

Network Working Group
Internet Draft
Intended status: Informational
Expires: December 18, 2015

S. Hartman
Painless Security
D. Zhang

M. Wasserman
Painless Security
Z. Qiang
Ericsson
M. Zhang
Huawei
June 18, 2015

Security Requirements of NV03
draft-ietf-nvo3-security-requirements-05

Abstract

The draft describes a list of essential requirements in order to benefit the design of NV03 security solutions. In addition, this draft introduces the candidate techniques which could be used to construct a security solution fulfilling these 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

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on December 18, 2015.

Copyright Notice

Copyright (c) 2015 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](http://trustee.ietf.org/license-info) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction.....	3
2. Terminology.....	3
3. NV03 Overlay Architecture.....	4
4. Threat Model.....	5
4.1. Capabilities of Outsiders.....	5
4.2. Capabilities of Insiders.....	5
4.3. Capabilities of Malicious TSes.....	6
5. Scope.....	6
6. Security Requirements.....	7
6.1. Control Plane of NV03 Overlay.....	7
6.2. NVE-NVE Data Plane.....	11
6.3. NVE-Hypervisor Data Plane.....	13
7. Candidate Techniques.....	14
7.1. Entity Authentication.....	15
7.2. Packet Level Security.....	15
7.3. Authorization.....	15
7.4. Automated Key Management.....	16
8. IANA Considerations.....	16
9. Security Considerations.....	16
10. Acknowledgements.....	17
11. References.....	17
11.1. Normative References.....	17
11.2. Informative References.....	17
Authors' Addresses.....	19

1. Introduction

As described in [[RFC7365](#)], the NV03 framework is intended to aid in standardizing protocols and mechanisms to support large-scale multi-tenancy data centers. In such kind data center, security is a key issue which needs to be considered during the network design. This document discusses the security risks that a NV03 network may encounter and tries to provide a list of essential security requirements that needs to be fulfilled. In addition, this document introduces the candidate techniques which could be potentially used to construct a security solution fulfilling the NV03 security requirements.

The remainder of this document is organized as follows. [Section 2](#) introduces several key terms used in this memo. [Section 3](#) gives a brief introduction of the NV03 network architecture. [Section 4](#) discusses the attack model of this document. [Section 5](#) lists the scope of the security considerations of this document. [Section 6](#) provides a list of security requirements as well as the associated justifications. In [Section 7](#), the candidate techniques are introduced.

2. Terminology

This document uses the same terminology as defined in the NV03 Framework document [[RFC7365](#)] and the Hypervisor to NVE Control Plane Requirements document [[I-D.ietf-nvo3-hpvr2nve-cp-req](#)]. The Followings are the additional terminologies that are used by this document.

Hypervisor: This memo uses the term "hypervisor" throughout when describing requirements at the Split-NVE scenario where part of the NVE functionality is off-loaded to a separate device from the "hypervisor" that contains a VM connected to a VN. In this context, the term "hypervisor" is meant to cover any device type where part of the NVE functionality is off-loaded in this fashion, e.g. a Network Service Appliance, Linux Container.

NV03 device: In this memo, the devices (e.g., NVE and NVA) work cooperatively to provide NV03 overlay functionalities are referred as NV03 devices.

