Network Working Group S. Hartman
Internet-Draft Painless Security
Intended status: Experimental D. Zhang
Expires: April 25, 2014 Huawei
M. Wasserman
Painless Security
October 22, 2013

Security Requirements of NVO3
draft-ietf-nvo3-security-requirements-01

Abstract

The draft provides a list of security requirements to benefit the design of NOV3 security mechanisms. In addition, this draft introduces the candidate techniques which could be used to fulfill such security requirements.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

Security is a key issue which needs to be considered in the design of a data center network. This document discusses the security risks that a NVO3 network may encounter and the security requirements that a NVO3 network needs to fulfill. In addition, this draft attempts to discuss the security techniques which could be applied to fulfill such requirements.

The remainder of this document is organized as follows. Section 2 introduces the terms used in this memo. Section 3 gives a briefly introduction of the NVO3 network architecture. Section 4 discusses the attack model of this work. Section 5 describes the essential security requirements which should be fulfilled in the generation of a NVO3 network.

2. Terminology

This document uses the same terminology as found in the NVO3 Framework document [I-D.ietf-nvo3-framework] and [I-D.kreeger-nvo3-hypervisor-nve-cp]. Some of the terms defined in the framework document have been repeated in this section for the convenience of the reader, along with additional terminology that is used by this document.

Tenant System (TS): A physical or virtual system that can play the role of a host, or a forwarding element such as a router, switch, firewall, etc. It belongs to a single tenant and connects to one or more VNs of that tenant.

End System (ES): An end system of a tenant, which can be, e.g., a virtual machine(VM), a non-virtualized server, or a physical appliance. A TS is attached to a Network Virtualization Edge(NVE) node.

Network Virtualization Edge (NVE): An NVE implements network virtualization functions that allow for L2/L3 tenant separation and tenant-related control plane activity. An NVE contains one or more tenant service instances whereby a TS interfaces with its associated instance. The NVE also provides tunneling overlay functions.

Virtual Network (VN): This is a virtual L2 or L3 domain that belongs to a tenant.

Network Virtualization Authority (NVA). A back-end system that is responsible for distributing and maintaining the mapping information for the entire overlay system. Note that the WG never reached consensus on what to call this architectural entity within the overlay system, so this term is subject to change.

NVO3 device: In this memo, the devices (e.g., NVE and NVA) work cooperatively to provide NVO3 overlay functionalities are called as NOV3 devices.

3. NVO3 Overlay Architecture

```
                      ..................................
                      .                                .
                      .                                .
                      .                                .
                    +-+--+                          +--+-++--------+
    +--------+      | NV |                          | NV || Tenant |
    | Tenant +------+Edge|      L3 Overlay          |Edge|| System |
    | System |      +-+--+       Network            +--+-++--------+
    +--------+        .                                .
                      .                                .
                      .                                .
                      ..................................
```

This figure illustrates a simple nov3 overlay example where NVEs provide a logical L2/L3 interconnect for the TSes that belong to a specific tenant network over L3 networks. A packet from a tenant system is encapsulated when they reach the egress NVE. Then encapsulated packet is then sent to the remote NVE through a proper tunnel. When reaching the ingress NVE, the packet is decapsulated and forwarded to the target tenant system. The address advertisements and tunnel mappings are distributed among the NVEs through either distributed control protocols or by certain centralized servers (called NVAs).

4. Threat Model

To benefit the discussion, in this analysis work, attacks are classified into two categories: inside attacks and outside attacks. An attack is considered as an inside attack if the adversary performing the attack (inside attacker or insider) has got certain privileges in changing the configuration or software of a NVO3 device and initiates the attack within the overlay security perimeter. In contrast, an attack is referred to as an outside attack if the adversary performing the attack (outside attacker or outsider) has no such privilege and can only initiate the attacks from compromised TSes (or the network devices of the underlying network which the overlay is located upon). Note that in a complex attack inside and outside attacking operations may be performed in a well organized way to expand the damages caused by the attack.

