Inter-correlation of Resource-/Flow-Level Visibility for APM Over OF@TEIN SDN-Enabled Multi-site Cloud

Muhammad Usman(✉), Aris Cahyadi Risdianto, Jungsu Han, and JongWon Kim

Networked Computing Systems Laboratory (NetCS), EECS, GIST, Gwangju, South Korea
{usman, aris, jshan, jongwon}@nm.gist.ac.kr

Abstract. Cloud computing and SDN technologies have potential to bring advanced capabilities to data centers and inter-connecting networks by maximizing resources utilization. However, SDN/Cloud composition also brings monitoring and visibility challenges, for operators to diagnose p+vResources (server, network and storage), flows and cloud based applications for end-to-end performance management, because legacy tools lack adequate visibility support for such kind of dynamic environments. In this paper, by developing and integrating open source tools, we present a unified visibility solution to inter-correlate Resource-/Flow-level visibility over OF@TEIN to assist cloud-based APM (Application Performance Management). Furthermore, a verification environment has been setup, to show the feasibility of our inter-correlation approach.

Keywords: SDN · Cloud · Visibility · Visualization · DevOps automation

1 Introduction

Unlike traditional network-focused testbeds, Future Internet testbeds should provide experimental networking facility without the limitations of number of simultaneous users with varying resources requirements, number of supported services, types of applications and most importantly in the deployed network topology. Thanks, to the open and programmable nature of emerging SDN (Software-Defined Networking) paradigm that encouraged the construction and operation of SDN-enabled testbeds over international Research & Education networks. In 2012, we launched OF@TEIN project [1], aligned with Future Internet testbed projects like GENI and FIRE to build an SDN-enabled Multi-site Cloud testbed over TEIN (Trans-Eurasia Information Network). As of now, OF@TEIN connects 10 international sites spread across 9 countries (i.e. Korea, Malaysia, Thailand, Indonesia, India, Vietnam, Pakistan, Taiwan and Philippines) as shown in Fig. 1.

This SDN/Cloud integration over multiple international sites with diverse networks, and heterogeneous hardware/software compositions at OF@TEIN provides diverse resource combinations for developers. However, SDN/Cloud integrated environment at

J.-H. Lee and S. Pack (Eds.): QShine 2016, LNICST 199, pp. 478–484, 2017.
DOI: 10.1007/978-3-319-60717-7_48

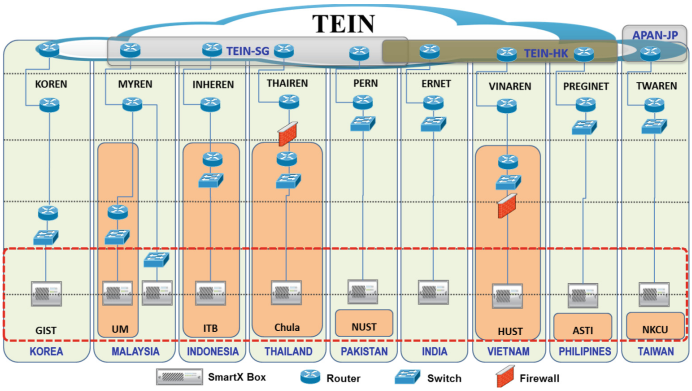

Fig. 1. OF@TEIN SDN-enabled multi-site cloud physical infrastructure

OF@TEIN also brings new challenges for operators like operational complexity and lack of visibility. Furthermore, OF@TEIN developer's also face APM issues; but they don't really know, where exactly problem is happening (i.e. application, p+vBox (physical/virtual), overlay/underlay network). As a result, it becomes quite challenging to provide a highly available and sustainable environment to OF@TEIN developers.

So, in order to deal with visibility challenges associated with cloud-based APM, we need further troubleshooting in multiple-levels of visibility (e.g. Resource-/Flow-/Service-level), to enable the operators with minimum efforts, to identify root cause of application performance degradation. Thus, to address the visibility challenges, in this paper, we focus on the inter-level correlation of Resource-level [2] and Flow-level [3] visibility, to provide adequate actionable information for APM to ease the job of operators. Furthermore, in our initial prototype, we verified the usability of operational data captured from multiple systems and identified missing data elements which can help to further optimize overall performance of OF@TEIN playground.