3. NV03 Overlay Architecture

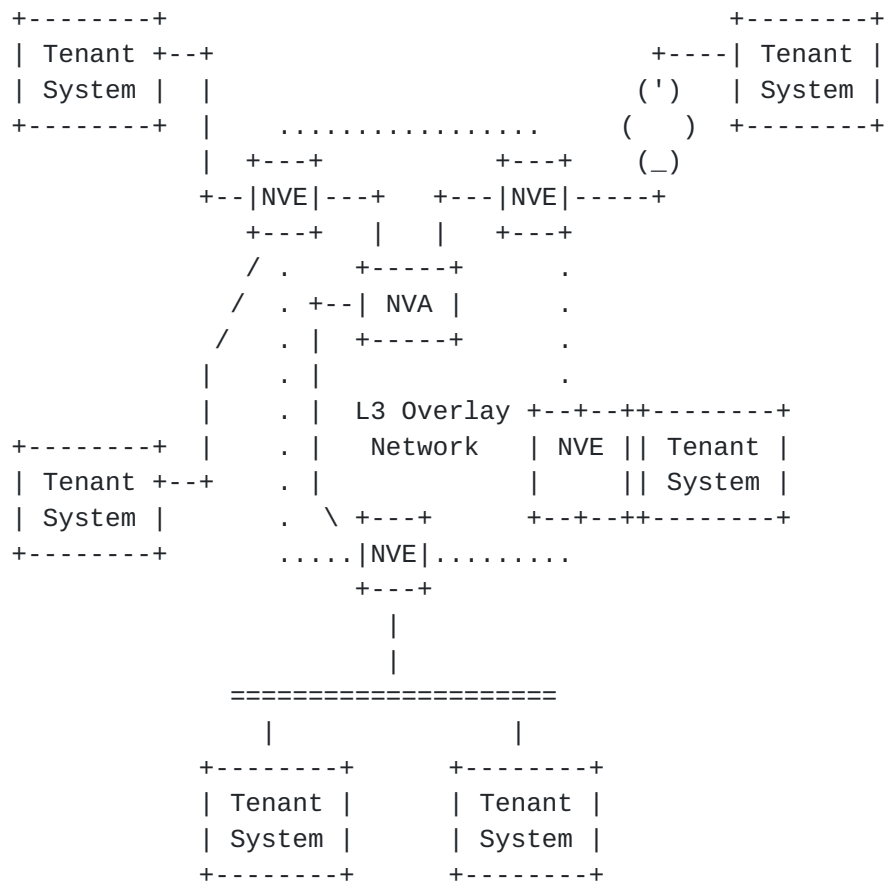


Figure 1 : Generic Reference Model for DC Network Virtualization Overlays [[RFC7365](#)]

This figure illustrates a generic reference model for NV03 overlay where NVEs provide a logical L2/L3 interconnect for the TSEs that belong to a specific tenant network over a L3 networks. A packet received from a tenant system is encapsulated by the ingress NVE. Then encapsulated packet is then sent to the remote NVE through a proper tunnel. When reaching the egress NVE of the tunnel, the packet is decapsulated and forwarded to the target tenant system. The address mappings and other related information are distributed to the NVEs by a logically centralized Network Virtualization Authority (NVA).

4. Threat Model

To benefit describing the threats a NV03 network may have to face, the attacks considered in this document are classified into three categories: the attacks from compromised NV03 devices (inside attacks), the attacks from compromised tenant systems, and the attacks from underlying networks (outside attacks).

The adversaries performing the first type of attack are called as insiders or inside attackers because they need to get certain privileges in changing the configuration or software of NV03 devices beforehand and initiate the attacks within the overlay security perimeter. In the second type of attack, an attacker (e.g., a malicious tenant, or an attacker who has compromised a virtual machine of an innocent tenant) has got certain privileges in changing the configuration or software of tenant systems and attempts to manipulate the controlled tenant systems to interfere with the normal operations of the NV03 overlay. The third type of attack is referred to as the outside attack since adversaries do not have to obtain any privilege on the NV03 devices or tenant systems in advance in order to perform this type attack, and thus the adversaries performing outside attacks are called as outside attackers or outsiders.

4.1. Capabilities of Outsiders

In practice, an outside attacker may perform attacks by intercepting packets, deleting packets, and/or inserting bogus packets. With a successful outside attack, an attacker may be able to:

- A) Analyze the traffic pattern within the network by performing passive attacks;
- B) Disrupt the network connectivity or degrade the network service quality (e.g., by performing DoS attacks); or
- C) Access the contents of the data/control packets which are not properly encrypted.

