Experimental Demonstration of Building and Operating QoS-aware Survivable vSD-EONs with Transparent Resiliency

Jie Yin[1], **Jiannan Guo**[1], **Bingxin Kong**[1], **Heqing Yin**[1], **and Zuqing Zhu**[1,*]

[1]*University of Science and Technology of China, Hefei, Anhui, 230027, China*
[*]*zqzhu@ieee.org*

Abstract: Software-defined elastic optical networks (SD-EONs) provide operators more flexibility to customize their optical infrastructures dynamically. By leveraging infrastructure-as-a-service (IaaS), virtual SD-EONs (vSD-EONs) can be realized to further enhance the adaptivity of SD-EONs and shorten the time-to-market of new services. In this paper, we design and demonstrate the building and operating of quality-of-service (QoS) aware survivable vSD-EONs that are equipped with transparent data plane (DP) resiliency. Specifically, when slicing a vSD-EON, our network hypervisor (NHV) chooses to use "1:1" virtual link (VL) protection or on-demand VL remapping as the DP restoration scheme, according to the service-level agreement (SLA) between the vSD-EON's operator and the infrastructure provider (InP). Then, during an actual substrate link (SL) failure, the NHV realizes automatic DP restoration that is transparent to the controllers of vSD-EONs. We build a network testbed to demonstrate the creation of QoS-aware survivable vSD-EONs, the activation of lightpaths in the vSD-EONs to support upper-layer applications, and the automatic and simultaneous QoS-aware DP restorations during an SL failure. The experimental results indicate that our vSD-EON slicing system can build QoS-aware survivable vSD-EONs on-demand, operate them to set up lightpaths for carrying real application traffic, and facilitate differentiated DP restorations during SL failures to recover the vSD-EONs' services according to their SLAs.

OCIS codes: (060.4250) Networks; (060.4510) Optical communications.

References and links

1. Z. Zhu, W. Lu, L. Zhang, and N. Ansari, "Dynamic service provisioning in elastic optical networks with hybrid single-/multi-path routing," J. Lightw. Technol. **31**, 15-22 (2013).
2. O. Gerstel, M. Jinno, A. Lord, and B. Yoo, "Elastic optical networking: a new dawn for the optical layer?" IEEE Commun. Mag. **50**, S12-S20 (2012).
3. P. Lu, L. Zhang, X. Liu, J. Yao, and Z. Zhu, "Highly-efficient data migration and backup for big data applications in elastic optical inter-datacenter networks," IEEE Netw. **29**, 36-42 (2015).
4. W. Lu, and Z. Zhu, "Malleable reservation based bulk-data transfer to recycle spectrum fragments in elastic optical networks," J. Lightw. Technol. **33**, 2078-2086 (2015).
5. L. Liu, R. Munoz, R. Casellas, T. Tsuritani, R. Martinez, and I. Morita, "OpenSlice: an OpenFlow-based control plane for spectrum sliced elastic optical path networks," Opt. Express **21**, 4194-4204 (2013).
6. Z. Zhu, X. Chen, C. Chen, S. Ma, M. Zhang, L. Liu, and B. Yoo, "OpenFlow-assisted online defragmentation in single-/multi-domain software-defined elastic optical networks," J. Opt. Commun. Netw. **6**, 901-910 (2015).
7. R. Casellas, R. Martinez, R. Munoz, R. Vilalta, L. Liu, T. Tsuritani, and I. Morita, "Control and management of flexi-grid optical networks with an integrated stateful path computation element and OpenFlow controller," J. Opt. Commun. Netw. **5**, A57-A65 (2013).
8. C. Chen, X. Chen, M. Zhang, S. Ma, Y. Shao, S. Li, M. Suleiman, and Z. Zhu, "Demonstrations of efficient online spectrum defragmentation in software-defined elastic optical networks," J. Lightw. Technol. **32**, 4701-4711 (2014).
9. Z. Zhu, C. Chen, S. Ma, L. Liu, X. Feng, and B. Yoo, "Demonstration of cooperative resource allocation in an OpenFlow-controlled multidomain and multinational SD-EON testbed," J. Lightw. Technol. **33**, 1508–1514 (2015).
10. S. Li, W. Lu, X. Liu, and Z. Zhu, "Fragmentation-aware service provisioning for advance reservation multicast in SD-EONs," Opt. Express **23**, 25804-25813 (2015).
11. L. Gong and Z. Zhu, "Virtual optical network embedding (VONE) over elastic optical networks," J. Lightw. Technol. **32**, 450-460 (2014).