4.1. Outsider Capabilities

The following capabilities of outside attackers MUST be considered in the design of a NOV3 security mechanism:

1. Eavesdropping on the packets,

2. Replaying the intercepted packets, and

3. Generating illegal packets and injecting them into the network.

With a successful outside attack, an attacker may be able to:

1. Analyze the traffic pattern within the network,

2. Disrupt the network connectivity or degrade the network service quality, or

3. Access the contents of the data/control packets if they are not properly encrypted.

4.2. Insider Capabilities

It is assumed that an inside attacker can perform any types of outside attacks from the inside or outside of the overlay perimeter. In addition, in an inside attack, an attacker may use already obtained privilege to, for instance,

1. Interfere with the normal operations of the overlay as a legal entity, by sending packets containing invalid information or with improper frequencies,

2. Perform spoofing attacks and impersonate another legal device to communicate with victims using the cryptographic information it obtained, and

3. Access the contents of the data/control packets if they are encrypted with the keys held by the attacker.

4.3. Security Issues In Scope and Out of Scope

During the specification of security requirements, the following security issues needs to be considered:

1. Insecure underlying network. It is normally assumed that a underlying network connecting NOV3 devices (NVEs and NVAs) is secure if it is located within a data center and cannot be directly accessed by tenants. However, in a virtual data center scenario, a NVO3 overlay scatters across different sites which are connected through the public network. Outside attacks may be raised from the underlying network.

2. Insider attacker. During the design of a security solution for a NVO3 network, the inside attacks raised from compromised NVO3 devices (NVEs and NVAs) needs to be considered.

3. Insecure tenant network. It is reasonable to consider the conditions where the network connecting TSes and NVEs is accessible to outside attackers.

The following issues are out of scope of cosideration in this document:

1. In this memo it is assumed that security protocols, algorithms, and implementations provide the security properties for which they are designed; attacks depending on a failure of this assumption are out of scope. As an example, an attack caused by a weakness in a cryptographic algorithm is out of scope, while an attack caused by failure to use confidentiality when confidentiality is a security requirement is in scope.

2. In practice an attacker controlling an underlying network device may break the communication of the overlays by discarding or delaying the delivery of the packets passing through it. However, this type of attack is out of scope.

5. Security Requirements and Candidate Approaches

This section introduces the security requirements and candidate solutions.

5.1. Control/Data Traffic within Overlay

This section analyzes the security issues in the control and data plans of a NVO3 overlay.

5.1.1. Control Plane Security

REQ1: A NVO3 security solution MUST enable two NOV3 devices (NVE or NVA) to perform mutual authentication before exchanging control packets.

This requirement is used to prevent an attacker from impersonating a legal NVO3 device and sending out bogus control packets without being detected.

The authentication between devices can be performed as a part of automated key management protocols (e.g., IKEv2[RFC5996], EAP[RFC4137], etc.). After such an authentication procedure, an device can find out whether its peer holds valid security credentials and is the one who it has claimed. Additionally, the keys shared between the devices can be also used for the authentication purpose. For instance, assumed a NVE and a NVA have shared a secret key without known by any other third parties.

The NVE can ensure that a device that it is communicating with is the NVA if the device can prove that it possesses the shared key.

a: The identity of the network devices SHOULD be verified during authentication.

In some authentication mechanisms, instead of verifying the peers' identities, the authentication result can only prove that a device joining the authentication is a legal member of a group. However, for a better damage confining capability to insider attacker, it is recommended to verify the devices' identities during authentication. Therefore, an insider attacker cannot impersonate others, even when it holds legal credentials or keys.

REQ2: Before accepting a control packet, the device receiving the packet MUST verify whether the packet comes from one which has the privilege to send that packet.

This is an authorization requirement. A device needs to clarify the roles (e.g., a NVE or a NVA) that its authentication peer acts as in the overlay. Therefore, if a compromised NVE uses it credentials to impersonate a NVA to communicate with other NVEs, it will be detected. In addition, authorization is important for enforcing the VN isolation, a device only can distribute control packets within the VNs it is involved within. If a control packet about a VN is sent from a NVE which is not authorized to support the VN, the packet will not be accepted

Normally, it is assumed that the access control operations are based on the authentication results. The simple authorization mechanisms (such as ACLs which filters packets based on the packet addresses) can be used as auxiliary approaches since they are relatively easy to bypass if attackers can access to the network and modify packets.

REQ3: Integrity, confidentiality, and origin Authentication protection for Control traffics

It is the responsibility of a NVO3 overlay to protect the control packets transported over the overlay against the attacks raised from the underlying network.

a: The integrity and origin authentication of the packets MUST be guaranteed.

With this requirement, the receiver can ensure that the packets are from the legitimate sender, not replayed, and not modified during the transportation.

b: The signaling packets SHOULD be encrypted.

On many occasions, the signaling packets can be transported in plaintext. However, In the cases where the information contained within the signaling packets are sensitive or valuable to attackers , the signaling packets related with that tenant need be encrypted.

To achieve such objectives, when the network devices exchange control plane packets, integrated security mechanisms or underlying security protocols need to provided. In addition, cryptographic keys need to be deployed manually in advance or dynamically generated by using certain automatic key management protocols (e.g., TLS [RFC5246]). The keys are used to generate digests for or encrypt control packets.