4.2. Capabilities of Insiders

Besides intercepting packets, deleting packets, and/or inserting bogus packets, an inside attacker may use already obtained privilege to,

- A) Interfere with the normal operations of the overlay as a legal NV03 device, by sending packets containing invalid information or with improper frequencies;

- B) Perform spoofing attacks and impersonate another legal NV03 device to communicate with victims using the cryptographic information it obtained; and
- C) Access the contents of the data/control packets if they are encrypted with the keys held by the attacker.

4.3. Capabilities of Malicious TSes

It is assumed that the attacker performing attacks from compromised TSes is able to intercept packets, delete packets, and/or insert bogus packets. In addition, after compromising a TS, an attacker may be able to:

- A) Interfere with the normal operations of the overlay as a legal TS, by sending packets containing invalid information or with improper frequencies to NVEs;
- B) Perform spoofing attacks and impersonate another legal TS or NVE to communicate with victims (other legal NVEs or TSes) using the cryptographic information it obtained; and
- C) Access the contents of the data/control packets if they are encrypted with the keys held by the attacker.

5. Scope

The following security issues are in the scope of the NV03 security requirement consideration of this document:

- A) The NV03 connections may be considered as secured if there is a security solution supported by the underlying network. However such kind security solution normally only can protect the NV03 network from outsider attacker.
- B) During the design of a security solution for a NV03 network, the attacks raised from compromised NVEs and hypervisors needs to be considered.
- C) It is reasonable to consider the conditions where the network connecting TSes and NVEs is accessible to outside attackers.

The following security issues are out of scope of the NV03 security requirement consideration of this document:

- A) In this document, 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. For instance, 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.

B) 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. The security consideration to prevent this type of attack is out of scope of this document.

C) NVAs are centralized servers and play a critical role in NV03 overlay network. A NVE will believe in the mapping information obtained from its NVA. After compromising a NVA, the attacker can distribute bogus mapping information to NVEs under the management of NVA. The security requirements discussed in this document is to protect a NVA from any security risk. And if a NVA is attacked, it should be detected. However, this document does not consider how to deal with the problem after a NVA is compromised.

D) Because this document only tries to provide the most essential high

level requirements, some important issues in designing concept security mechanisms are not covered in the requirements. Such issues include:

- How to manage keys/credentials during their life periods
- How to support algorithm agility
- How to provide accountability
- How to secure the management interfaces
- Use underlying security protocols versus design integrated security extensions

6. Security Requirements

6.1. Control Plane of NV03 Overlay

In this section, the security requirements associated with following control plane are described:

- The NVE-NVA control plane: allows a NVE to obtain information about the location and status of other TSs with which it needs to communicate; to provide updates to the NVA about the attached TSs;

and to report any communication errors. In this case, the term "NV03 device" is referring to a NVA or a NVE.

- The NVA-NVA control plane: Multiple NVAs may be deployed in a NV03 overlay for better scalability and fault tolerance capability. The NVAs may use unicast and/or multicast to exchange signaling packets within the control plane. In this case, the term "NV03 device" is referring to a NVA.
- The NVE-NVE control plane: As specified in [[RFC7365](#)], in order to obtain reachability information, NVEs may exchange information directly between themselves via a control-plane protocol. In this case, the term "NV03 device" is referring to a NVE.
- The NVE-Hypervisor control plane: In the Split-NVE scenario, the NVE and hypervisors may also need to exchange signaling packets over network in order to facilitate, e.g., VM online detection, VM migration detection, or auto-provisioning/service discovery as described in [[RFC7365](#)]. In this case, the term "NV03 device" is referring to a Hypervisor or a NVE.

REQ 1. The security solution for NV03 MUST enable the two NV03 devices to mutually authenticate each other before exchanging any control packets.

Entity authentication can protect a NV03 device against imposter attacks and then reduce the risk of DoS / DDoS attacks and man-in-the-middle attacks. In addition, a successful authentication normally results in the distribution key materials for the security protection for subsequent communications. More detailed discussions are provided in [Section 6.1](#).

REQ 2. The security solution of NV03 MUST be able to provide integrity protection, replay protection, and packet origin authentication for the control packets exchanged between two NV03 devices.

