


Taxonomic Modeling of Security Threats in Software Defined Networking

Recent advances in software defined networking (SDN) provide an opportunity to create flexible and secure next-generation networks. Many companies have expressed the interest in SDN utilization. Although much has been said about the ability of SDN to solve persistent network security problems, our current knowledge on SDN vulnerabilities, threats, and attacks is limited. In this paper, I use the threat modeling approach to develop a novel SDN threat model that provides a foundation for identifying possible threats to SDNs.



Contents

1	Introduction	3
2	SDN Attack Surface	3
3	SDN Threat Model	5
4	Attacks Examples	8
4.1	Unauthorized Access Using Password Brute-Forcing or Password-Guessing Attacks.....	8
4.2	Unauthorized Access Using Remote Application Exploitation Attacks	9
4.3	Unauthorized Disclosure of Information Using RAM Scraping Attacks	9
4.4	Unauthorized Disclosure of Information Using API Exploitation Attacks	10
4.5	Unauthorized Destruction Using API Exploitation Attacks.....	11
4.6	Unauthorized Access Using Remote or Local Application Exploitation Attacks	11
4.7	Unauthorized Disclosure of Information Using Traffic Sniffing Attacks	12
4.8	Unauthorized Modification Using Identity Spoofing Attacks	12
4.9	Disruption of Service Using Remote Application Exploitation Attacks.....	13
4.10	Unauthorized Disclosure of Information Using Side Channel Attacks	13
4.11	Disruption of Service Using Flooding Attacks	14
4.12	Unauthorized Modification Using Data Forging Attacks	14
5	SDN Threat Mitigation.....	15
6	Conclusion	16

1 Introduction

Software Define Networking (SDN) is an emerging networking paradigm that aims to change the limitations of the traditional networking. The value of SDN lies in its ability to provide consistent policy enforcement and deliver greater scalability and control over the entire network by means of centralized management and network programmability. The next generation of security solutions will take advantage of the wealth of network usage information available in SDN to improve policy enforcement and traffic anomaly detection and mitigation.

Although much has been said about the benefits of SDN to solve persistent network security problems, our current knowledge on SDN threats and attacks is limited. The new systems required to carry out SDN functions may themselves come under malicious attacks. While some attacks will be common to existing networks, others will be more specific to SDN. Adversaries will inevitably exploit SDN systems if a successful network compromise could be achieved through such exploitations. A vulnerable SDN network could therefore undermine security and availability of the entire network.

In order to secure SDN networks, all of the potential security threats and attacks must be anticipated before attackers take advantage of vulnerabilities. In this paper, I describe a novel SDN threat model, which looks at SDN from an adversary's perspective to identify potential threats to and attacks on SDN at the architectural level, regardless of whether or not these threats can be successfully carried out. I further apply the threat model to synthesize a set of potential real-life attacks adversaries could launch against SDN networks. Finally, I provide security recommendations to address the threats.

2 SDN Attack Surface

The SDN threat modeling requires understanding of the SDN architectural components and their interconnections. Figure 1 illustrates a typical SDN architecture to reveal the main SDN building blocks, which include the SDN planes and interfaces. Any of these architectural components could contain vulnerabilities and be exploited by attackers to compromise a SDN network. The overall SDN attack surface is represented by the sum of all of these components.