The rest of the paper is organized as follows. In Sect. 2, we discuss Multi-level Visibility requirements and related visibility solutions. In Sect. 3, we present our initial design and implementation details of Inter-level Correlator. In Sect. 4, we discuss verification environment. In Sect. 5, we conclude the paper.

2 Multi-level Visibility and Related Work

2.1 Multi-level Visibility Requirements and Support Tools

2.1.1 Multi-level Visibility Requirements

The OF@TEIN Multi-level Visibility framework should provide an innovative solution that should bridge Cloud, SDN and physical/virtual infrastructure for end-to-end visibility for automated troubleshooting. Furthermore, this, Multi-level visibility with centralized access must assure integration of independent and isolated operation data, captured from multiple visibility levels, from multiple sources. Thus, in this paper, we focused on leveraging multiple visibility solutions that have an integrated and synchronized awareness of Resource-/Flow-level visibility to assist APM.

2.1.2 Resource-/Flow-Level Visibility Tools

In order to fulfill Multi-level visibility requirements, we are developing SmartX Visibility Center as component solution of SmartX DevOpsTower as shown in Fig. 2. SmartX DevOpsTower is a centralized location for operators to fully monitor and control the operation of OF@TEIN playground. As an initial step towards Multi-level visibility we developed standalone Resource-level [2] and Flow-level [3] visibility tools to assist both Operators and Developers. OF@TEIN Resource-level visibility tools provide detailed information of p+vBox and status of inter-connecting links/paths. OF@TEIN flow-centric visibility solution covers flow monitoring, inspection and visualization capabilities by redirecting specific packets from OpenFlow-enabled switches to SmartX Visibility Server for further analysis.

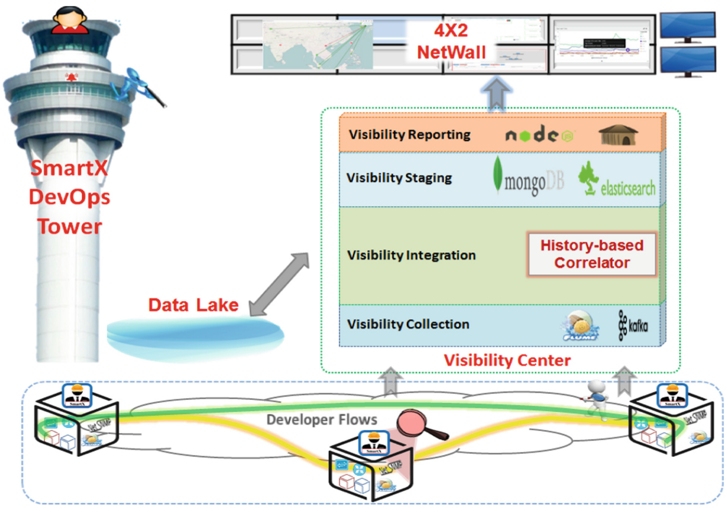

Fig. 2. SmartX DevOpsTower for enabling multi-level visibility support for OF@TEIN

2.2 Related Work

There are few, unified visibility solutions available that uses different approaches to collect, transform and inter-correlate visibility data. Gigamon Visibility Fabric [4] and ThousandEyes [5] are well-known visibility solutions that can process, inspect and filter data packets from internal and external networks by correlating metrics to optimize performance. Unfortunately, both Gigamon Visibility Fabric and ThousandEyes are commercial products and require dedicated hardware and specific software licenses which makes the Multi-level visibility solution too expensive to be implemented in open networking environment. ONUG (Open Networking User Group) [6] also focused its effort on what is needed to deliver an open network state collection, correlation and analytics service to the enterprise market without providing any details to how to achieve this task. Furthermore, most studies related to multi source data correlation, mainly focuses on security related aspects [7–9], typically for anomalies detection problem, instead of multi-level visibility aspect of the system.