12. R. Munoz, R. Vilalta, R. Casellas, R. Martinez, T. Szyrkowiec, A. Autenrieth, V. Lopez, and D. Lopez, "Integrated SDN/NFV management and orchestration architecture for dynamic deployment of virtual SDN control instances for virtual tenant networks," J. Opt. Commun. Netw. **7**, B62-B70 (2015).
13. Y. Ran, "Considerations and suggestions on improvement of communication network disaster countermeasures after the wenchuan earthquake," IEEE Commun. Mag. **40**, 44-47 (2011).
14. H. Jiang, Y. Wang, L. Gong, and Z. Zhu, "Availability-aware survivable virtual network embedding in optical datacenter networks," J. Opt. Commun. Netw. **7**, 1160-1171 (2015).
15. J. Yin, J. Guo, B. Kong, and Z. Zhu, "Demonstration of survivable vSD-EON slicing with automatic data plane restoration to support reliable video streaming," in Proc. of OFC, (2017), 1-3.
16. A. Jarray, Y. Song, and A. Karmouch, "Resilient virtual network embedding," in Proc. of ICC, (2013), 3461-3465.
17. B. Guo, C. Qiao, J. Wang, H. Yu, Y. Zuo, J. Li, Z. Chen, and Y. He, "Survivable virtual network design and embedding to survive a facility node failure," J. Lightw. Technol. **32**, 483-493 (2014).
18. H. Jiang, L. Gong, and Z. Zhu, "Efficient joint approaches for location-constrained survivable virtual network embedding," in Proc. of GLOBECOM, (2014), 1810-1815.
19. C. Meixner, F. Dikbiyik, M. Tornatore, C. Chuah, and B. Mukherjee, "Disaster-resilient virtual-network mapping and adaptation in optical networks," in Proc. of ONDM, (2013), 107-112.
20. W. Xie, J. Jue, Q. Zhang, X. Wang, Q. She, P. Palacharla, and M. Sekiya, "Survivable virtual optical network mapping in flexible-grid optical networks," in Proc. of ICNC, (2014), 221-225.
21. G. Luo, H. Ding, J. Zhou, J. Zhang, Y. Zhao, B. Chen, and C. Ma, "Survivable virtual optical network embedding with probabilistic network-element failures in elastic optical networks," in Proc. of ICOCN, (2014), 1-4.
22. Z. Ye, A. Patel, P. Ji, and C. Qiao, "Survivable virtual infrastructure mapping with dedicated protection in transport software-defined networks," J. Opt. Commun. Netw. **7**, A183-A189 (2015).
23. D. Siracusa, F. Pederzolli, M. Gerola, A. Zanardi, D. Fauci, and G. Galimberti, "Demonstration of a hybrid S-DN/GMPLS control plane for optical virtual private networks with restoration capabilities," in Proc. of ECOC, (2016), 1-3.
24. J. Yin, H. Yin, S. Liu, Z. Zhou, X. Chen, and Z. Zhu, "On-demand and reliable vSD-EON provisioning with correlated data and control plane embedding," in Proc. of GLOBECOM, (2016), 1-6.
25. OpenVirteX [Online]. http://ovx.onlab.us/.
26. Y. Wang, D. Lin, C. Li, J. Zhang, P. Liu, C. Hu, and G. Zhang, "Application driven network: Providing on-demand services for applications," in Proc. of ACM SIGCOMM, (2016), 617-618.
27. Optical Tranport Protocol Extensions [Online]. https://www.opennetworking.org/images/stories/downloads/sdn-resources/onf-specifications/openflow/Optical Transport Protocol Extensions V1.0.pdf.
28. F. Ji, X. Chen, W. Lu, J. Rodrigues, and Z. Zhu, "Dynamic p-cycle protection in spectrum-sliced elastic optical networks," J. Lightw. Technol. **32**, 1190-1199 (2014).
29. X. Chen, M. Tornatore, F. Ji, W. Zhou, C. Chen, D. Hu, S. Zhu, L. Jiang, and Z. Zhu, "Flexible availability-aware differentiated protection in software-defined elastic optical networks," J. Lightw. Technol. **33**, 3872-3882 (2015).
30. X. Chen, S. Zhu, L. Jiang, and Z. Zhu, "On spectrum efficient failure-independent path protection p-cycle design in elastic optical networks," J. Lightw. Technol. **33**, 3719-3729 (2015).
31. Open Network Operating System (ONOS) [Online]. http://onosproject.org/.
32. OpenFlow Switch Specificaltion version 1.3.4 [Online]. https://www.opennetworking.org/images/stories/downloads /sdn-resources/onf-specifications/openflow/openflow-switch-v1.3.4.pdf.
33. Open vSwitch [Online]. http://www.openvswitch.org/.
34. EasyDarwin [Online]. http://www.easydarwin.org/.
35. Official website of Rio 2 produced by the the Blue Sky Studio [Online]. http://www.riomovies.com/.

1. Introduction

Flexible-grid elastic optical networks (EONs) rely on agile bandwidth management [1] in the optical layer to make optical networks more adaptive and not be restricted by the fixed spectrum grids or rigid optical transmission setup (*e.g.*, bit-rate and modulation selections) [2, 3]. Therefore, it is believed that with EONs, network operators have more freedom to configure their optical infrastructures according to customer demands [4]. In the meantime, the additional freedom can make service provisioning more complex, and thus a powerful network control and management (NC&M) mechanism would be required to address this issue. One possible way to realize such an NC&M mechanism is to build software-defined EONs (SD-EONs) that leverage software-defined networking (SDN) [5, 6]. Previously, people have already demonstrated that SD-EONs can achieve programmable and application-aware optical networking with enhanced service support [7–10]. On the other hand, the network virtualization that realizes infrastructure-as-a-service (IaaS) and enables multiple tenants to share substrate optical infrastructure efficiently can further enhance the adaptivity of SD-EONs and shorten the time-to-market of new services [11, 12]. Specifically, IaaS slices virtual SD-EONs (vSD-EONs) on a shared substrate EON for the tenants, each of which can deploy its own services in a vSD-EON.

However, it is known that optical networks are not always intact. For example, the 2008 Wenchuan earthquake had destroyed $30,000$ kilometers of fiber-optic cables and caused long service interruption [13]. Note that, the situation can become even worse for vSD-EONs. Specifically, since IaaS makes multiple vSD-EONs share the same substrate elements (*i.e.*, optical switches and fiber links), the failure of a single substrate element might bring down multiple vSD-EONs simultaneously [14]. Hence, it is relevant to consider how to realize survivable vSD-EON slicing that can guarantee the intactness of network services during substrate failures. Moreover, since the controllers of vSD-EONs normally do not interact with substrate elements directly (*i.e.*, through a network hypervisor (NHV)), if the data plane (DP) restoration can be automatic and transparent to the controllers, simplified vSD-EON management and fast service recovery can be achieved [15].

Previously, the studies in [16–18] have addressed the problem of survivable virtual network embedding (S-VNE) in generic substrate networks, and proposed various algorithms to slice virtual networks that can survives from substrate failures. Meanwhile, people have also tried to solve the problem of S-VNE in substrate networks that are fixed-grid wavelength-division multiplexing (WDM) networks [14, 19] or EONs [20–22]. Specifically, these investigations developed algorithms that consider the wavelength/spectrum assignment in substrate optical networks for realizing S-VNE. However, all these studies only tackled the problem of slicing survivable virtual networks from the perspective of algorithm design and evaluated the proposed algorithms with numerical simulations, but did not validate their proposals in practical network systems. In [23, 24], the authors demonstrated the control plane (CP) operations of slicing survivable virtual optical networks in software-emulated network systems, while the related DP operations, such as failure detection and path switching, were omitted. Note that, for effectively evaluating how the DP restoration schemes would affect the quality-of-service (QoS) of upper-layer applications in vSD-EONs, we should incorporate the “full stack” investigation that not only considers the building of survivable vSD-EONs with DP resiliency but also addresses the operating of them to carry real application traffic during DP restoration.