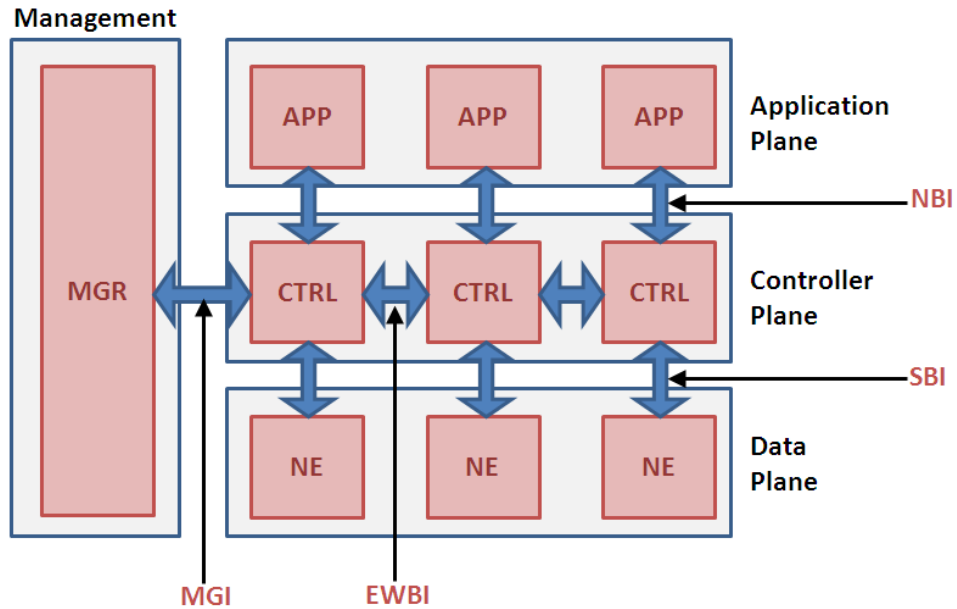


Figure 1. The SDN components include applications (APP), controllers (CTRL), network elements (NE), management consoles (MGR), the northbound interface (NBI), the southbound interface (SBI), east/west bound interface (EWBI), and the management interface (MGI).

The following planes are identified in SDN:

- the data plane comprising network elements for traffic forwarding or processing
- the controller plane comprising a set of controllers, which control network elements in the data plane
- the application plane comprising applications with access to resources exposed by controllers
- the management plane comprising management consoles for applications, controllers, and network elements and supporting remote management tasks.

The following interfaces are identified in SDN:

- the east/west bound interfaces required by distributed controllers for importing/exporting data between controllers and monitoring/notification capabilities
- the northbound interfaces enabling the communication between the controller plane and the application plane

- the southbound interfaces providing the link between the controller plane and the data plane
- management interfaces performing management functions on applications, controllers, and network elements in each plane.

3 SDN Threat Model

In order to systematically identify threats affecting SDN, I decompose each threat into 3 elements:

- threat source – a source triggering the vulnerability
- vulnerability source – a SDN component where the vulnerability arises
- threat action – an action by which the threat is carried out

I classify the threat sources into the following (see Figure 2):

- a non-SDN component – system that is not a part of the SDN architecture
- a rogue SDN component – an unauthorized SDN system within a SDN network engaged in unauthorized activities
- a malicious SDN application (a compromised application or a user engaged in malicious activities using the application)
- a malicious controller (a compromised controller or a user engaged in malicious activities on the controller)
- a malicious network element (a compromised network element or a user engaged in malicious activities using the network element)
- a malicious management console (a compromised management console or a user engaged in malicious activities using the console)

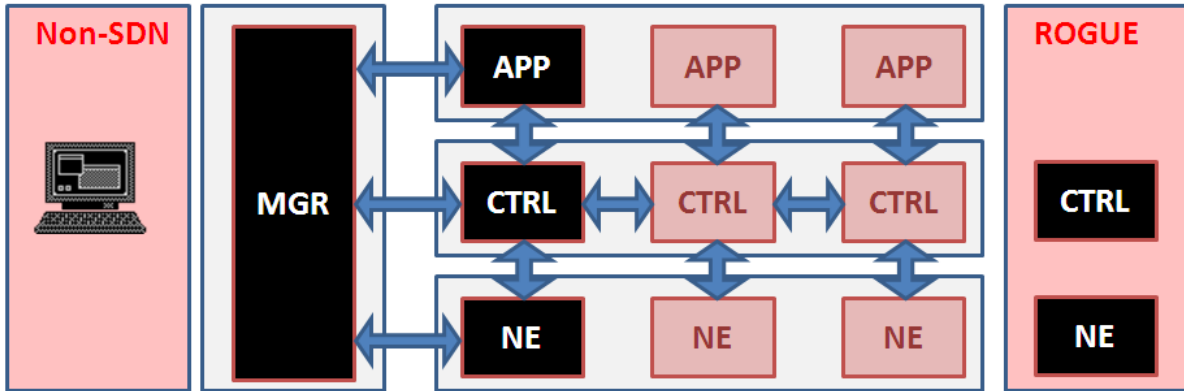


Figure 2. The threat sources shown in black include management consoles (MGR), applications (APP), controllers (CTRL), network elements (NE), non-SDN elements, and rogue devices.

