusions and Future Work

This study has proposed and evaluated a dual-uCPE architecture consisting of two servers, two WANs, and two gateways for high-availability (HA) computing and networking services with the objective of minimizing the downtime during failure or routine maintenance. The evaluation results have shown that the failover and failback times for both the POS and NVR computing services are on the order of unit seconds, which is significantly faster than manual recovery, which can take tens of minutes. According to the results of a three-tier HA networking evaluation, AP failure results in the longest failover and failback periods of 40.14 and 30.02 seconds, respectively, due to WDS reconnection and forwarding table reconstruction. As the backup uCPE must wait for three consecutively missing advertisements before reassigning the VIP, failover in networking took longer time than failback. Using load balancing on WAN connections and the LAN network, HA networking with dual-WAN and dual-gateway increased network throughput by four times over a single uCPE with a single WAN. Generally, load balancing balances traffic directed to the server from clients. For the NVR application considered in the present study, the program running on the uCPE is a client application that initiates connections to multiple IP cameras. To balance the

NVR traffic, the original NVR application must be modified in some way. The traffic load balancing scheme employed in the present study assumes that each link has the same weight, and the scheme operates in a simple round-robin way. However, such an approach precludes certain vital services from being assigned a high priority. Consequently, future studies may examine the feasibility of integrated a priority-based load balancing method with the designed dual-uCPE system.

Acknowledgment

This work and its proof-of-concept system were supported by a project from Lanner Inc.

References

[1] J. Proença *et al.*, "Virtualization of Residential Gateways: A Comprehensive Survey," *IEEE Commun. Surveys & Tutorials*, vol. 21, no. 2, 2018, pp. 1462–82.

[2] F. Meneses, D. Corujo, and R. L. Aguiar, "Virtualization of Customer Equipment: Challenges and Opportunities," *Internet Technology Letters*, vol. 3, no. 6, e183, 2020.

[3] B. Han *et al.*, "Network Function Virtualization: Challenges and Opportunities for Innovations," *IEEE Commun. Mag.*, vol. 53, no. 2, 2015, pp. 90–97.

[4] S. Troia *et al.*, "SD-WAN: How the Control of the Network Can be Shifted from Core to Edge," *2021 Int'l. Conf. Optical Network Design and Modeling (ONDM)*, 2021, pp. 1–3.

[5] S. Cheruvu *et al.*, "IoT Vertical Applications and Associated Security Requirements," *Demystifying Internet of Things Security*, Springer, 2020, pp. 413–62.

[6] M. Choi *et al.*, "Duplex Communication Method for Railway Vehicle Communication System," *IEEE Int'l. Conf. Information and Communication Tech. Convergence (ICTC)*, IEEE, 2020, pp. 823–27.

[7] N. Herbaut *et al.*, "Migrating to a NFV-based Home Gateway: Introducing a Surrogate vNF Approach," *Int'l. Conf. Network of the Future (NOF)*, IEEE, 2015, pp. 1–7.

[8] Z. Bronstein and E. Shraga, "NFV Virtualisation of the Home Environment," *11th IEEE Consumer Commun. and Networking Conf. (CCNC)*, 2014, pp. 889–904.

[9] K. Suksomboon, M. Fukushima, and M. Hayashi, "Optimal Virtualization of Functionality for Customer Premise Equipment," *IEEE Int'l. Conf. Commun. (ICC)*, 2015, pp. 5685–90.

[10] J. Proença *et al.*, "Building an NFV-Based vRGW: Lessons Learned," *14th IEEE Annual Consumer Communications & Networking Conf. (CCNC)*, 2017, pp. 653–58.

[11] N. Huang *et al.*, "A Novel vCPE Framework for Enabling Virtual Network Functions with Multiple Flow Tables Architecture in SDN Switches," *19th Asia-Pacific Network Operations and Management Symp. (APNOMS)*, 2017, pp. 64–69.