In this paper, we extend our preliminary work in [15] and experimentally demonstrate the building and operating of QoS-aware survivable vSD-EONs that are equipped with transparent DP resiliency. We first design and implement a QoS-aware survivable vSD-EON slicing system that utilizes an OpenVirteX [25] based NHV to realize automatic DP restoration that is transparent to the controllers of vSD-EONs. Specifically, when slicing a vSD-EON, the NHV determines whether to use “1:1” virtual link (VL) protection or on-demand VL remapping to restore its DP services when substrate link (SL) failure happens, according to the service-level agreement (SLA) between the vSD-EON’s operator and the infrastructure provider (InP). The SLA is normally based on the actual upper-layer applications that the vSD-EON carries [26]. For instance, if the vSD-EON is mainly used to support delay-insensitive data-oriented applications (*e.g.*, data backup), on-demand VL remapping would be good enough for its DP restoration. Otherwise, if the vSD-EON carries delay-sensitive flow-oriented applications (*e.g.*, real-time video streaming), “1:1” VL protection would be required. Then, when an SL failure happens, the NHV incorporates the corresponding DP restoration schemes to automatically recover the services of all the vSD-EONs that are affected. Therefore, the DP restoration is made transparent to the controllers of vSD-EONs. Here, we also develop an SL monitoring subsystem that checks the operating status of all the SLs continuously and would send an alert message to the NHV when an SL failure is detected.

For the experimental demonstrations, we build a network testbed that consists of commercial optical transmission chassis (OTC), bandwidth-variable wavelength-selective switches (BV-

WSS'), and high-performance servers. With the testbed, we demonstrate the creation of QoS-aware survivable vSD-EONs, the activation of lightpaths in the vSD-EONs to support two types of upper-layer applications (*i.e.*, bulk file transfer and real-time video streaming), and the automatic and simultaneous QoS-aware DP restorations for the vSD-EONs during an SL failure. The experimental results indicate that our vSD-EON slicing system can build QoS-aware survivable vSD-EONs on-demand, operate them to set up lightpaths for carrying application traffic, and facilitate differentiated DP restorations during SL failures to recover the vSD-EONs' services according to their SLAs. To the best of our knowledge, this is the first experimental demonstration of building and operating of QoS-aware survivable vSD-EONs with transparent resiliency, which covers both CP and DP operations and incorporates real application traffic (*i.e.*, the "full stack" investigation).

The rest of the paper is organized as follows. Section 2 describes the architecture, functional modules and communication protocols of our QoS-aware survivable vSD-EON slicing system. The operation procedure of the proposed system is presented in Section 3, and we discuss the experimental demonstrations in Section 4. Finally, Section 5 summarizes the paper.

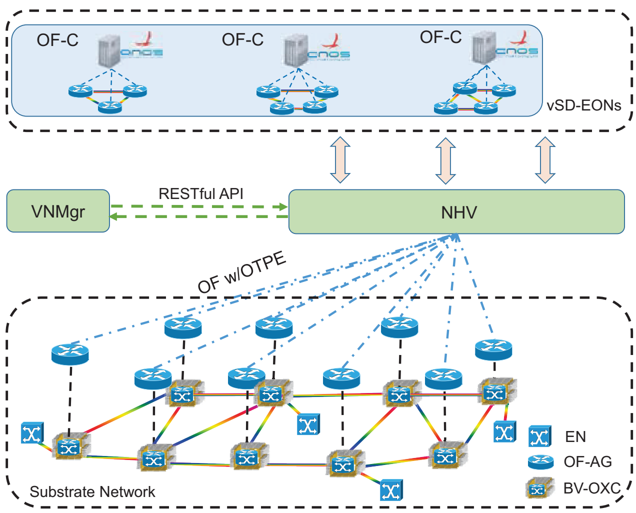

Fig. 1. Architecture of our QoS-aware survivable vSD-EON slicing system, OF-C: ONOS-based OpenFlow controller, VNMgr: Virtual network manager, NHV: Network hypervisor, RESTful API: RESTful application programming interface, OF w/ OTPE: OpenFlow with optical transport protocol extensions, EN: End-node, OF-AG: OpenFlow agent, BV-OXC: Bandwidth-variable optical cross-connect.

2. QoS-aware Survivable vSD-EON Slicing System with Transparent Resiliency

2.1. Network Architecture

Figure 1 shows the network architecture of our proposed QoS-aware survivable vSD-EON slicing system. It can be seen that the substrate network is an EON that consists of a few bandwidth-variable optical cross-connects (BV-OXCs) connected by fiber links. We have an OpenFlow agent (OF-AG) locally attached to each BV-OXC for controlling its operation, *i.e.*, optical spectrum management. There are also some end-nodes in the substrate network, from which the customers can launch upper-layer applications. To manage the substrate network for vSD-EON slicing, the InP incorporates a virtual network manager (VNMgr) and a network hypervisor (N-HV). Specifically, when a vSD-EON operator sends a vSD-EON request to the InP, the VNMgr calculates the S-VNE scheme (*i.e.*, the node mapping and link mapping results) for its virtual

DP (vDP) based on the SLA of the vSD-EON operator, and then forwards the scheme to the NHV by using the RESTful API. The NHV communicates with the OF-AGs using OpenFlow protocol with optical transport protocol extensions (OF w/ OTPE) [27], realizes control message interpretation between the virtual and substrate network elements, and enables the OpenFlow controller (OF-C) of each vSD-EON to realize transparent NC&M. In the meantime, the NHV monitors the BV-OXCs' working status proactively, and will invoke DP restoration automatically to recover the services of affected vSD-EONs in case of SL failures. Therefore, transparent DP resiliency can be realized for all the vSD-EONs. In this work, we consider two DP restoration schemes to satisfy the SLAs of vSD-EONs. Specifically, "1:1" VL protection is utilized to support the vSD-EONs that carry delay-sensitive flow-oriented applications to minimize the service recovery latency, while for the vSD-EONs that are built for carrying delay-insensitive data-oriented applications, on-demand VL remapping is used to save substrate resources. Note that, these two DP restoration schemes are chosen for proof-of-concept demonstrations, and by modifying the VNMgr's program, we can easily realize more sophisticated schemes [28–30] to further balance the trade-off between service recovery latency and substrate resource utilization.

2.2. Design of Function Modules and Protocols

To explain the design of function modules and communication protocols, we zoom in the vertical structure of a vSD-EON in Fig. 2. We implement the OF-C of each vSD-EON based on the ONOS platform [31]. Specifically, we leverage OF v1.3.4 [32] and OTPE [27] and expand the functionality of related ONOS modules to support the routing and spectrum assignment (RSA) for lightpath setup on BV-OXCs. For each vSD-EON, the OF-C is instantiated in high-performance servers on demand with network function virtualization (NFV). The OF-C communicates with the OF-AG on each BV-OXC using OF w/ OTPE through the NHV, which is realized by modifying OpenVirteX [25] to support OF w/ OTPE. Note that, the design of OpenVirteX determines that a single NHV based on it can support a maximum number of 255 vSD-EONs [25]. Considering the fact that due to the constraints on spectrum resources, a practical EON cannot support too many vSD-EONs, this upper limit is sufficient and would not restrict the scalability of our proposed system. The NHV translates the OF messages from the OF-C for its virtual optical switches into those that can be understood by the OF-AGs. Hence, the NHV realizes the vSD-EON slicing.