REQ4: The toleration of DOS attacks

a: Frequency in distributing control packets within in the overlay MUST be limited.

The issues within DOS attacks also need to be considered in designing the overlay control plane. For instance, in the VXLAN solution[I-D.mahalingam-dutt-dcops-vxlan], an attacker attached to a NVE can try to manipulate the NVE to keep multicasting control packets by sending a large amount of ARP packets to query the inexistent VMs. In order to mitigate this type of attack, the NVEs SHOULD be only allowed to send signaling packet in the overlay with a limited frequency. When there are centralized servers (e.g., the backend oracles providing mapping information for NVEs[I-D.ietf-nvo3-overlay-problem-statement], or the SDN controllers) are located within the overlay, the potential security risks caused by DDOS attack on such servers can be more serious.

b: Mitigation of amplification attacks SHOULD be provided.

During the design of the control plane, it is important to consider the amplification effects. For instance, if NVEs may generate a large response to a short request, an attacker may send spoofed requests to the NVEs with the source address of a victim. Then the NVEs will send the response to the victim and result in DDOS attacks.

If the amplification effect cannot be avoided in the control protocol, the requirements 1,2,3, and 4a can all be used to benefit the mitigation of this type of attacks.

REQ5: The key management solution MUST be able to confine the scope of key distribution and provide different keys to isolate the control traffic according to different security requirements.

a: It SHOULD be guaranteed that different keys are used to secure the control packets exchanged within different tenant networks.

This requirement can be used to provide a basic attack confinement capability. The compromise of a NVE working within a tenant will not result in the key leakage of other tenant networks.

b: It SHOULD be guaranteed that different keys are used to secure the control packets exchanges with different VNs.

This requirement can be used to provide a better attack confinement capability for the control plane. The compromise of a NVE working within a VN will not result in the key leakage of other VNs. However, since there is only a single key used for securing the data traffic within a VN, an attacker which has compromised a NVE within the VN may be able to impersonate any other NVEs within the VN to send out bogus control packets. In addition, the key management overheads introduced by key revocation also need to be considered[RFC4046]. When a NVE stops severing a VN, the key used for the VN needs to be revoked, and a new key needs to be distributed for the NVO3 devices still within the VN.

If we expect to provide a even stronger confinement capability and prevent a compromised NVE from impersonating other NVEs even when they are in the same VN, different NVEs working inside a VN need to secure their signaling packets with different keys.

If there is automated key management deployed, the authentication and authorization can be used to largely mitigate the isolation issues. When a NVE attempts to join a VN, the NVE needs to be authenticated and prove that it have sufficient privileges. Then, a new key (or a set of keys) will be generated to secure its control packet exchanged with this VN.

5.1.2. Data Plane

[I-D.ietf-nvo3-framework] specifies a NVO3 overlay needs to generate tunnels between NVEs for data transportation. When a data packet reaches the boundary of a overlay, it will be encapsulated and forwarded to the destination NVE through a proper tunnel.

REQ6: Integrity, confidentiality, and origin authentication protection for data traffics a: The integrity and origin authentication of data traffics MUST be guaranteed when the underlying network is not secure.

During the transportation of data packets, it is the responsibility of the NVO3 overlay to deal with the attacks from the underlying network. For instance, an inside attacker compromising a underlying network device may intercept an encapsulated data packet transported within a tunnel, modify the contents in the encapsulating tunnel packet and, transfer it into another tunnel without being detected. When the modified packet reaches a NVE, the NVE may decapsulated the data packet and forward it into a VN according to the information within the encapsulating header generated by the attacker. Similarly, a compromised NVE may try to redirect the data packets within a VN into another VN by adding improper encapsulating tunnel headers to the data packets.

Under such circumstances, in order to enforce the VN isolation property, underlying security protocols need to provided. Signatures or digests need to be generated for both data packets and the encapsulating tunnel headers in order to provide data origin authentication and integrity protection.

b: The confidentiality protection of data traffics SHOULD be provided, when the underlying network is not secure.

If the data traffics from the TSes is sensitive, they needs to be encrypted during the tunnels. However, if the data traffics is not valuable and sensitive, the encryption is not necessary.

REQ7: Different tunnels SHOULD be secured with different keys

This requirement can be used to provide a basic attack confinement capability. When different tunnels secured with different keys, the compromise of a key in a tunnel will not affect the security of others.