3 Inter-level Correlation: Design and Implementation

3.1 Inter-level Correlation: Design

By leveraging available visibility tools, we are designing Inter-level Correlator for OF@TEIN to integrate resource and flow information at one place. The initial design of Inter-level Correlator is shown in Fig. 3. In the initial design, we focused on operational data verification, and defining keys for inter-level correlation. We defined five tuple (source/destination instance ip, source/destination TCP/UDP ports and protocol type), timestamp and OpenDaylight Controller IP as correlation-key. Based on correlation-key, Inter-level Correlator filters and extracts flows information from Data Lake and calculates critical flow stats. Then, Inter-level Correlator extracts OpenDayLight Configurations from Data Lake for finding associated Operation/Developer Flows, SDN topology information and related performance statistics. Followed by, instance configurations and performance data filtering based on UUID (Universally Unique Identifier) search-key from Data Lake. After that, Inter-level Correlator determines pBox where instance is deployed and extracts CPU, Memory and Disk performance data from Data Lake. Finally, this unified Resource-/Flow-level data is stored in Inter-level Data Lake and visualized.

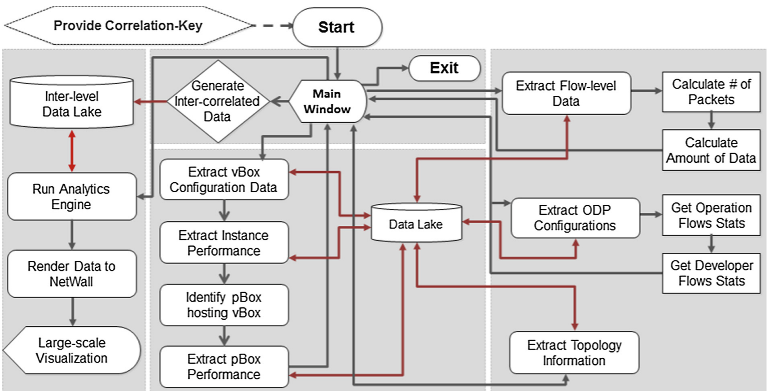

Fig. 3. Initial design of Inter-level correlator for OF@TEIN

3.2 Implementation of Inter-level Correlation

In this paper we focused on inter-level correlation of Resource-/Flow-level visibility for two hyper-convergent SmartX Boxes.

Inter-level Correlator: We implemented Inter-level Correlator by using Java. We also, extensively used OpenDayLight REST API's for extracting SDN Topology/Configurations information and OpenStack Nova REST API's for extracting instance configurations.

Inter-level Data Storage: For storing Inter-level Correlation data generated by Inter-level Correlator we used Elasticsearch [10] which is distributed, scalable and real-time search engine using index-based approach for data storage and retrieval.

Inter-level Data Visualization: Finally, this unified Inter-level Visibility data is visualized by using Kibana [10] Visualization Engine (by creating relevant visualizations and searches) over SAGE-enabled [2] NetWall.

4 Inter-level Correlation: Verification

To verify our unified visibility approach, we created verification environment with two OpenStack instances in two regions (GIST and ID). Followed by, assigning specific Operator and Developer Controller's to manage the SDN topology. Our simplified, verification environment is shown in Fig. 4.

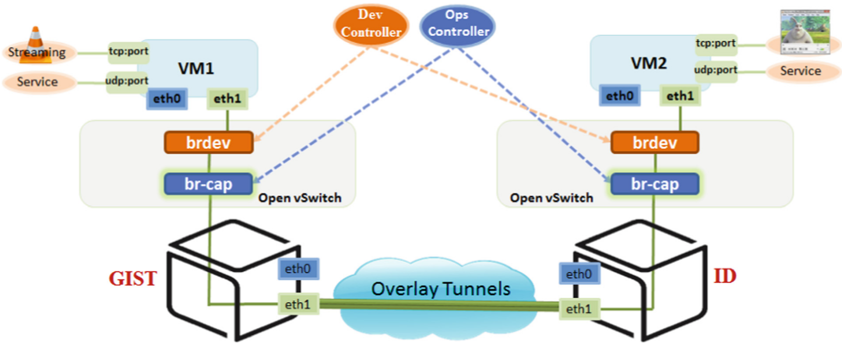

Fig. 4. Inter-level correlation verification environment for video streaming application

We performed video streaming service tests by selecting instance in GIST site as streaming server and instance in ID site as streaming client. During the execution of video streaming service, we observed low video quality problem. Application performance statistics showed that 105 frames were lost. So, after the execution of video streaming tests, we executed Inter-level Correlator by providing correlation-key for inter-correlating captured data, to figure out root cause of problem.