The in-house developed VNMgr is programmed to calculate the S-VNE schemes for vSD-EONs. As we have explained above, it inspects the vSD-EON operators' SLAs and uses either "1:1" VL protection or on-demand VL remapping to ensure the survivability of vSD-EONs accordingly. Here, for "1:1" VL protection, the VNMgr uses an S-VNE algorithm that is modified from the one that we developed in [14], which leverages the shared-path protection to calculate both the working and backup substrate paths for each VL. For on-demand VL remapping, the backup substrate path is obtained by finding the shortest feasible path based on the current network status, when the working substrate path of a VL is affected by an SL failure. The communication between the NHV and VNMgr is based on the RESTful API. The OF-AGs are used to configure the BV-OXCs according to the flow-entries from the OF-C, which has been translated by the NHV, and implement the required RSA for lightpath setup. An OF-AG consists of two parts, *i.e.*, the OF client and equipment controller. The OF client is programmed based on OpenvSwitch [33] to parse the OF messages for lightpath management, while the equipment controller is in-house developed and can configure the bandwidth-variable wavelength-selective switches (BV-WSS') in the BV-OXC through a serial port according to the instructions from the OF client. In our experimental testbed, we run the VNMgr, NHV, and OF-AGs on high-performance Linux servers.

As shown in Fig. 2, the BV-OXC is realized with Finisar 1×9 BV-WSS'. On each of its input ports, we implement a link monitoring module (LMM), which monitors the corresponding SL

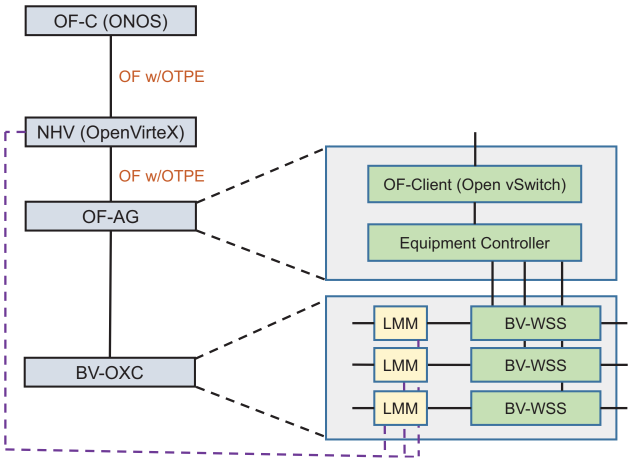

Fig. 2. Zoom-in view of the vertical structure of a vSD-EON, OF Client: OpenFlow client, BV-WSS: Bandwidth-variable wavelength-selective switch, LMM: Link monitoring module.

proactively. Specifically, if an LMM detects that the input optical power is below -35 dBm, it will generate an alert message to report the SL failure to the NHV. Then, the NHV invokes DP restoration automatically to recover the services of the affected vSD-EONs. Note that, if the vSD-EON uses "1:1" VL protection, the path switching is conducted directly by the NHV, while if it uses on-demand VL remapping, the NHV will need to communicate with the VNMgr for obtaining the remapping schemes of the affected VLs. Hence, the recovery latency is different for vSD-EONs with different QoS requirements.

3. Operation Procedure

To fully demonstrate the functionalities of our vSD-EON slicing system, we consider the creation of QoS-aware survivable vSD-EONs, the activation of lightpaths in the vSD-EONs to support two types of upper-layer applications (*i.e.*, bulk file transfer and real-time video streaming), and the automatic and simultaneous QoS-aware DP restorations for the vSD-EONs during an SL failure. The detailed operation procedure is illustrated in Fig. 3. First of all, when a vSD-EON operator sends a vSD-EON request to the InP, VNMgr calculates the S-VNE scheme of vDP embedding based on the information in it, *i.e.*, the vDP topology, bandwidth requirement, SLA, and end-node locations. Then, the vDP embedding is realized with the NHV, which maps the virtual nodes (VNs) and VLs to substrate BV-OXCs and paths as suggested. Specifically, for the VL mapping, the NHV reserves enough spectra for the vSD-EON on the selected substrate paths, and if the vSD-EON uses "1:1" VL protection, enough spectral resources are reserved on both working and backup substrate paths for each VL. Meanwhile, an OF-C is instantiated on the NFV server for the vSD-EON. At this moment, the InP accomplishes the vSD-EON creation, and it then hands over the OF-C to the vSD-EON operator, which will operate the vSD-EON and set up lightpaths in it to support the upper-layer applications from end-nodes.

When an end-node needs to launch an application, it forwards a lightpath request to the OF-AG that is on its local BV-OXC. The OF-AG adds the end-node's tenant ID to the request and sends it to the NHV. Upon receiving the request, the NHV determines which vSD-EON it belongs to based on the tenant ID, and forwards it to the corresponding OF-C in the NFV server. Then, the OF-C calculates an RSA scheme for the lightpath based on the vDP topology that it is aware of, uses the flow-entries in *FlowMod* messages to encode the scheme, and sends the messages to the NHV. The NHV translates the flow-entries into what they should be in

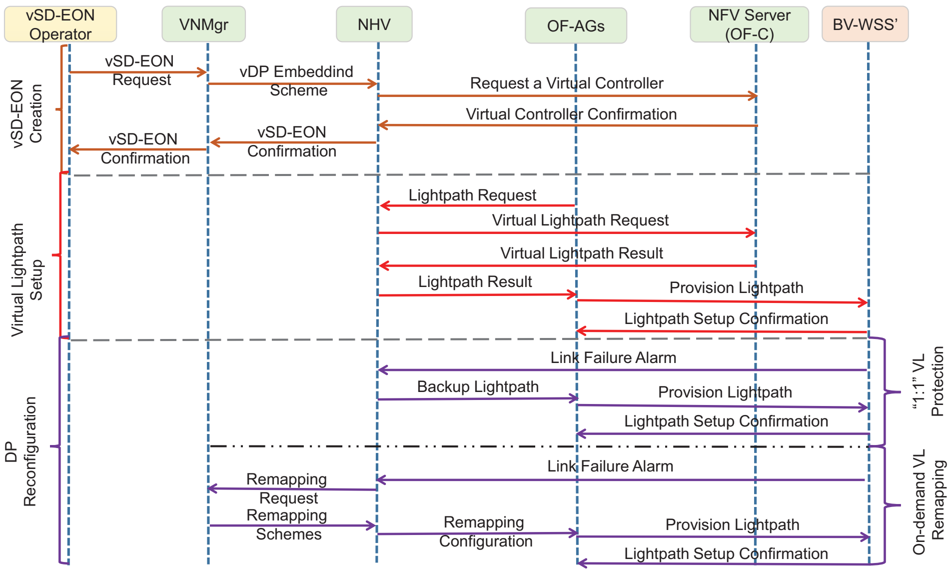

Fig. 3. Operation procedure for vSD-EON creation, lightpath establishment, and transparent DP restoration.