5.2. Control/Data Traffic between NVEs and Hypervisors

Assume there is a VNE providing a logical L2/L3 interconnect for a set of TSes. Apart from data traffics, the NVE and certain TSes (i.e., Hypervisors) also need to exchange signaling packets in order to facilitate, e.g., VM online detection, VM migration detection, or auto-provisioning/service discovery [I-D.ietf-nvo3-framework].

The NVE and its associated TSes can be deployed in a distributed way (e.g., a NVE is implemented in an individual device, and VMs are located on servers) or in a co-located way (e.g., a NVE and the TSes it serves are located on the same server).

5.2.1. Distributed Deployment of NVE and Hypervisor

In this case, the data and control traffic between the NVE and the TSes are exchanged over network.

5.2.1.1. Control Plane

REQ8: Mutual authentication MUST be performed between a NVE and a TS at the beginning of their communication, if the network connecting them is not secure.

Mutual authentication is used to guarantee that an attacker cannot impersonate a legal NVE or a hypervisor without being detected.

There are various ways to perform mutual authentication. If there are auto key management mechanism (e.g., IKEv2, EAP), the NVE and the TS can use their credential to perform authentication. If there a key pre-distributed between a NVE and a TS, an entity can also use the key verify the identity of is remote peer.

If practice, a NVE and a TS may simply use IP or MAC addresses to identify each other. This type of technique can be used as a complementary approach although it may becomes vulnerable if attackers can inject bogus control packets the network and modify the packets transported between the NVE and TS.

REQ9: Before accepting a control packet, the receiver device MUST verify whether the packet comes from one which has the privilege to send that packet.

This is an authorization requirement. A device needs to clarify the roles (e.g., a TS or a NVE) of the device that it is communicating with. Therefore, if a compromised TS attempts to use it credentials to impersonate a NVE to communicate with other TSes, it will be detected.

Authorization is very important to guarantee the isolation property. For instance, if a compromised hypervisor tries to elevate its privilege and interfere the VNs that it is not supposed to be involved within, its attempt will be detected and rejected.

Normally, it is assumed that the access control operations are based on the authentication results. The simple authorization mechanisms (such as ACLs which filters packets based on the packet addresses) can be used as complementary solutions.

REQ10: Integrity, Confidentiality, and Origin Authentication for Control Packets

a:The security solution of a NVE network MUST be able to provide integrity protection and origin authentication for the control packets exchanged between a NVE and a TS if they have to use an insecure network to transport their packet.

This requirement can prevent an attacker from illegally interfere with the normal operations of NVEs and TSes by injecting bogus control packets into the network.

b:The confidentiality protection for the control packet exchange SHOULD be provided.

When the contents of the control packets (e.g., the location of a ES, when a VM migration happens) are sensitive to a tenant, the control packet needs to be encrypted.

There are various security protocols (such as IPsec, SSL, and TCP-AO) can be used for transport control packets. In addition, it is possible to define integrated security solutions for the control packets.

In order to secure the control traffic, cryptographic keys need to be distributed to generate digests or signatures for the control packets. Such cryptographic keys can be manually deployed in advance or dynamically generated with certain automatic key management protocols (e.g., TLS [RFC5246]).

REQ11: The key management solution MUST be able to confine the scope of key distribution and provide different keys to isolate the control traffic according to different security requirements.

a: If assuming TSes (hypervisors) will not be compromised, the TSes belonging to different Tenants MUST use different keys to secure the control packet exchanges with their NVE.

This requirement is used to enforce the security boundaries of different tenant networks. Since different tenants belong to different security domains and may be competitive to each other, the control plane traffics need to be carefully isolated so that an attacker from a tenant cannot affect the operations of another tenant network.

b: If assuming the hypervisors can be compromised, the TSes belonging to different VNs MUST use different keys to secure the control packets exchanges with their NVE.

Therefore, if a key used for a VN is compromised, other VNs will not be affected. This requirement is used to ensure the VN isolation property.

5.2.1.2. Data Plane

REQ12: The data traffic isolation of different VNs MUST be guaranteed.

In [I-D.ietf-nvo3-overlay-problem-statement], the data plane isolation requirement amongst different VNs has been discussed. The traffic within a virtual network can only be transited into another one in a controlled fashion (e.g., via a configured router and/or a security gateway). Therefore, if the NVE supports multiple VNs concurrently, the data traffic in different VNs MUST be isolated.

a:The security solution of a NVE network MUST be able to provide integrity protection and origin authentication for the data packets exchanged between a NVE and a TS if they have to use an insecure network to transport their data packet.