Message authentication is performed on each incoming packet. Packet level security protection can prevent an attacker from illegally interfere with the normal operations of NV03 device by injecting bogus control packets into the network. Through message authentication, the NV03 device receiving a control packet can verify whether the packet is generated by a legitimate NV03 device, is not antique, and is not tampered during transportation.

Such protection must be deployed if there is any possibility that the control packets could be accessed by outside attackers. This

protection can prevent an attacker locating in the middle between the NV03 devices and modifying the information in the control packet so as to redirect the traffic as wished. In addition, with the support of properly distributed keys, these level protections can also benefit the detection of spoofing attacks raised from insiders.

REQ 3. The security solution of a NV03 network SHOULD provide confidentiality protection for the control packets exchanged between two NV03 devices.

On many occasions, the control packets can be transported in plaintext. However, if the information contained within the control packets is considered to be sensitive or valuable, it is recommended to encrypt the packets in order to prevent outsiders from accessing the sensitive data, especially when the underlying network is not secured enough. Note that encryption will impose additional overhead in processing control packets and make NV03 devices more vulnerable to DoS / DDoS attacks.

REQ 4. Node authorization procedure MUST be supported before processing any received control packets in the NV03 device

When receiving a control packet, besides authentication, authorization needs to be carried out by the receiver to identify the role that the packet sender acts as in the overlay and then assess the sender's privileges. If a compromised NV03 device tries to illegally elevate its privilege, it will be detected and rejected. For instance, a compromised NV03 device may use its credentials to communicate with other NVEs as a NVA, or attempting to access or update the mapping information of the VNs which it is not authorized to serve.

REQ 5. The security solution of NV03 SHOULD be able to provide distinct cryptographic keys for each NV03 device to protect the unicast control traffics exchanged between different NV03 devices respectively.

During the exchange of control packets, keys are critical in authenticating the packet senders. The purpose of this requirement is to provide a basic capability to confine the damage caused by inside attacks. After compromising a NV03 device, an attacker may be able to use the keys it obtained to exchange control traffics with other NV03 devices. But it will not be able to use the keys it obtained to breach the security of the control traffics exchanged between other NV03 devices.

REQ 6. The security solution of NV03 SHOULD be able to assign distinct cryptographic group keys for each multicast group to protect the multicast packets exchanged among the NV03 devices within the group.

In order to provide an essential packet level security protection specified for integrity and confidentiality, at least one group key may need to be shared among the NV03 devices in a same multicast group. It is recommended to use different keys for different multicast groups.

REQ 7. The resistance at DOS/DDoS attack MUST be considered in the design of NV03 control plane

Any NV03 devices may be used by an attacker to initiate a DOS/DDoS. One example is that in a NV03 overlay, NVAs can be the valuable targets of DoS / DDoS attacks, and large amount of NVEs can be potentially used as reflectors in reflection attacks. Therefore, the DoS / DDoS risks needs be considered during designing the control planes for NV03. The following requirements, but not limited to this listed, are used to benefit the migration of DoS/DDoS issue.

REQ 7.a. A NV03 device MUST have a frequency limitation at sending its control packets and processing any received control packets.

Without this limitation, an attacker can attempt to perform DoS / DDoS attacks to exhaust the limited computing and memory resources of a target NV03 device by manipulating a compromised NV03 device to generate a significant amount of control plane packets in a short period.

REQ 7.b. The amplification effect MUST be avoided

A distributed denial-of-service attack may involve sending forged requests of some type to a very large number of NV03 devices that will reply to the requests. If in certain conditions, the responses generated by a NV03 device are a much longer process than the received requests. An attacker may take advantage of this amplification effect procedure, which the NV03 device is used as a reflector to carry out DoS / DDoS attacks towards a victim NV03 device.

For instance, the attacker may send request messages to a NV03 device with a spoofed source address set to the targeted victim. In that case, all the replies generated by the NV03 device will be sent (and flooded) to the target. Another example is that as discussed in [I-

D.ietf-nvo3-arch], a NVE may wish to query the NVA about individual mapping when receiving a packet with unknown destination address. This query procedure may also be triggered at ARP / ND message handling or when NVE-NVE interaction message is received. An attacker may take advantage of this query procedure which the NVE is used as a reflector to carry out DoS / DDoS attacks towards the NVA.