the substrate network and forwards the translated messages to the related OF-AGs. With the *FlowMod* messages, the OF-AGs configure their BV-OXCs to establish the lightpath in the substrate network. Then, the lightpath is used to carry application traffic in the vSD-EON.

The operation of the vSD-EON can be interrupted by SL failures. In that case, the NHV will restore its vDP automatically, and the whole restoration process is transparent to the OF-C. We use Fig. 4 to explain the transparent DP restoration schemes whose operation procedures are shown in Fig. 3. As the first step, when an SL failure happens, the NHV gets informed immediately. In Fig. 4(a), if the vSD-EON uses on-demand VL remapping, the NHV updates the substrate topology in the VNMgr and communicates with it to obtain the remapping schemes for the affected VLs. Then, the NHV implements the remapping schemes by updating the vSD-EON's embedding scheme in its database and sending corresponding *FlowMod* messages to the related OF-AGs. Hence, the vSD-EON's lightpath(s) that are affected by the SL failure will be rerouted in the substrate network. On the other hand, as shown in Fig. 4(b), if the vSD-EON

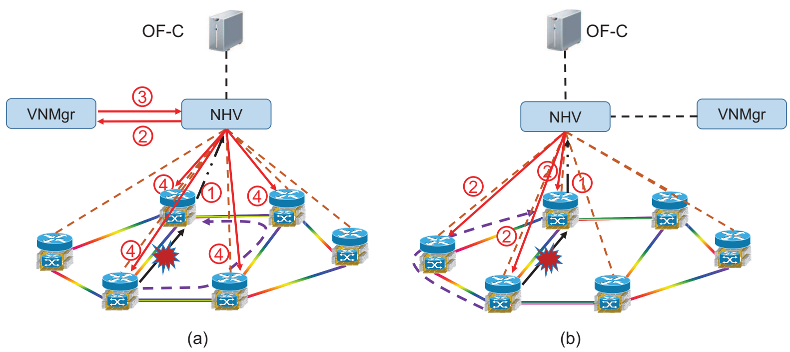

Fig. 4. Transparent DP restoration schemes for vSD-EONs using (a) on-demand VL remapping, and (b) "1:1" VL protection.

uses “1:1” VL protection, the NHV will switch the affected VLs to their backup substrate paths directly, and there is no need to communicate with the VNMgr.

4. Experimental Demonstrations

This section discusses the experimental demonstrations that cover the creation of QoS-aware survivable vSD-EONs, the activation of lightpaths in the vSD-EONs to support upper-layer applications, and the automatic and simultaneous QoS-aware DP restorations for the vSD-EONs during an SL failure. The network testbed is shown in Fig. 5. We implement the CP elements of the vSD-EON slicing system (*i.e.*, the OF-Cs, VNMgr, NHV, and OF-AGs) in high-performance Linux servers. The testbed also includes two commercial OTCs (Huawei Optix OSN3500), several Finisar 1×9 BV-WSS’, erbium-doped optical fiber amplifiers (EDFAs), *etc.*, for DP operations. The OTCs are used to aggregate/deaggregate application traffic. Specifically, an OTC can aggregate the Ethernet traffic (at 100 Mb/s or 1 Gb/s) from clients into an STM-64 optical signal (at 9.95 Gb/s), or the other way around. The spectrum range of each BV-WSS is $[1528.43, 1566.88]$ nm in the C-band, and its minimum bandwidth allocation granularity (*i.e.*, a frequency slot (FS)) is 12.5 GHz.

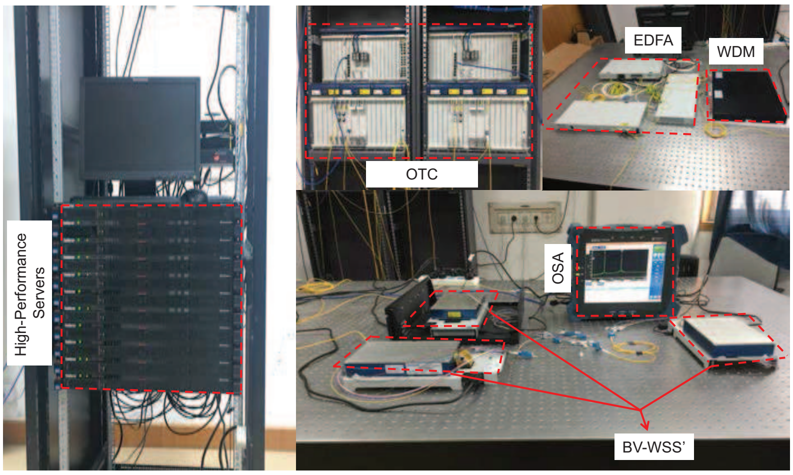

Fig. 5. Experimental testbed, OTC: Optical transmission chassis, OSA: Optical spectrum analyzer, EDFA: Erbium-doped optical fiber amplifier, WDM: Wavelength-division multiplexer/demultiplexer.

