

Towards an Accountable Software-Defined Networking Architecture

Benjamin E. Ujcich^{*†}, Andrew Miller^{*†}, Adam Bates[‡], and William H. Sanders^{*†}

^{*}Information Trust Institute, [†]Department of Electrical and Computer Engineering, [‡]Department of Computer Science
University of Illinois at Urbana–Champaign
Urbana, Illinois USA

Email: {ujcich2, soc1024, batesa, whs}@illinois.edu

Abstract—Software-defined networking (SDN) overcomes many limitations of traditional networking architectures because of its programmable and flexible nature. Security applications, for instance, can dynamically reprogram a network to respond to ongoing threats in real time. However, the same flexibility also creates risk, since it can be used against the network. Current SDN architectures potentially allow adversaries to disrupt one or more SDN system components and to hide their actions in doing so. That makes assurance and reasoning about past network events more difficult, if not impossible. In this paper, we argue that an SDN architecture must incorporate various notions of accountability for achieving systemwide cyber resiliency goals. We analyze accountability based on a conceptual framework, and we identify how that analysis fits in with the SDN architecture’s entities and processes. We further consider a case study in which accountability is necessary for SDN network applications, and we discuss the limits of current approaches.

I. INTRODUCTION

Software-defined networking (SDN) has emerged as a new networking architecture that attempts to overcome some of the limitations of traditional networking. SDN is distinguished by a logically centralized but physically distributed programmable control plane in which decisions about forwarding are decoupled from the traffic being forwarded [1]. This flexibility has encouraged SDN adoption in enterprise, campus, cloud, mobile, and telecommunication networks, among others [1].

At first glance, SDN enhances the ability of security services to protect the network’s end hosts from threats. The architecture’s global perspective [2] allows the control plane to monitor traffic flows and abstract such information for network security applications’ use. Further, the architecture’s programmable nature allows administrators or security applications to rapidly reconfigure the network’s forwarding behavior to adjust to threats in real time.

However, such insight and flexibility are not without costs, as the increased attack surface and centralized programmatic control could be used against operators and administrators. If we assume that an attacker has already compromised an SDN system’s components to control network behavior and that the attacker lies, equivocates, and hides or destroys evidence of such actions [3], how can we be assured of past event integrity? Furthermore, how can we trust the provenance of such events, attribute them to the responsible entities, and take actions

against them to hold them responsible and meet systemwide cyber resiliency goals?

In this paper, we argue that *accountability by design* [4], [5] is necessary for SDN. As SDN architectures become increasingly complex distributed network operating systems [6], we see a need for ensuring accountable practices at multiple levels among various entities and stakeholders. Achieving accountability is not strictly a technical problem—legal, regulatory, and policy frameworks all help guide accountable systems design—but accountability requires data assurances to correctly identify responsible entities when taking responsive actions to support resiliency goals.

Our paper’s contributions include application of a conceptual framework of accountability [7] to SDN to identify the agents, system entities, processes, and standards involved in network accountability assurance. We find that while previous work has considered some of these aspects, no complete design solution exists today that systematically incorporates all of them. We illustrate the need for accountability through a practical case study scenario of SDN network applications.

II. SDN BACKGROUND

SDN separates a network into a *data plane* where traffic among end hosts is forwarded through *switches*, a *control plane* where forwarding decisions are made by a set of *controllers*, and an *application plane* where *network applications* can query or set high-level *intents* about network state.

Controllers communicate via application program interfaces (APIs). The *northbound API* allows network applications to query controllers for network state abstractions or set intents about network policy. The *southbound API* allows controllers to set low-level forwarding rules in switches and to query about network state. The *eastbound/westbound API* allows controllers to communicate among themselves or with other SDN systems to exchange distributed state.

SDN architectures introduce new and augmented security and dependability vulnerabilities, since controllers logically centralize control, and network applications programmatically set policies [8]. Kreutz et al. [8] cite the need for a reliable SDN forensics system to diagnose faults (attacks). Scott-Hayward et al. [9] consider accountability as a system-level SDN security challenge within the authentication, authorization, and accounting (AAA) framework.