Specifically, the attacker can concurrently send out a large amount of spoofed short request messages to multiple NV03 devices which the amplification effect can be enlarged which may overwhelm the victim's processing capability quickly.

REQ 8. The security solution of a NV03 SHOULD be able to provide different security levels of protections for the control traffics and data traffics exchanged between NV03 devices.

In NVE-NVE interface and NVE-Hypervisor interface, the same security solution may be used to protect both the control plane and data plane traffic. In many cases, the control and data traffics between NV03 devices may be transported over the same path or even within the same security channel. However, the control traffics and data traffics may have different levels of security sensitivity. Therefore, the protection on the traffic needs be distinguished. In this case, the security solution may need to provide different security channels for control traffics and data traffics respectively and protect the data traffics and control traffics exchanged between NV03 devices with different keys and ciphers.

6.2. NVE-NVE Data Plane

As specified in [[RFC7365](#)], a NV03 overlay needs to generate tunnels between NVEs for data packet transportation. When a data packet reaches the boundary of an overlay, the ingress NVE will encapsulate the packet and forward it to the destination egress NVE through a proper tunnel.

REQ 9. The security solution for NV03 MAY enable two NVEs to mutually authenticate each other before establishing a tunnel for data transportation.

This entity authentication requirement is used to protect a NVE against imposter attacks. Also, this requirement can help guarantee a data tunnel is generated between two proper NVEs and reduce the risk of man-in-the-middle attacks.

In order to protect the data packets transported over the overlay against the attacks raised from the underlying network, the NV03

overlay needs to provide essential security protection for data packets.

REQ 10. The security solution of NV03 SHOULD be able to provide integrity protection, replay protection, and packet origin authentication for data traffics exchanged between NVEs.

This requirement is used to prevent an attacker who has compromised underlying network devices on the path from replaying antique packets or injecting bogus data packets without being detected.

Such protection must be deployed if there is any possibility that the data packets could be accessed by outside attackers. This protection can prevent an attacker locating in the middle between the NVEs and modifying the tunnel address information in the data packet header so as to redirect the data traffic as wished.

REQ 11. The security solution of NV03 MAY be able to provide confidentiality protection for data traffics exchanged between NVEs, if information leaking is a concern.

If TS data traffic privacy is required, the TS data traffic needs to be encrypted when being transported within the overlay. In practice, tenants may select end-to-end security solutions to encrypt their sensitive data during transportation. Therefore this confidentiality requirement for data plane is an optional requirement.

REQ 12. The security solution of NV03 SHOULD be able to assign different cryptographic keys to protect the unicast tunnels between NVEs respectively.

This requirement is used to confine the damage caused by inside attacks. When different tunnels secured with different keys, the compromise of a key in a tunnel will not affect the security of other tunnels. In addition, if the key used to protect a tunnel is only shared by the NVEs on the both sides, the egress NVE receiving a data packet is able to distinctively prove the identity of the ingress NVE encapsulating the data packet during the message authentication.

REQ 13. If there are multicast packets, the security solution of NV03 SHOULD be able to assign distinct cryptographic group keys to protect the multicast packets exchanged among the NVEs within different multicast groups.

In NV03, a NVE may need to support data plane multicast capability. In order to provide an essential packet level security protection (including authentication, integrity, confidentiality) for the

multicast packets transferred within the group, at least one group key may need to be shared among the NVEs of the same multicast group. It is recommended to deploy different keys for different multicast groups, in order to confine the insider attacks on NVEs.

REQ 14. Upon receiving a data packet, an egress NVE MUST be able to verify whether the packet is sent from a proper ingress NVE which is authorized to forward that packet.