I classify the vulnerability sources into the following (see Figure 3):

- an application
- a controller
- a network element
- a management console
- the northbound interface
- the east/west bound interface
- the southbound interface
- the management interface




Figure 3. The vulnerability sources shown in purple include applications (APP), controllers (CTRL), network elements (NE), management consoles (MGR), the northbound interface, the southbound interface, east/west bound interface, and the management interface.

The threat sources can take one or more of the following actions against a SDN component:

- unauthorized access to a SDN component where a threat source accesses a SDN component for which it does not have the proper access level
- unauthorized disclosure of information where a threat source obtains information for which it is not explicitly authorized to have access to (loss of confidentiality)
- unauthorized modification where a threat source alters data which it is not explicitly authorized to modify (loss of integrity)
- destruction where a threat source destroys an essential SDN function or data (loss of integrity)
- disruption of service where a threat source disrupts a function of a SDN component (loss of availability)

Figure 4 shows the integration of all of the threat elements and SDN components into the SDN threat model.

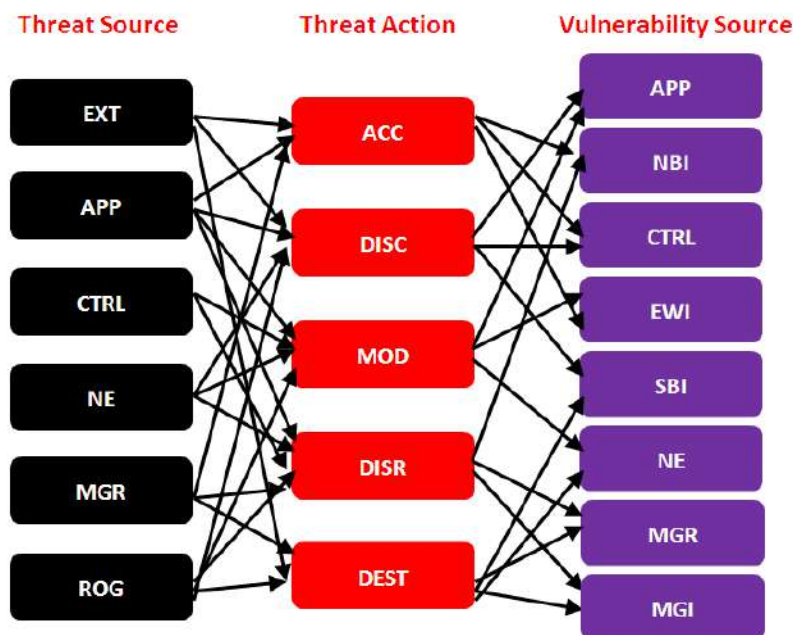


Figure 4. The SDN threat model. The threat sources shown in black include management consoles (MGR), applications (APP), controllers (CTRL), network elements (NE), non-SDN elements (EXT), and rogue devices (ROG). The threat events shown in red include unauthorized access (ACC), unauthorized disclosure of information (DISC), unauthorized modification (MOD), disruption of service (DISR), and destruction (DEST). The vulnerability sources shown in purple include applications (APP), controllers (CTRL), network elements (NE), management consoles (MGR), the northbound interface (NBI), the southbound interface (SBI), east/west bound interface (EWBI), and the management interface (MGI).

The model accounts for many-to-many relationships between the threat sources and threat events as well as the threat events and vulnerability sources. Each relationship illustrates a potential attack path starting at some threat source and taking an action to exploit vulnerabilities in vulnerability sources. While some of the paths are merely conceptual and are unlikely to materialize into actual attacks, some paths may represent real-life attack scenarios posing realistic dangers to SDN. A threat source may take multiple iterations through various attack paths to compromise a series of SDN components. For instance, after an attacker gains access to a vulnerable component, she may use the compromised component as a stepping stone to attack other components by accessing sensitive data stored in it or causing it to perform a function that could be detrimental to network availability and integrity.

