OpenFlow Network Virtualization on Multiple Administration Infrastructures

Hiroaki Yamanaka Eiji Kawai Shuji Ishii Shinji Shimojo
National Institute of Information and Communications Technology

1 Introduction

Wide-area deployment of OpenFlow networks is expected. Since the capital expenditure (CAPEX) and the operational expenditure (OPEX) are large for wide-area network infrastructures (e.g., numbers of high-grade switches and long-haul links), multi-tenancy in wide-area OpenFlow networks is necessary. Multi-tenancy enables multiple tenants to use their exclusive virtual OpenFlow networks in a single OpenFlow infrastructure. In order to accept many tenants, virtualization of flow space (i.e., the header space to define flows) is necessary for virtual OpenFlow networks. Flow space virtualization allows tenants to use the overlapped flow spaces while flows in the physical OpenFlow networks are isolated. As a result, no agreement to divide the flow space is required for tenants.

In a flow space virtualization technique [1], a proxy, which is between OpenFlow switches and tenant controllers, translates (e.g., adding VLAN tag) the matching fields in OpenFlow protocol messages to isolate the flow space in the physical OpenFlow network. For virtualization in a data plane, the proxy controls the ingress OpenFlow switches to translates (e.g., adding VLAN tag) the headers of data packets from end-hosts of the tenants.

To build virtual OpenFlow networks on wide-area network infrastructures, a technique for flow space virtualization over multiple administration infrastructures is necessary. A straightforward approach is to deploy a global mapping management system over multiple administrators. However, it is costly and disadvantageous in terms of system scalability.

HyperFlow [2] firstly proposed a distributed implementation of an OpenFlow controller. The technique enables to control large-scale OpenFlow networks. However, the technique does not consider multi-tenancy. FlowN [3] proposed a duplicate database system that stores the global mapping of the header space for flow space virtualization. FlowN enables a virtualization on a large-scale physical OpenFlow network. However, FlowN did not assume the collaboration among multiple physical infrastructure administrators.

In this research, we propose a technique to build virtual OpenFlow networks by lightweight collaboration among multiple infrastructure administrators. Administrators individually manage the mapping of the header spaces between its own infrastructure and tenants. Since data packets are transferred between OpenFlow switches of different administrators, different physical infrastructure administrators need to share the mapping of the flow spaces for translation. In the proposed technique, necessary information for translation is exchanged between the administrators when data packets are transferred between different administration physical OpenFlow networks. The proposed collaboration scheme does not require any agree for the mapping of the header space and large communication overhead among multiple physical administrators.

2 Proposal

We describe a technique to build flat virtual OpenFlow networks on multiple administration physical OpenFlow networks. The proposed technique conceals the heterogeneity of administration in physical OpenFlow networks from tenants. Meanwhile, the technique does not require large communication overhead among different administrators.

2.1 System overview

Figure 1 shows the system overview. The physical OpenFlow infrastructure are composed of multiple administration physical OpenFlow networks. Different administration physical OpenFlow networks are connected by data plane links, and their proxies can communicate in a control plane. A tenant has a tenant controller and end-hosts.

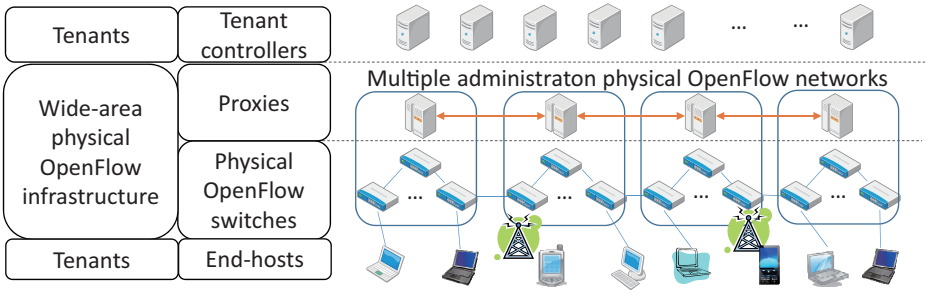

Figure 1: System overview.

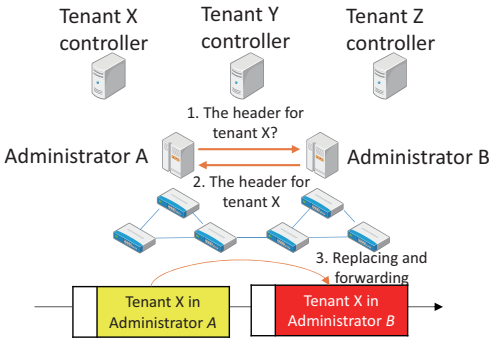

Figure 2: Data packet forwarding between different administration physical OpenFlow networks.

2.2 Implementation of collaborative flow space virtualization

A single administrator manages the local mapping of the header space in its physical OpenFlow network and in the tenants' view. According to the local mapping, a proxy translates OpenFlow protocol messages from tenant controllers, and data packet headers at the switch ports connecting to end-hosts.

When a data packet its header is replaced so that the proxy of the receiver administrator can handle the data packet without any consideration. As shown in Figure 2, the proxy of administrator A queries with the keys of the tenant ID and the header. Then, the other proxy of administrator B replies with the corresponding header in its physical OpenFlow network. The proxy of administrator A replaces the header with the replied value and forwards it. The proxy of administrator B can straightforwardly translate the header space according to its local mapping. Note that this query process is not necessary when the sender administrator has already known the corresponding header in the receiver administrator.

Furthermore, a proxy in an administration physical OpenFlow network can be implemented in a distributed and hierarchy manner by applying existing techniques: HyperFlow [2], FlowN [3], and Onix [4]. This improves the scalability of an administration physical OpenFlow network infrastructure.

2.3 Advantages of proposal

The proposed technique does not require any global mapping management system (e.g., a third party management system or distributed management system by administrators) for flow space virtualization. Instead, administrators can manage the mapping by its own management scheme. Furthermore, for administrators, the communication overhead with other administration systems is almost same as the communication overheard in case that a global mapping management system is used. Even if there is a global mapping system, an administrator needs to query the corresponding header in physical OpenFlow networks to handle data packets or OpenFlow protocol messages from tenants. We believe these characteristics of the proposed technique improve the scalability for physical infrastructure and acceptable tenants in OpenFlow network virtualization.

References

[1] E. Salvadori, R. Corin, A. Broglio, and M. Gerola, "Generalizing virtual network topologies in openflow-based networks," in *Global Telecommunications Conference (GLOBECOM 2011), 2011 IEEE*, 2011, pp. 1–6.

[2] A. Tootoonchian and Y. Ganjali, "HyperFlow: A distributed control plane for openflow," in *Proceedings of the 2010 internet network management conference on Research on enterprise networking (INM/WREN'10)*. USENIX Association, 2010.

[3] D. Drutskoy, E. Keller, and J. Rexford, "Scalable network virtualization in software-defined networks," *IEEE Internet Computing*, vol. 17, no. 2, pp. 20–27, 2012.

[4] T. Koponen, M. Casado, N. Gude, J. Stribling, L. Poutievski, M. Zhu, R. Ramanathan, Y. Iwata, H. Inoue, T. Hama, and S. Shenker, "Onix: A distributed control platform for large-scale production networks," in *Proceedings of the 9th Symposium on Operating Systems Design and Implementation (OSDI '10)*, 2010.