4.1. Creation of QoS-aware Survivable vSD-EONs

We first demonstrate the creation of QoS-aware survivable vSD-EONs in the substrate network that shows in Fig. 6(a), which consists of 6 substrate nodes. We assume that there are two vSD-EON requests for *Tenants* 1 and 2. Each tenant requires a triangle vDP topology as illustrated in Fig. 6(a). Here, we assign the SLA value as 1 to indicate that the vSD-EON uses on-demand VL remapping, while if the SLA value is 2, “1:1” VL protection is used. *Tenant* 1 has a bandwidth requirement of 180 GHz (*i.e.*, 15 FS’) and its SLA value is 2. Its embedding scheme is shown in Fig. 6(a) with purple solid lines. It can be seen that the three VLs are mapped onto substrate paths 2-3, 3-5, and 5-2. Since the vSD-EON uses “1:1” VL protection, the backup substrate path is also determined for each VL. Specifically, the backup paths for 2-3, 3-5, and 5-2 are 2-1-3, 3-4-5, and 5-6-4-2, respectively. Fig. 6(b) shows the Wireshark capture of a control message from the VNMgr to the NHV for embedding a VL of *Tenant* 1. On both working and backup substrate paths, the FS-block $[190, 204]$, which covers the FS’ whose center wavelengths range from 1546.37 nm to 1547.77 nm, is reserved for *Tenant* 1. On the other hand, the bandwidth requirement of *Tenant* 2 is 100 GHz (*i.e.*, 8 FS') and its SLA value is 1. Fig. 6(a) shows its embedding scheme with green dash lines, *i.e.*, substrate paths 1-2, 2-3, and 3-1 are used to carry its three VLs. The vSD-EON uses on-demand VL remapping and thus no backup substrate path needs to be allocated during its creation. A control message from the VNMgr to embed a VL for *Tenant* 2 is shown in Fig. 6(c). The FS-block $[344, 351]$ (*i.e.*, FS' with center wavelengths ranging within $[1531.85, 1532.53]$ nm) is reserved for *Tenant* 2.

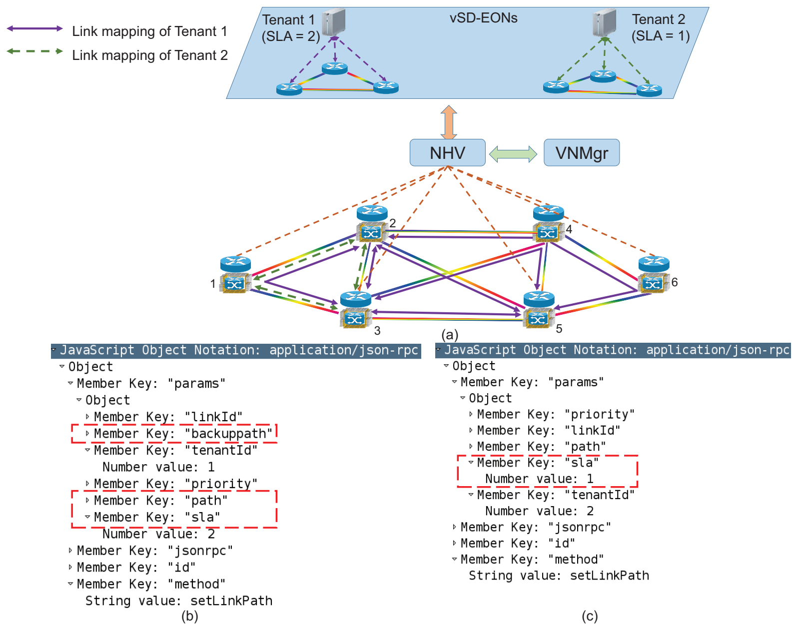

Fig. 6. Experiments of vSD-EONs creation, (a) substrate topology, and the control messages from VNMgr to NHV for embedding (b) a VL with "1:1" VL protection and (c) a VL with on-demand VL remapping.

Figure 7 illustrates the Wireshark capture of the control messages used for creating the two vSD-EONs, which provides the overall procedure of the vSD-EON slicing. Specifically, to create each vSD-EON, the slicing system needs to assign the tenant ID and accomplish node mapping, port mapping and link mapping. It can be seen that the vSD-EON creation for *Tenant* 1 takes 383 msec, while the same latency of *Tenant* 2 is 252 msec, which is shorter because no backup substrate path needs to be allocated.

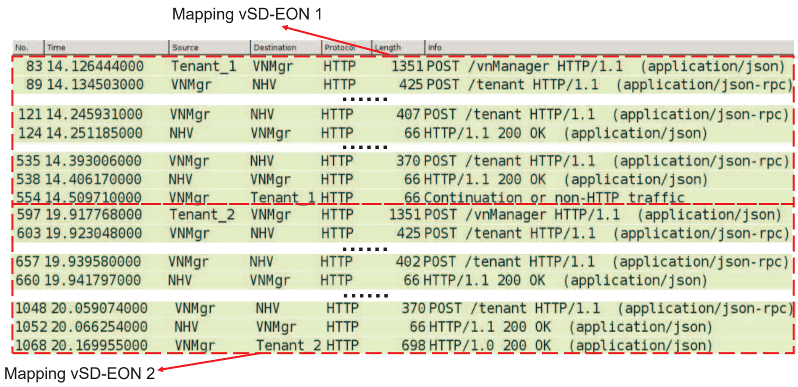

Fig. 7. Wireshark capture of control messages used for creating the vSD-EONs of *Tenants* 1 and 2.

4.2. Activation of Lightpaths in the vSD-EONs

After the vSD-EONs have been created, their end-nodes can send service requests to the OF-Cs, which will activate lightpaths to support the applications. Fig. 8(a) shows the Wireshark capture of messages used to set up two lightpaths for *Tenants* 1 and 2, respectively. The first request is from *Tenant* 1 and it asks for a lightpath for real-time high-definition (HD) video streaming. Specifically, the request is generated by the end-node that is locally connected to the BV-OXC on *Substrate Node (SN)* 3, and the lightpath should be routed to *SN* 2. The *PacketIn* message for the request is formulated and sent to the NHV, which determines it belongs to *Tenant* 1 based on the tenant ID, translates it and forwards the translated message to the OF-C of *Tenant* 1. The OF-C then calculates the RSA of the lightpath in its vSD-EON and encodes the result in *FlowMod* messages. Next, the *FlowMod* messages are translated by the NHV and forwarded to the BV-WSS' on *SNs* 2 and 3 to activate the lightpath, which occupies one FS.