4 Attacks Examples

In this section, I discuss several attack examples derived from the SDN threat model. These attacks could be launched on SDN components to achieve unauthorized access, unauthorized disclosure of information, unauthorized modification, destruction, and/or disruption of service. Most scenarios involve less privileged threat sources targeting more privileged functions or components.

4.1 Unauthorized Access Using Password Brute-Forcing or Password-Guessing Attacks

An attacker residing on a non-SDN element may achieve unauthorized access to a SDN component. For instance, an attacker may access a management console through random or systematic guessing of passwords (see Figure 5). The compromised console would empower the attacker with resources to launch attacks on the controller and the network managed by the controller.

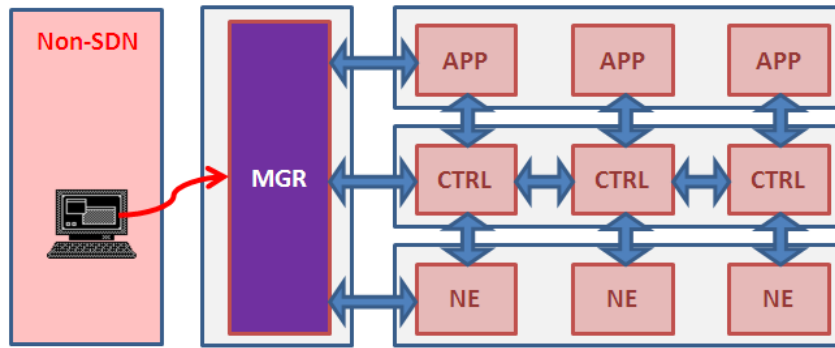


Figure 5. An unauthorized access attack is illustrated using the EXT->ACC->MGR path.

4.2 Unauthorized Access Using Remote Application Exploitation Attacks

By exploiting a software vulnerability in a SDN component, an attacker may be able to gain unauthorized access to the component. For instance, an attacker with access to a management console may exploit a buffer overflow vulnerability in a SDN application server to gain access to SDN applications (see Figure 6).

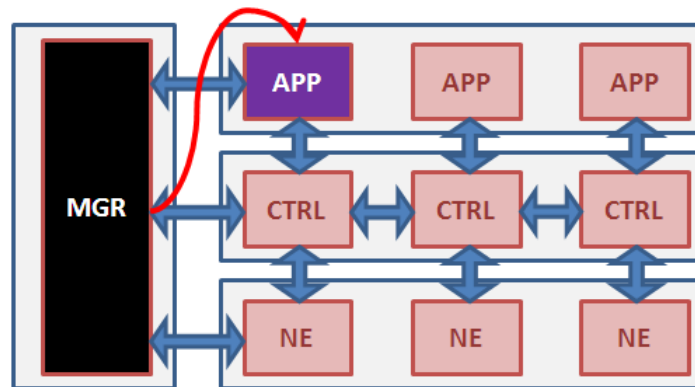


Figure 6. An unauthorized access attack is illustrated using the MGR->ACC->APP path.

4.3 Unauthorized Disclosure of Information Using RAM Scraping Attacks

In a traditional RAM scraper attack, the attacker first gains unauthorized access to a victim system, and scans its physical memory in the hope to extract sensitive information. If an attacker is successful in gaining access to a SDN application server, the attacker may scan SDN-related process memory to extract flow rules that may have previously been sent in an API request to a controller (see Figure 7).




Figure 7. An unauthorized disclosure of information attack is illustrated using the APP->DISC->APP path.

4.4 Unauthorized Disclosure of Information Using API Exploitation Attacks

The controller plane provides APIs for applications to query SDN network information and to monitor and react to network changes. The APIs or applications built using the APIs may contain vulnerabilities that could allow an attacker to perform a variety of information disclosure attacks. For example, an unauthorized disclosure of information case could involve a malicious user harvesting information regarding flow rules by exploiting the northbound interface (see Figure 8).