In cooperation with authentication, authorization enables an egress NVE to detect the data packets which violate certain security policies, even when they are forwarded from a legal NVE. For instance, if the remote NVE is not authorized to forward data packet of a given VN, the packet needs to be detected and discarded without processing. Note that the detection of an invalid packet may not indicate that the system is under a malicious attack. Mis-configuration or byzantine failure of a NVE may also result in such invalid packets.

6.3. NVE-Hypervisor Data Plane

As described in the NV03 architecture draft [[I-D.ietf-nvo3-arch](#)], in split-NVE scenario, a number of link types are possible between NVE and hypervisor. One simple deployment scenario may have a simple L2 Ethernet link. A more complicated scenario may have the server and NVE separated by a bridged access network, such as when the NVE resides on a ToR, with an embedded switch residing between servers and the ToR.

In any of above deployment scenarios, the data link between NVE and hypervisor may be potentially accessible to attackers, e.g. with a shared link. In that case, security solutions, including integrity protection and confidentiality protection, may be needed to secure the data link.

REQ 15. The security solution of NV03 SHOULD be able to provide integrity protection, replay protection and origin authentication for the data packets exchanged between a NVE and a hypervisor.

Packet level security protection can prevent an attacker from illegally interfere with the normal operations of NVEs and hypervisors by injecting bogus packets into the network. Because it is assumed that the network connecting the NVE and the hypervisor is potentially accessible to attackers, security solutions need to prevent an attacker locating in the middle between the NVE and the

hypervisor from modifying the information in the data packet headers so as to redirect the traffic as wished.

REQ 16. The security solution of a NV03 network MAY provide confidentiality protection for the data traffics exchanged between a NVE and a hypervisor.

If TS data packet privacy is required, the data packet needs to be encrypted. The security solution of a NVE network may need to provide confidentiality for the data packets exchanged between a NVE and a hypervisor if they have to use an insecure network to transport their data packet.

REQ 17. The security solution of a NV03 network MAY be able to provide different cryptographic keys to secure the unicast data traffic exchanged between different hypervisors and their NVEs respectively.

This requirement is used to benefit the damage confinement of inside attacks. For instance, data traffic may be forwarded over a shared link between a NVE and a hypervisor. In that case, the compromise of a hypervisor or a NVE will not be able to affect the security of data traffics exchanged between different hypervisors and their NVEs.

REQ 18. The security solution of NV03 MAY be able to assign distinct cryptographic group keys to protect the multicast traffic exchanged between different hypervisors and their NVEs respectively within different multicast groups.

If there are multicast data traffic between hypervisors and their NVE, in order to provide an essential packet level security protection (including authentication, integrity, confidentiality) for the multicast packets transferred within the multicast group, at least one group key may need to be shared among the hypervisors and their NVE of the same multicast group. It is recommended to deploy different keys for different multicast groups, in order to confine the insider attacks on the hypervisors and their NVE.

7. Candidate Techniques

This section introduces the techniques which can potentially be used to fulfill the security requirements introduced in [Section 6](#).

7.1. Entity Authentication

Entity authentication is normally performed as a part of automated key management, and a successful authentication may result in the key materials used in subsequent communications.

In the circumstance where no authentication protocols are applied, the communicating entities could use message authentication mechanisms to verify each other's identity.

The widely adopted protocols supporting entity authentication include: IKE [[RFC2409](#)], IKEv2 [[RFC5996](#)], EAP [[RFC4137](#)], TLS [[RFC5246](#)] and etc.

It is recommended to cryptographically verify the devices' identities during authentication. Therefore, an inside attacker cannot use the keys or credentials got from the compromised device to impersonate other victims.

7.2. Packet Level Security

There are requirements about protecting the integrity, confidentiality, and provide packet origin authentication for control/ data packets. Such functions can be provided through using the underlying security protocols, e.g., IPsec AH [[RFC4302](#)], IPsec ESP [[RFC4303](#)], TLS [[RFC5246](#)], or MACsec [[802.1AE](#)]. Also, when designing the control protocols, people can select to provide embedded security approaches (just like the packet level security mechanism provided in OSPFv2 [[RFC2328](#)]). The cryptographic keys can be manually deployed or dynamically generated by using certain automatic key management protocols. Note that when using manual key management, the replay protection mechanism of IPsec will be switched off.