Inter-level Correlator, linked and filtered flows data, from multi sites. Inter-Correlator output shows, that numbers of packets were dropped in tunnel (e.g. number of packets captured from operator bridge brcap in GIST site was 1,999 but numbers of packets received in ID site operator bridge brcap were only 1,356). Then, Inter-level Correlator analyzed SDN topology over the period of time, but there were no changes, also no changes were recorded in Operation/Developer flows. Further, Inter-level Correlator filtered vBox configurations without any changes and pBox performance stats were also satisfactory. So, the problem found to be in tunnel network. In Fig. 5, some of simplified visualizations generated by Inter-level Correlator execution are shown. Similarly we also performed streaming tests by setting wrong configurations in brdev (Developer Bridge) in client site. Again Inter-level Correlator was able to identify that packets received up to brcap but cannot reach to OpenStack instance and there were some changes in brdev configurations over the period of time. Further, Inter-level Correlator has capabilities to figure out application problems due to p+vBox performance issues.

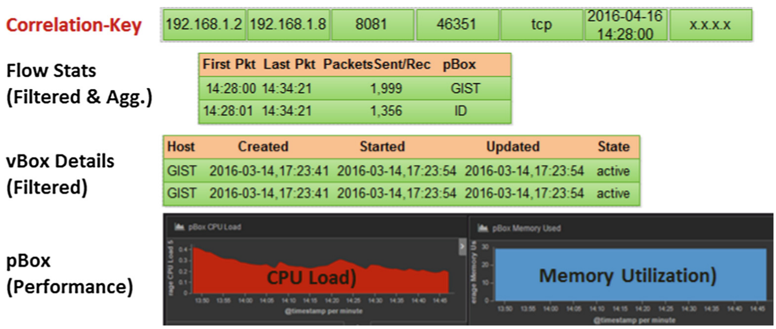

Fig. 5. Results of inter-level correlator execution

5 Conclusion and Future Work

In this paper, we have shown an early effort for inter-correlating Resource-/Flow-level visibility to assist cloud-based APM. In, the future, we are planning to extend our work towards multi-level visibility.

Acknowledgment. This work was supported in part by Institute for Information & communications Technology Promotion (IITP) grant funded by the Korea government (MSIP) (No. R7117-16-0218, Development of Automated SaaS Compatibility Techniques over Hybrid/Multisite Clouds). This work was also supported in part by Institute for Information & communications Technology Promotion (IITP) grant funded by the Korea government (MSIP) (No. B190-15-2012, Global SDN/NFV Open-Source Software Core Module/Function Development).

References

1. Risdianto, A.C., Kim, J.: Prototyping media distribution experiments over OF@TEIN SDN-enabled testbed. APANNRW **38**, 12–18 (2014). Nantou
2. Usman, M., Risdianto, A.C., Kim, J.: Resource monitoring and visualization for OF@TEIN SDN-enabled multi-site cloud. In: ICOIN, Kota Kinablu, pp. 427–429 (2016)

3. Risdianto, A.C., Kim, J.: Flow-centric visibility tools for OF@TEIN OpenFlow-based SDN testbed. In: CFI, Seoul, pp. 46–50 (2015)
4. Gigamon unified visibility fabric. https://www.gigamon.com/best-practices
5. ThousandEyes. https://www.thousandeyes.com/
6. Network State Collection, Correlation and Analytics Product/RFI Requirements. White paper, Open Networking User Group (2015)
7. Pierazzi, F., Casolari, S., Colajanni, M., Marchetti, M.: Exploratory security analytics for anomaly detection. J. Comput. Secur. **56**, 28–49 (2015)
8. Jagadeesan, L., Bride, A.M., Gurbani, V.K., Yang, J.: Cognitive security: security analytics and autonomics for virtualized networks. In: IPTComm, Chicago, pp. 43–50 (2015)
9. Cao, P., Badger, E.C., Kalbarczyk, Z.T., Iyer, R.K., Withers, A., Slagell, A,J.: Towards an unified security testbed and security analytics framework. In: HostSoS, Illinois, pp. 24–25 (2015)
10. ELK Stack. https://www.elastic.co/webinars/introduction-elk-stack