On the end-nodes, we implement the streaming media software EasyDarwin [34] to realize real-time HD video streaming on the lightpath, which uses the user datagram protocol (UDP) protocol in the transport layer. Specifically, the software multiplexes/demultiplexes 50 video streams (each of which has the resolution of 1920×1080) and generates a traffic flow whose data-rate is around 1 Gbps. The video traffic is sent/received using the Gigabit Ethernet (1GbE) ports on the high-performance servers on *SNs* 2 and 3, between which the optical transmission is realized by the OTC. Specifically, we connect the servers' 1GbE ports to/from the 1GbE ports on the OTC, where the video traffic is groomed onto or de-groomed from an STM-64 optical signal whose center wavelength is 1547.71 nm. The second request is from *Tenant* 2 for bulk file transfer. The lightpath, which also uses one FS, is set up with the similar procedure as discussed above, and the OTC transmits the data traffic using a center wavelength of 1532.47 nm. Fig. 8(b) shows the spectra of the two established lightpaths, which indicates that the lightpaths in the two vSD-EONs share the same substrate path 2-3 with different spectrum assignments. Note that, to ensure an apple-to-apple comparison on the recovery latency, we also implement the file transfer using the UDP protocol and make sure that only the DP restoration scheme in the physical layer would affect the recovery latency.

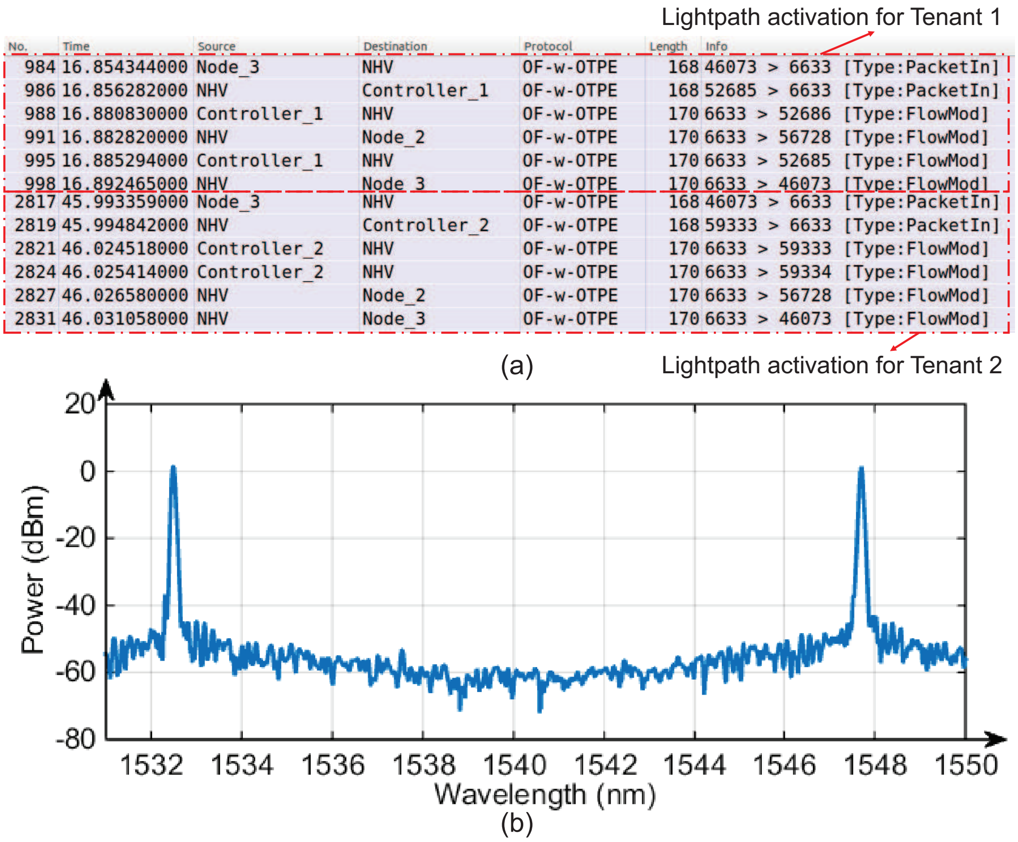

No.	Time	Source	Destination	Protocol	Length	Info
984	16.854344000	Node_3	NHV	OF-w-OTPE	168	46073 > 6633 [Type:PacketIn]
986	16.856282000	NHV	Controller_1	OF-w-OTPE	168	52685 > 6633 [Type:PacketIn]
988	16.880830000	Controller_1	NHV	OF-w-OTPE	170	6633 > 52686 [Type:FlowMod]
991	16.882820000	NHV	Node_2	OF-w-OTPE	170	6633 > 56728 [Type:FlowMod]
995	16.885294000	Controller_1	NHV	OF-w-OTPE	170	6633 > 52685 [Type:FlowMod]
998	16.892465000	NHV	Node_3	OF-w-OTPE	170	6633 > 46073 [Type:FlowMod]
2817	45.993359000	Node_3	NHV	OF-w-OTPE	168	46073 > 6633 [Type:PacketIn]
2819	45.994842000	NHV	Controller_2	OF-w-OTPE	168	59333 > 6633 [Type:PacketIn]
2821	46.024518000	Controller_2	NHV	OF-w-OTPE	170	6633 > 59333 [Type:FlowMod]
2824	46.025414000	Controller_2	NHV	OF-w-OTPE	170	6633 > 59334 [Type:FlowMod]
2827	46.026580000	NHV	Node_2	OF-w-OTPE	170	6633 > 56728 [Type:FlowMod]
2831	46.031058000	NHV	Node_3	OF-w-OTPE	170	6633 > 46073 [Type:FlowMod]

Fig. 8. Experimental results for lightpath activation, (a) Wireshark capture of messages used, and (b) spectra of the established lightpaths in two vSD-EONs.

4.3. Automatic and Simultaneous QoS-aware DP Restoration for the vSD-EONs

Finally, we demonstrate the automatic and simultaneous QoS-aware DP restorations for the two vSD-EONs. Specifically, when the application traffic on the two established lightpaths are ongoing, we disconnect the fiber that connects *SNs* 2 and 3 to emulate an SL failure. Fig. 9(a) captures the messages used for the DP restorations whose procedure is explained as follows. In **Step 1**, the LMM sends an alert message to the NHV immediately after it detects the SL failure. Then, in **Step 2**, the NHV first tries to recovery the VL of *Tenant* 1, since its SLA value is higher. Basically, since the backup substrate path is precalculated (*i.e.*, 2-1-3), the NHV directly sends *FlowMod* messages to *SNs* 1, 2 and 3 and asks them to perform the path switching to restore the lightpath. Then, the real-time HD video streaming is recovered with a relatively short latency. Next, in **Step 3**, the NHV starts to recover the lightpath of *Tenant* 2 with on-demand VL remapping. Specifically, it communicates with the VNMgr to report the SL failure and obtain the VL remapping scheme. Then, the NHV implements the scheme by sending corresponding *FlowMod* messages to related SNs as shown in **Step 4**. Here, the restoration substrate path for the VL remapping is 2-4-5-3. Finally, in **Step 5**, the NHV communicates with the VNMgr and LMM to inform them that the DP restorations are accomplished and update the network status.