7.3. Authorization

Without any cryptographic supports, the authorization mechanisms (e.g., packet filters) could be much easier to be bypassed by attackers, and thus the authorization mechanisms deployed on NV03 devices should interoperate with entity authentication and other packet level security mechanisms, and be able to make the access control decisions based on the cryptographically proved results.

An exception is packet filtering. Because packet filters are efficient and can effectively drop some un-authorized packets before they have to be cryptographically verified, it is worthwhile to use packet filters as an auxiliary approach to dealing with some simple

attacks and increasing the difficulties of DoS / DDoS attacks targeting at the security protocol implementations.

For instance, a NVE may maintain an authorization NVE table. This table may be distributed by a trusted entity, e.g. NVA, in combination with the inner-outer address mapping table. And NVE may use this table to filter the received control / data packets over NVE-NVE interface. The NVE may effectively drop any packets received from an unauthorized NVE before processing it, e.g. cryptographically verification procedure.

7.4. Automated Key Management

Because entity authentication and automated key distribution are normally performed in the same process, the requirements of entity authentication have already implied that it is recommended to use automated key management in the security solutions for NV03 networks. In the cases where there are a large amount of NVEs working within a NV03 overlay, manual key management becomes infeasible. First, it could be tedious 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](#)]. Moreover, some security protocols need the support of automated key management in order to perform certain security functions properly. As mentioned above, the replay protecting mechanism of IPsec will be turned off without the support of automated key management mechanisms.

8. IANA Considerations

This document makes no request of IANA.

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

9. Security Considerations

This is a requirement document which provides security requirements for the NV03 network and in itself does not introduce any new security concerns.

10. Acknowledgements

Many people have contributed to the development of this document and many more will probably do so before we are done with it. While we cannot thank all contributors, some have played an especially prominent role. The followings have provided essential input: Melinda Shore and Makan Pourzandi.

11. References

11.1. Normative References

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

11.2. Informative References

[I-D.ietf-nvo3-arch] Black, D., Narten, T., et al, "An Architecture for Overlay Networks (NV03)", [draft-narten-nvo3-arch](#), work in progress.

[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-hpvr2nve-cp-req] Yizhou, L., Yong, L., Kreeger, L., Narten, T., and D. Black, "Hypervisor to NVE Control Plane Requirements", [draft-ietf-nvo3-hpvr2nve-cp-req-01](#) (work in progress), November 2014.

[RFC2409] Harkins, D. and D. Carrel, "The Internet Key Exchange (IKE)", [RFC 2409](#), November 1998.

[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.

[RFC4302] Kent, S., "IP Authentication Header", [RFC 4302](#), December 2005.

- [RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", [RFC 4303](#), December 2005.
- [RFC4306] Kaufman, C., "Internet Key Exchange (IKEv2) Protocol", [RFC 4306](#), 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.
- [RFC7364] Narten, T., Gray, E., Black, D., Fang, L., Kreeger, L., and M. Napierala, "Problem Statement: Overlays for Network Virtualization", [RFC 7364](#), October 2014.
- [RFC7365] Lasserre, M., Balus, F., Morin, T., Bitar, N., and Y. Rekhter, "Framework for Data Center (DC) Network Virtualization", [RFC 7365](#), October 2014.
- [802.1AE] 802.1AE - Media Access Control (MAC) Security

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
Chaoyang Dist. Beijing
P.R. China

Email: dacheng.zhang@gmail

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>

Zu Qiang
Ericsson
8400 Decarie Blvd.
Town of Mount Royal, QC, H4P 2N2
Canada

Phone: +1 514 345 7900 x47370
Email: Zu.Qiang@ericsson.com

Mingui Zhang
Huawei Technologies
No. 156 Beiqing Rd. Haidian District,
Beijing 100095
P.R. China

Email: zhangmingui@huawei.com