TABLE I
SUMMARY OF DESIGN CONSIDERATIONS AND PROPERTIES FOR SOFTWARE-DEFINED NETWORKING ACCOUNTABILITY

Who is accountable to whom	What one is accountable for	Assurance mechanisms	Standards	Effects of breach
<ul style="list-style-type: none"> Software process level <ul style="list-style-type: none"> – Switch–switch – Controller–switch – Controller–application – Controller–controller User level <ul style="list-style-type: none"> – Network administrators – Security administrators – End users Organizational level <ul style="list-style-type: none"> – Clients–providers – Peers 	<ul style="list-style-type: none"> Forwarding / topology Intent / policy <ul style="list-style-type: none"> – Network resources – Constraints – Criteria – Instructions Configuration Authorization / access <ul style="list-style-type: none"> – Permissions and roles – Authentication and access 	<ul style="list-style-type: none"> Data provenance Authenticated logging <ul style="list-style-type: none"> – Tamper-proof – Non-repudiable Fault tolerance <ul style="list-style-type: none"> – Byzantine fault tolerance – Graphical modeling – Blockchains Roots of trust 	<ul style="list-style-type: none"> Legal Regulatory Policy Contractual 	<ul style="list-style-type: none"> Deterrence <ul style="list-style-type: none"> – Loss of money – Loss of reputation Resiliency <ul style="list-style-type: none"> – Response – Recovery

III. DESIGNING AN ACCOUNTABLE SDN ARCHITECTURE

Accountability is the “security goal that generates the requirement for actions of an entity to be traced uniquely to that entity” and which supports “nonrepudiation, deterrence, fault isolation, intrusion detection and prevention, and after-action recovery and legal action” [10]. We borrow concepts from the public policy domain, in particular the accountability framework for designing accountable systems proposed by Mashaw [7], to analyze accountability as it applies to SDN architectures and entities. Table I summarizes the design considerations and desirable properties for SDN accountability.

A. Who is accountable to whom?

We follow Mashaw in describing accountability as relationships between entities. To say “A is accountable to B” means that A behaves according to processes guided by standards by which B can correctly attribute A’s actions and take responses against A if A deviates from such processes and standards. (The meanings of *processes* and *standards* are explained in detail throughout the remainder of Section III.)

We organize those relationships at three levels: software process, user, and organizational accountability.

1) *Software process level*: SDN software components should keep each other accountable for low-level network state changes so as to attribute actions to particular software instances for troubleshooting (e.g., fault isolation) or for forensics. We identify the following accountability-critical classes of inter-process relationships:

a) *Switch–switch*: Switches should keep each other accountable for their data plane forwarding actions. In particular, they should ensure that packets traverse the correct switches to enforce isolation guarantees and forwarding accountability [11].

b) *Controller–switch*: Controllers should keep switches accountable for their actions to ensure that network intents are followed. For instance, switches should attest to their current forwarding behavior state and report it to controllers.

c) *Controller–application*: Controllers should keep network applications accountable to ensure that conflicting policies are mitigated according to a permissions model [12]. To ensure trustworthy network applications [8], it is necessary to

hold network applications accountable for actions they take that affect the network state [12]. Application developers and publishers should be held accountable, too.

d) *Controller–controller*: In contrast to a single centralized controller, distributed controllers should provide high availability, scalability, and fault tolerance properties [6]. Distributed controller instances share copies of the network state and may act as clients in reading from a distributed data store [13], [14]. Given that the network’s intelligence is logically centralized in the controllers, they should keep each other accountable for network state changes.

2) *User level*: Network and security administrators should keep each other accountable for decisions that affect network state, particularly if the administrators have the potential to collude or are assumed to be individually untrusted [15]. Administrators should keep the network’s end users accountable for their actions on the network, such as equitably sharing network resources based on policy.

3) *Organizational level*: Organizations should keep other organizations accountable for network resources when considering client–provider or peer models. In cloud computing, for instance, a provider should use a telemetry service to account for network resources used by clients, and clients should be able to audit providers to ensure that the services requested are being provided in practice [16]. For an autonomous system (AS), each AS should make other ASes accountable for their Border Gateway Protocol (BGP) advertisements or for the inter-AS network resources (e.g., bandwidth) they use related to peering agreements.

B. What is one accountable for?

An accountable architecture accounts for the system’s “state” and state changes via events or actions taken by system entities. Here, we identify several notions of state.

1) *Forwarding behavior and topology*: From the data plane perspective, the network’s state consists of the forwarding behavior (e.g., flow table entries) and the topology (e.g., ports, links, switches, hosts). In OpenFlow-based SDNs, the forwarding behavior is defined by flow tables that consist of flow entries with matching attributes and a set of actions to take for matching packets [17]. The topological information

is based on switch configuration [17] and also from auxiliary protocols such as the Link Layer Discovery Protocol (LLDP) and the Address Resolution Protocol (ARP) [18].