[12] Ericsson, "Virtual CPE and Software Defined Networking," Ericsson Inc., Tech. Rep., 2015; http://www.ericsson.com/res/docs/2014/virtual- cpeand-software-defined-networking.pdf.

[13] F. Meneses *et al.*, "Dynamic Modular vCPE Orchestration in Platform as a Service Architectures," *8th IEEE Int'l. Conf. Cloud Networking (CloudNet)*, 2019, pp. 1–6.

[14] J. Proença *et al.*, "Virtualizing Customer Premises Equipment," *IEEE Conf. Network Function Virtualization and Software Defined Networks (NFV-SDN)*, 2021, pp. 104–05.

[15] L. Zhou *et al.*, "IoT Gateway Edge VNFs on uCPE," *IEEE Conf. Network Function Virtualization and Software Defined Networks (NFV-SDN)*, 2018, pp. 1–2.

Biographies

YING-DAR LIN (ydlin@cs.nctu.edu.tw) is a Chair Professor of computer science at National Yang Ming Chiao Tung University (NYCU), Taiwan. He received his Ph.D. in computer science from the University of California at Los Angeles (UCLA) in 1993. His research interests include network softwarization, cybersecurity, and wireless communications. His work on multi-hop cellular was the first along this line, and has been cited over 1000 times. He is an IEEE Fellow and IEEE Distinguished Lecturer. He has served or is serving on the editorial boards of several IEEE journals and magazines, and was the Editor-in-Chief of IEEE Communications Surveys and Tutorials (COMST), during 2016–2020.

KUAN-YU HO (kyho.cs08g@nctu.edu.tw) received his M.S. degree in computer science at National Yang Ming Chiao Tung University (NYCU), Taiwan, in 2022. His research interests include software-defined networking, network function virtualization, edge computing and high-availibility system.

WIDHI YAHYA (widhi.yahya@ub.ac.id) received his Ph.D degree in Electrical Engineering and Computer Science (EECS), National Yang Ming Chiao Tung University (NYCU), Taiwan in 2023. He is a lecturer and researcher in computer science at Universitas Brawijaya, Indonesia. His research interest areas are in mobile edge computing, software-defined networking, optimization with meta-heuristics and Reinforcement learning approaches

CHI-YU LI [M] (chiyuli@cs.nctu.edu.tw) received the bachelor's and master's degrees from the Department of Computer Science, National Chiao Tung University (NCTU), Hsinchu, Taiwan, and the Ph.D. degree in computer science from the University of California, Los Angeles (UCLA), Los Angeles, CA, USA, in 2015. He is currently an Associate Professor with the Department of Computer Science, National Yang Ming Chiao Tung University (NYCU). His research interests include wireless networking, mobile networks and systems, and network security. He received the Award of MTK Young Chair Professor (2016), MOST Young Scholar Research Award (2017–2020), and Best Paper Award in IEEE CNS 2018.

YUAN-CHENG LAI (laiyc@cs.ntust.edu.tw) received his Ph.D. degree from the Department of Computer and Information Science from National Chiao Tung University in 1997. He joined the faculty of the Department of Information Management at National Taiwan University of Science and Technology in August 2001 and has been a professor since February 2008. His research interests include performance analysis, network security, 5G mobile networks and Internet of Things.

JEANS H. TSENG (jeans_tseng@lannerinc.com) is SVP, CTO and GM for Telecom Applications BU at Lanner Electronics Inc. He leads the development of vCPE benchmarking and NFV PoC with various 3rd parties including Verizon, Wind River and NTT Lab. He has over 30 years of experience in sales and technical management in IT, Networking, and Telecom industries. Prior to Lanner, he was with Alcatel-Lucent from 1999-2013 as country General Manager, VP of sales and President of Taiwan Alcatel Technology Co. Ltd.. Prior to that, he was with Xylan Corp. as country general manager, D-Link and Siemens Telecommunication Systems as Director of the R\&D Center and software engineer.