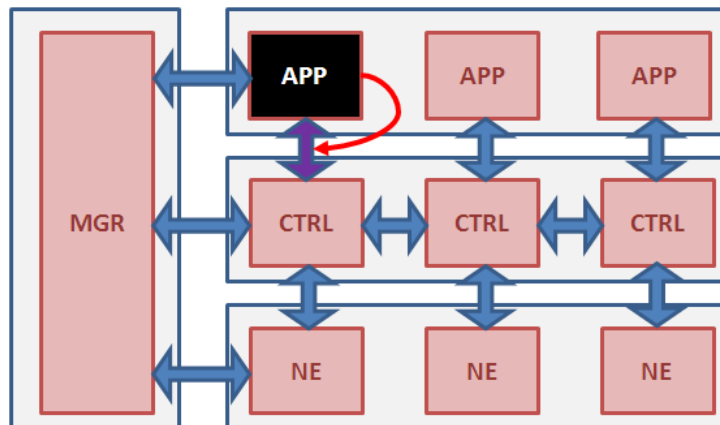


Figure 8. An unauthorized disclosure of information attack is illustrated using the APP->DISC->NBI path.

4.5 Unauthorized Destruction Using API Exploitation Attacks

An unauthorized destruction case could involve a malicious user deleting network flows to prevent traffic from reaching its destination by exploiting a vulnerability in the northbound interface (see Figure 9).




Figure 9. An unauthorized destruction attack is illustrated using the APP->DEST->NBI path.

4.6 Unauthorized Access Using Remote or Local Application Exploitation Attacks

By exploiting a software vulnerability in a SDN component, a malicious user or an attacker with access to the component may be able to escalate her privileges. For instance, an attacker could exploit a session management vulnerability to execute commands on behalf of a more privileged user (see Figure 10).




Figure 10. An unauthorized access attack is illustrated using the APP->ACC->APP path.

4.7 Unauthorized Disclosure of Information Using Traffic Sniffing Attacks

An attacker could conduct a sniffing attack to take advantage of unencrypted communications or communications using weak encryption to intercept configuration data. For instance, a legitimate controller may periodically communicate with peer controllers to synchronize the network state information. Using a compromised network tapping application, the attacker could eavesdrop on clear text communications between two controllers (see Figure 11). The intercepted communications may potentially include information about flows in use and traffic permitted across the network.




Figure 11. An unauthorized disclosure of information attack is illustrated using the APP->DISC->EWI path.

4.8 Unauthorized Modification Using Identity Spoofing Attacks

An attacker may spoof the identity of a legitimate controller to attempt to interact with a network element to instantiate flows into the network element's flow table (see Figure 12).

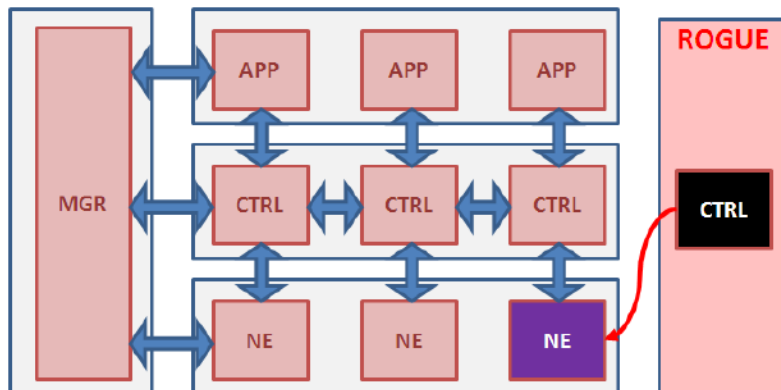


Figure 12. An unauthorized modification attack is illustrated using the ROG->MOD->NE path.

4.9 Disruption of Service Using Remote Application Exploitation Attacks

Disruption of service threats could be used to reduce availability of a SDN component. For instance, an adversary may exploit an improper input validation vulnerability in the controller or network element software to cause them to become unavailable (see Figure 13).




Figure 13. A disruption of service attack is illustrated using the MGR->DISR->CTRL and MGR->DISR->NE paths.