In practice, the data traffics in different VNs can be isolated physically or by using VPN technologies. If the network connecting the NVE and the TSes is potentially accessible to attackers, security solutions need to be considered to prevent an attacker locating in the middle between the NVE and TS from modifying the VN identification information in the packet headers so as to manipulate the NVE to transport the data packets within a VN to another. The security protocols such as IPsec and TCP-AO, can be used to enforce the isolation property if necessary.

The key management requirement R11 can be applied here for data traffic

5.3. Key Management

REQ13: A security solution for NVO3 SHOULD provide automated key management mechanisms.

In the cases where there are a large amount of NVEs working within a NVO3 overlay, manual key management may become infeasible. First, it could be burdensome to deploy pre-shared keys for thousands of NVEs, not to mention that multiple keys may need to be deployed on a single device for different purposes. Key derivation can be used to mitigate this problem. Using key derivation functions, multiple keys for different usages can be derived from a pre-shared master key. However, key derivation cannot protect against the situation where a system was incorrectly trusted to have the key used to perform the derivation. If the master key were somehow compromised, all the resulting keys would need to be changed [RFC4301]. In addition, VM migration will introduce challenges to manual key management. The migration of a VM in a VN may cause the change of the NVEs which are involved within the NV. When a NVE is newly involved within a VN, it needs to get the key to join the operations within the VN. If a NVE stops supporting a VN, it should not keep the keys associated with that VN. All those key updates need to be performed at run time, and difficult to be handled by human beings. As a result, it is reasonable to introduce automated key management solutions such as EAP [RFC4137] for NVO3 overlays.

Without the support automated key management mechanisms, some security functions of certain security protocols cannot work properly. For instance, the anti-replay mechanism of IPsec is turned off without the support of automated key management mechanisms. Therefore, if IPsec is selected to protect the control packets. In this case, the system may suffer from the replay attacks.

6. IANA Considerations

This document makes no request of IANA.

Note to RFC Editor: this section may be removed on publication as an RFC.

7. Security Considerations

TBD

8. Acknowledgements

Thanks a lot for the comments from Melinda Shore, and Zu Qiang.

9. References

9.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

9.2. Informative References

[I-D.ietf-ipsecme-ad-vpn-problem]
Manral, V. and S. Hanna, "Auto Discovery VPN Problem Statement and Requirements", draft-ietf-ipsecme-ad-vpn-problem-09 (work in progress), July 2013.

[I-D.ietf-nvo3-framework]
Lasserre, M., Balus, F., Morin, T., Bitar, N., and Y. Rekhter, "Framework for DC Network Virtualization", draft-ietf-nvo3-framework-03 (work in progress), July 2013.

[I-D.ietf-nvo3-overlay-problem-statement]
Narten, T., Gray, E., Black, D., Fang, L., Kreeger, L., and M. Napierala, "Problem Statement: Overlays for Network Virtualization", draft-ietf-nvo3-overlay-problem-statement-04 (work in progress), July 2013.

[I-D.kreeger-nvo3-hypervisor-nve-cp]
Kreeger, L., Narten, T., and D. Black, "Network Virtualization Hypervisor-to-NVE Overlay Control Protocol Requirements", draft-kreeger-nvo3-hypervisor-nve-cp-01 (work in progress), February 2013.

[I-D.mahalingam-dutt-dcops-vxlan]
Mahalingam, M., Dutt, D., Duda, K., Agarwal, P., Kreeger, L., Sridhar, T., Bursell, M., and C. Wright, "VXLAN: A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks", draft-mahalingam-dutt-dcops-vxlan-05 (work in progress), October 2013.

[RFC4046] Baugher, M., Canetti, R., Dondeti, L., and F. Lindholm, "Multicast Security (MSEC) Group Key Management Architecture", RFC 4046, April 2005.

[RFC4137] Vollbrecht, J., Eronen, P., Petroni, N., and Y. Ohba, "State Machines for Extensible Authentication Protocol (EAP) Peer and Authenticator", RFC 4137, August 2005.

[RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.

[RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, August 2008.

[RFC5996] Kaufman, C., Hoffman, P., Nir, Y., and P. Eronen, "Internet Key Exchange Protocol Version 2 (IKEv2)", RFC 5996, September 2010.

Authors' Addresses

Sam Hartman
Painless Security
356 Abbott Street
North Andover, MA 01845
USA

Email: hartmans@painless-security.com
URI: http://www.painless-security.com

Dacheng Zhang
Huawei
Beijing
China

Email: zhangdacheng@huawei.com

Margaret Wasserman
Painless Security
356 Abbott Street
North Andover, MA 01845
USA

Phone: +1 781 405 7464
Email: mrw@painless-security.com
URI: http://www.painless-security.com