2) *Intent and policy*: From the application plane perspective, the network's state consists of the policies implemented by intents. ONOS [6], for instance, defines intents by network resources (e.g., ports, links), constraints (e.g., bandwidth), criteria (e.g., matching headers), and instructions (e.g., header modifications, output). Intents contrast with specific protocols like OpenFlow [17] by abstracting the implementation details.

3) *Configuration*: From a system administrator's perspective, each network component requires configuration. OF-Config [19] configures switches, changes port states, and changes security certificates, among other functions.

4) *Authorization and access*: From a security administrator's perspective, the network requires a system for authorizing users' actions based on roles and permissions. Such a system needs to record state modification, permissions, and authentication and access events, among other records [12].

C. What process assures accountability mechanisms?

We now consider four necessary classes of assurance mechanisms, including their uses to date in SDN and other fields.

1) *Data provenance*: Arguably, the most important property of data assurances for accountability is their ability to answer questions about where data came from and why the data came to be [20]. Data provenance answers these questions by attributing data to their sources in order to support audit trail generation [21]. Data provenance has been used in database systems [20] and in distributed systems for identifying which system components took specific actions [22]. Dwaraki et al. [23] model SDN forwarding state changes through a distributed version control system to answer provenance queries about network state, though the architecture does not assume an adversarial setting.

2) *Authenticated logging*: While data provenance explains data origins, an adversarial setting requires assurance that the stored data cannot be tampered with. Authenticated data structures (e.g., Merkle trees) provide a way of implementing tamper-proof logs [24]. Each log entry is associated with a cryptographic hash, and tampering with previous log entries makes tampering evident. Among different entities, each entity can digitally sign entries it makes to the log and thus cannot repudiate previous entries it has signed. Porras et al. [12] extend the Floodlight SDN controller to include an application audit module for associating logged events with their sources.

3) *Fault tolerance*: Accountability does not begin only after the system as a whole has already failed. Many systems are designed to tolerate (or mask) failed components. When some component fails, the failed component should be held accountable for its actions and for events attributed to it.

Byzantine fault tolerant (BFT) protocols are a practical way to make many systems more robust, and the security they offer is evaluated by the strength of their guarantees (i.e., the weakness of their assumptions). PBFT [25], for instance, guarantees that a network of $N = 3f + 1$ replicas can tolerate up to

f corrupted instances that behave arbitrarily or maliciously. Furthermore, most BFT protocols guarantee liveness and high availability under very weak assumptions about the ability of the uncorrupted nodes to communicate; they also guarantee consistency even in a completely asynchronous network.

While the system as a whole should exhibit some degree of fault tolerance (as previously explored in the SDN context in [13], [14]), the subsystems used to ensure accountability should be especially fault-tolerant. BFT protocols must typically rely on at least a majority of the nodes to ensure safety and liveness, though secure network provenance (SNP) can rely on an even weaker assumption [26]. SNP uses a provenance graph to capture events in a distributed system, and minimally requires only one correct node to have witnessed an event in order to attribute it. Cryptocurrencies have recently popularized the use of widely distributed BFT protocols to provide a transparent and publicly verifiable transaction log known as a “blockchain.” Blockchain updates are relatively expensive and slow, but this trade-off may be appropriate for accountability-critical information.

4) *Roots of trust*: By adopting accountability as an explicit design goal, we strive to reduce the amount of trust in the system. However, no design is perfect, and we believe practical architectures will still require some “roots of trust” upon which the system’s security relies. Explicit descriptions of these roots of trust will be essential to evaluating accountability designs. To validate a design, we must identify the trusted entities and justify their trustworthiness. In the SDN context, for instance, Jacquin et al. [27] propose a trusted SDN architecture by placing trust in trusted platform modules (TPMs).

D. By what standards should accountability be judged?

Given that the network infrastructure is a central component of many institutions, and that it can “see” everything (including sensitive data) [2], an accountable SDN may be necessary in practice for meeting or aiding legal, regulatory, policy, or contractual requirements. In this context, standards set the accountability requirements that must be met.

Accountability standards are derived through laws, regulations, and policies. In the U.S., for instance, federal laws and regulations set domain-specific accountability standards as they apply to health records (e.g., HIPAA), educational records (e.g., FERPA), and financial records (e.g., Gramm-Leach-Bliley Act), among other domains.