4.10 Unauthorized Disclosure of Information Using Side Channel Attacks

An attacker may conduct a side channel attack to determine if a flow rule already exists by detecting a difference in the time required for a new network connection to be established (see Figure 14).

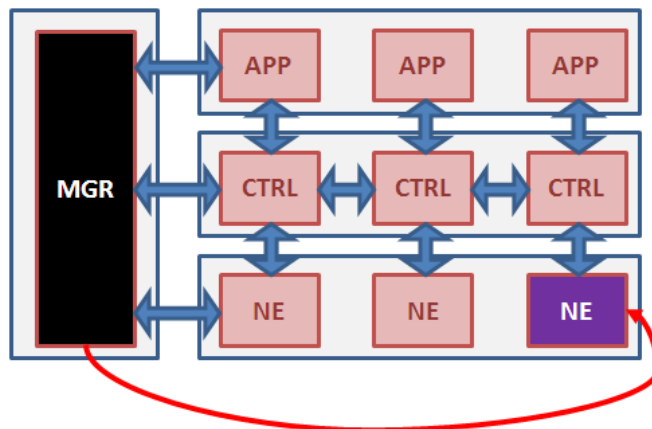


Figure 14. An unauthorized disclosure of information attack is illustrated using the MGR->DISC->NE path.

4.11 Disruption of Service Using Flooding Attacks

An attacker could leverage a compromised network element to flood a controller with network messages to exhaust its resources (see Figure 15).




Figure 15. A disruption of service attack is illustrated using the NE->DISR->CTRL path.

4.12 Unauthorized Modification Using Data Forging Attacks

A compromised network element could forge network data to poison a controller's view of the network topology (see Figure 16). This attack could be leveraged to further carry out a variety of other attacks on the network, for instance, for diverting traffic flows in the attacker's direction for eavesdropping.




Figure 16. An unauthorized modification attack is illustrated using the NE->MOD->CTRL path.

5 SDN Threat Mitigation

Using the threat model, a risk assessment could be further conducted to determine which of the threats pose realistic dangers to a particular SDN deployment and thus, must be mitigated. For each threat to be mitigated, security requirements should be specified outlining a set of goals to achieve the completeness of the SDN security. Some examples of security requirements and mitigation techniques for the threats and attacks presented in the previous section may include the following:

- 1) Threat: Conduct brute force login attempts/password guessing attacks against the management console

Security requirement: A management console shall not allow any user to successfully use a password guessing attack to gain unauthorized access

Mitigation techniques: All vendor default passwords for management consoles should be changed

- 2) Threat: Exploit a known information disclosure vulnerability in the NBI

Security requirement: An application shall not allow any user to successfully exploit a vulnerability to access information which the user is not authorized to access

Mitigation techniques: All application server patches should be applied in a timely manner

- 3) Threat: Conduct communications interception attack against the EWI

Security requirement: The east/west bound interface shall not allow unauthorized users to eavesdrop on network communications between the controllers

Mitigation techniques: The east/west bound communication channels should be protected using strong cryptography

- 4) Threat: Cause a denial of service on a controller

Security requirement: A controller shall not allow any network element to successfully use a denial of service attack to reduce its availability

Mitigation techniques: Rate limiting and packet dropping at the controller plane to avoid denial of service attacks. Specific rules should be installed on the network elements where the attack is being originated

6 Conclusion

The SDN threat model presented in this paper could be used to identify all the potential threats to SDNs before the respective attacks are conceptualized and mounted. The model characterizes the threats at the architectural level only and makes no assumptions about the actual SDN implementation. The approach could be further extended by incorporating the details of specific SDN implementations to identify further threats. The model could also serve as a foundation for planning and carrying out SDN penetration tests.

The attack paths revealed by the model provide a way for enterprises to determine if adequate protection mechanisms have been built into specific SDN designs. If a given threat is not addressed by the SDN design, a malicious attacker could exploit a weakness to compromise SDN and possibly, the entire network. The threat model and the attack examples discussed in the paper could help enterprises in developing realistic and meaningful SDN security requirements as well as mitigation measures to reduce the ability of malicious attackers to exploit SDN networks.