Then, we conduct experiments to measure the QoS parameters of the tenants' applications. Here, in order to make the measurements more accurate for the applications, each experiment only build a vSD-EON for one tenant and invoke DP restoration for it, *i.e.*, performing indepen-

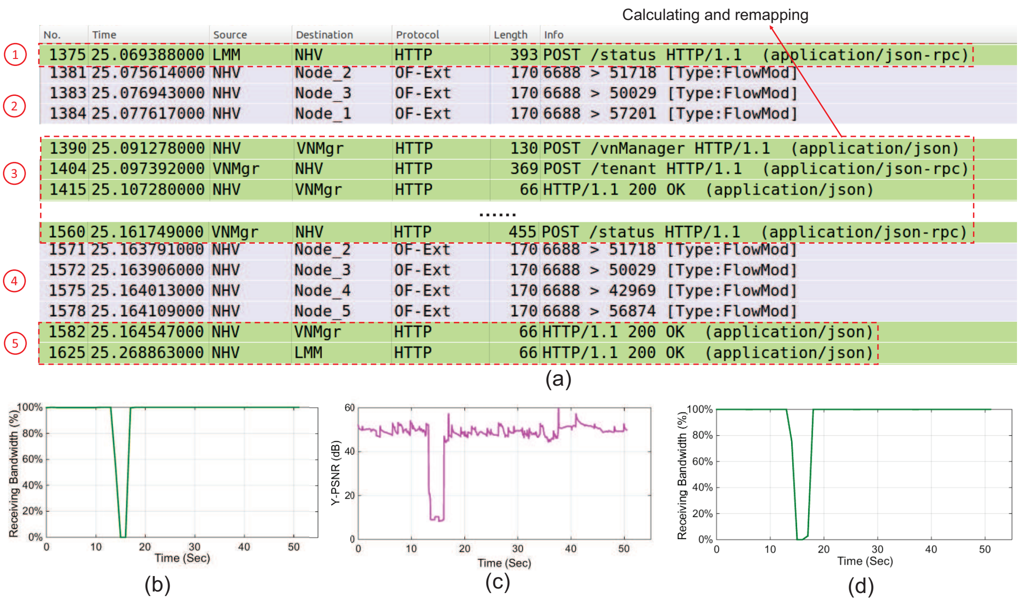

Fig. 9. Experimental results for DP restorations, (a) Wireshark capture of messages used, (b) receiving bandwidth of video streaming in *Tenant* 1, (c) Y-PSNR of video playback in *Tenant* 1, and (d) receiving bandwidth of bulk file transfer in *Tenant* 2.

dent instead of simultaneous DP restorations. The receiving bandwidth of the video streaming in *Tenant* 1 is illustrated in Fig. 9(b). Here, the results are obtained by comparing the receiving bandwidth with the sending one at each time instant. It can be seen that at $t = 13$ seconds, the SL failure happens and starts to impact the video traffic, while only after 2.459 seconds, the service of video streaming is recovered. Here, the recovery latency includes the time used for detecting the SL failure, sending/receiving control messages, reconfiguring related BV-WSS', and resuming the upper-layer video streaming session. To verify that the video streaming service does get restored, we randomly select one from the 50 video streams and plot the luminance components peak signal-to-noise-ratio (Y-PSNR) of its playback in Fig. 9(c). The results on Y-PSNR confirm that the video's playback quality is recovered after ~ 3 seconds. Moreover, we take screen-shots for the video (*i.e.*, Rio 2 produced by the Blue Sky Studio [35]) before the SL failure, during the DP restoration, and after the DP restoration, and show them in Fig. 10 to demonstrate the playback quality explicitly. It can be seen that even though the SL failure can cause packet losses and degrade the video's playback quality, the service gets recovered to its original state after the DP restoration.

The receiving bandwidth of the file transfer in *Tenant* 2 is plotted in Fig. 9(d), which indicates that the data traffic is affected by the SL failure at $t = 13$ seconds and its service gets restored after 5.188 seconds. Here, the recovery latency includes the time used for detecting the SL failure, sending/receiving control messages, recalculating the VL mapping scheme, reconfiguring related BV-WSS', and resuming the upper-layer file transfer session. At last, we repeat aforementioned DP restoration experiments for 20 times and obtain the average recovery latencies for *Tenants* 1 and 2 as shown in Table 1. Note that, in a practical EON that covers a relatively large geographical area, the difference on the recovery latency would be larger, since the on-demand VL remapping would take longer time to distribute the control messages and have more BV-WSS' to configure. Moreover, as the backup paths are not reserved in advance, the VNMgr might not be able to obtain a feasible VL remapping scheme when the traffic load in the EON is relatively high, which would cause even longer recovery latency.

(a) (b) (c)

Fig. 10. Screen-shots of video playback, (a) Before failure (normal transmission), (b) during reconfiguration (packet losses and degradation), (c) after restoration (return to normal transmission).

Table 1. Average recovery latencies for vSD-EONs carrying different applications.

Tenant ID	DP Restoration Scheme	Recovery Latency (second)
1	“1:1” VL Protection	3.091
2	On-demand VL Remapping	4.887

5. Conclusion

In this paper, we designed and demonstrated the building and operating of QoS-aware survivable vSD-EONs that are equipped with transparent DP resiliency. We first designed and implemented the vSD-EON slicing system to realize automatic DP restoration that is transparent to the controllers of vSD-EONs. Then, we built a network testbed to experimentally demonstrate the creation of QoS-aware survivable vSD-EONs, the activation of lightpaths in the vSD-EONs to support two types of upper-layer applications, and the automatic and simultaneous QoS-aware DP restorations during an SL failure. The experimental results indicated that our vSD-EON slicing system can build QoS-aware survivable vSD-EONs on-demand, operate them to set up lightpaths for carrying real application traffic, and facilitate differentiated DP restorations during SL failures to recover the vSD-EONs’ services according to their SLAs.

Funding

NSFC Project 61371117; Natural Science Research Project for Universities in Anhui (KJ2014ZD38); The Strategic Priority Research Program of the Chinese Academy of Sciences (XDA06011202); The New Generation Broadband Wireless Mobile Communication Network Key Project under Grant No. 2017ZX03001019-004.

Acknowledgments

We would like to thank the anonymous reviewers for their valuable comments and suggestions.