We can also apply accountability in the context of other policies, such as network neutrality. The Council of the European Union recently passed network neutrality regulations for European Union member states, noting that “a significant number of end-users are affected by traffic management practices which block or slow down specific applications or services” [28]. Here, accountability includes customers and regulators who keep network providers accountable for their network management practices, as the regulation states that “reasonable traffic management measures... should be transparent, non-discriminatory and proportionate, and should not be based on commercial considerations” [28]. Accountable

designs can help ensure compliance with these regulations and can support their enforcement.

Outside of established legal regulations, any two parties can decide to enter into contractual service level agreements (SLA) regarding network resources, and the agreement's terms can set the standards that determine which entities and processes assure accountability and the effects of breaching the standards.

E. What are the effects of breaching standards?

Accountability can provide a natural deterrent against some classes of attacks, and can therefore have a passive effect of helping parties conform to agreed-upon standards. However, accountability can also play a more active role in system resiliency by supporting responses that restore the system to correct function after a failure.

1) Deterrence: An accountable SDN architecture may provide a disincentive to attack the network or deviate from agreed-upon rules, as such an attack or deviation could be attributed to the responsible entity [29]. As a result, the responsible entity has something to lose if it had previously pledged something of value in order to participate [30]. The loss may be monetary, as detailed by an agreed-upon SLA or smart contract [34]. Alternatively, the loss may be implicit and reputational, such that other entities choose not to participate with the responsible entity after discovery [30].

2) Resiliency: Such deterrence alone may not provide enough disincentive to stop an attacker from attacking a network, and in such cases, intrusion tolerance designed to meet system resiliency goals is necessary. The resiliency process is often defined as comprising detection, response, and recovery phases. Accountability clearly plays a role in detection, but it also provides essential information for effective responses.

Each entity may audit other entities to identify misbehavior and attribute it to the responsible entity. Upon detection, an entity may report to other entities to indicate that they should take some response action. For instance, a misbehaving entity might be isolated by peer entities so as to allow for human intervention (e.g., forensic analysis by a security administrator) while still meeting system service goals. Finally, for recovery, two existing mechanisms for SDN include partial configuration rollback [31] and elastic controller provisioning [32].

IV. CASE STUDY: ACCOUNTABLE SDN APPLICATIONS

We now consider a case study of applying the accountability process to SDN network applications.

A. Scenario

As controller software has become increasingly complex, there has been a proliferation of available network applications (apps) that network and security administrators can deploy. HP Enterprise, for instance, offers an SDN App Store [33] where users can download monitoring, security, optimization, orchestration, and visualization tools that coordinate with the HP VAN SDN controller. At the time of this writing, 35 of the 39 apps are provided for free, and 29 of the 39 apps were developed by third parties outside of HP Enterprise [33].

A natural security question arises: how can we trust network applications? Furthermore, how can we attribute actions that they take? Consider a simple example of three systemwide network applications used by a cloud provider: 1) an intrusion detection system (IDS) app monitoring all external-bound data plane traffic for intrusions, 2) a quality of service (QoS) app supporting network SLAs between the cloud provider and its clients, and 3) a firewall app protecting the data plane from external threats and isolating inter-client traffic.

Suppose the IDS app discovers a potential intrusion with systemic consequences in one of the client's resources, and the "discovery" is later determined to be a false positive. While the potential intrusion is still considered a real threat, the firewall receives the IDS alert and reactively reconfigures the network to block traffic that affects other clients' network resources. The QoS app determines it cannot provide any routing paths that support other clients' SLAs now that particular routes have been blocked, and the agreed terms from the SLAs are thus breached as a result of a nonexistent threat.

Each app, viewed independently, provided its respective services correctly, yet the actions taken on behalf of one client negatively affected other clients. Which entities should be held accountable for breaching the SLAs—the client whose resources caused the alert to be generated, the cloud provider whose security policies required IDS monitoring, the apps' developers whose software generated alerts and actuated the responses that breached the SLAs, or a combination thereof?

Porras et al. [12] consider the application coexistence problem as one in which multiple applications compete to make decisions that affect network behavior. They propose a mediation policy that includes minimal permissions levels, and they suggest implementing application accountability through a security audit service module. While their auditing solution attributes events to the applications that generated them, we posit that this is only one component in the accountability process. How the data can be used afterward to provide attribution, how the collected data relate to each other, and how the data can be used to enforce automatic penalties for the accountable entities must also be considered.

B. Analysis

Our framework, described previously in Section III, moves accountability from the view of what data are collected to a complete end-to-end view of how those data can be used to assign attribution and drive responsive decisions automatically. We highlight parts of the three-application scenario in which the conceptual framework of accountability can help in the SDN architecture design process.

1) Provider and clients: The cloud provider is accountable to its clients for providing acceptable levels of network service (e.g., bandwidth, latency, denial of service protections) as agreed upon contractually in an SLA. Its clients should be able to audit the provider to ensure that the service is provided in practice [16]. The provider and clients can agree on culpability when service levels are breached, and they can codify this logic with smart contracts that include monetary stakes. If a

client determines that service is not being provided, the client can receive monetary compensation [34].

2) *Controllers and apps*: The three apps are accountable to the SDN controllers for the high-level intents that they ask the controllers to implement. The SDN controllers are accountable to the apps for the low-level actions they take in response, as apps may require assurance that certain actions were taken or to provide evidence if disputes arise later. For instance, if the QoS app cannot change the network's routing to meet the SLA because of an action that the firewall app requested, the QoS app can use evidence of the firewall app's actions to declare its innocence as the root cause of the SLA failure. (The firewall app may do the same to declare its innocence as the root cause of failure vis-à-vis the IDS app.)

C. Remarks

Based on our accountability framework's considerations, SDN has made some progress in provenance [23], secure auditing [12], fault tolerance [13], [14], roots of trust [27], and resiliency [31], [32]. However, no one design captures *all* of the elements required for an accountable architecture. As illustrated in our case study, no solution to date has considered accountability as an end-to-end process, starting from assured data guarantees, continuing with auditing and detection of breaches, and ending with automated actions for deterrence and response that support resiliency goals.

V. CONCLUSION

SDN continues to be applied in a multitude of enterprises and domains [1], and its global perspective [2] can aid in providing detection and response mechanisms for systemwide cyber resiliency. In this paper, we argued that the security property of accountability must be considered in the architecture design so as to support detection assurances that ultimately inform responses. We provided a conceptual framework analysis of accountability as applied to SDN entities and processes, and we applied several notions of accountability in our network application case study. We hope that this paper spurs further interest in incorporating accountable networking by design.

ACKNOWLEDGMENT

The authors would like to thank Jenny Applequist for her editorial assistance and members of the PERFORM Group at the University of Illinois for feedback. This material is based upon work supported by the Army Research Office under Award No. W911NF-13-1-0086 and by the National Science Foundation under Grant No. CNS-1657534.

REFERENCES

- [1] D. Kreutz, F. Ramos, P. Veríssimo, C. Rothenberg, S. Azodolmolky, and S. Uhlig, "Software-defined networking: A comprehensive survey," in *Proceedings of the IEEE*, vol. 103, no. 1, Jan. 2015.
- [2] A. Bates, K. Butler, A. Haeberlen, M. Sherr, and W. Zhou, "Let SDN be your eyes: Secure forensics in data center networks," in *NDSS SENT*, 2014.
- [3] C. Röpke and T. Holz, "SDN rootkits: Subverting network operating systems of software-defined networks," in *RAID*, 2015.
- [4] A. R. Yumerefendi and J. S. Chase, "Trust but verify: Accountability for network services," in *ACM SIGOPS European Workshop*, 2004.
- [5] ———, "The role of accountability in dependable distributed systems," in *USENIX HotDep*, 2005.
- [6] P. Berde, M. Gerola, J. Hart, Y. Higuchi, M. Kobayashi, T. Koide, B. Lantz, B. O'Connor, P. Radoslavov, W. Snow, and G. Parulkar, "ONOS: Towards an open, distributed SDN OS," in *ACM HotSDN*, 2014.
- [7] J. L. Mashaw, "Accountability and institutional design: Some thoughts on the grammar of governance," in *Public Accountability: Designs, Dilemmas, and Experience*, M. W. Dowdle, Ed. Cambridge University Press, 2006, pp. 115–156.
- [8] D. Kreutz, F. M. Ramos, and P. Veríssimo, "Towards secure and dependable software-defined networks," in *ACM HotSDN*, 2013.
- [9] S. Scott-Hayward, S. Natarajan, and S. Sezer, "A survey of security in software defined networks," *IEEE Communications Surveys Tutorials*, vol. 18, no. 1, pp. 623–654, 2016.
- [10] R. Kissel, "Glossary of key information security terms," National Institute of Standards and Technology, Tech. Rep. NISTIR 7298, May 2013.
- [11] T. Sasaki, C. Pappas, T. Lee, T. Hoefer, and A. Perrig, "SDNsec: Forwarding accountability for the SDN data plane," in *IEEE ICCN*, 2016.
- [12] P. Porras, S. Cheung, M. Fong, K. Skinner, and V. Yegneswaran, "Securing the software-defined network control layer," in *NDSS*, 2015.
- [13] F. Botelho, A. Bessani, F. M. V. Ramos, and P. Ferreira, "On the design of practical fault-tolerant SDN controllers," in *IEEE EWSN*, 2014.
- [14] N. Katta, H. Zhang, M. Freedman, and J. Rexford, "Ravana: Controller fault-tolerance in software-defined networking," in *ACM SOSR*, 2015.
- [15] S. Matsumoto, S. Hitz, and A. Perrig, "Fleet: Defending SDNs from malicious administrators," in *ACM HotSDN*, 2014.
- [16] A. Haeberlen, "A case for the accountable cloud," *ACM SIGOPS Operating Systems Review*, vol. 44, no. 2, pp. 52–57, Apr. 2010.
- [17] Open Networking Foundation. OpenFlow v1.3.0. [Online]. Available: <https://www.opennetworking.org/images/stories/downloads/sdn-resources/onf-specifications/openflow/openflow-spec-v1.3.0.pdf>
- [18] Big Switch Networks, "Project Floodlight: Open source software for building software-defined networks," Jan. 2016. [Online]. Available: <http://www.projectfloodlight.org/floodlight/>
- [19] Open Networking Foundation. OF-CONFIG 1.2. [Online]. Available: <https://www.opennetworking.org/images/stories/downloads/sdn-resources/onf-specifications/openflow-config/of-config-1.2.pdf>
- [20] P. Buneman, S. Khanna, and W. C. Tan, "Why and where: A characterization of data provenance," in *ICDT*, 2001.
- [21] Y. L. Simmhan, B. Plale, and D. Gannon, "A survey of data provenance in e-science," *ACM SIGMOD Record*, vol. 34, no. 3, Sep. 2005.
- [22] A. Haeberlen, S. Kouznetsov, and P. Druschel, "PeerReview: Practical accountability for distributed systems," in *ACM SOSR*, 2007.
- [23] A. Dwaraki, S. Seetharaman, S. Natarajan, and T. Wolf, "GitFlow: Flow revision management for software-defined networks," in *ACM SOSR*, 2015.
- [24] S. A. Crosby and D. S. Wallach, "Efficient data structures for tamper-evident logging," in *USENIX SSYM*, 2009.
- [25] M. Castro and B. Liskov, "Practical byzantine fault tolerance," in *USENIX OSDI*, 1999.
- [26] W. Zhou, Q. Fei, A. Narayan, A. Haeberlen, B. T. Loo, and M. Sherr, "Secure network provenance," in *ACM SOSR*, 2011.
- [27] L. Jacquin, A. L. Shaw, and C. Dalton, "Towards trusted software-defined networks using a hardware-based integrity measurement architecture," in *IEEE NetSoft*, 2015.
- [28] Council of the European Union, "Regulation (EU) 2015/2120," in *Official Journal of the European Union*, 2015.
- [29] F. B. Schneider, "Accountability for perfection," *IEEE Security and Privacy*, vol. 7, no. 2, Mar. 2009.
- [30] C. E. Landwehr, "A national goal for cyberspace: Create an open, accountable Internet," *IEEE Security and Privacy*, vol. 7, no. 3, May 2009.
- [31] Y. Zhang, N. Beheshti, and R. Manghirmalani, "NetRevert: Rollback recovery in SDN," in *ACM HotSDN*, 2014.
- [32] A. Dixit, F. Hao, S. Mukherjee, T. Lakshman, and R. Kompella, "Towards an elastic distributed SDN controller," in *ACM HotSDN*, 2013.
- [33] Hewlett-Packard Enterprise. HPE SDN app store. (Accessed 7 Jan. 2017). [Online]. Available: <https://marketplace.saas.hpe.com/sdn>
- [34] Ethereum Project. Ethereum white paper: A next-generation smart contract and decentralized application platform. (Accessed 7 Jan. 2017). [Online]. Available: <https://github.com/ethereum/wiki/wiki/White-Paper>