

Network Working Group
Internet-Draft
Intended status: Experimental
Expires: July 28, 2014

S. Hartman
Painless Security
D. Zhang
Huawei
M. Wasserman
Painless Security
January 24, 2014

Security Requirements of NV03
draft-ietf-nvo3-security-requirements-02

Abstract

The draft describes a list of essential requirements in order to benefit the design of NOV3 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 July 28, 2014.

Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) 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	2
2.	Terminology	3
3.	NV03 Overlay Architecture	4
4.	Threat Model	4
4.1.	Capabilities of Outsiders	5
4.2.	Capabilities of Insiders	5
4.3.	Capabilities of Malicious TSeS	5
4.4.	Security Issues In Scope and Out of Scope	6
5.	Security Requirements	7
5.1.	Control/Data Plane of NV03 Overlay	7
5.1.1.	NVE-NVA Control Plane	7
5.1.2.	NVA-NVA Control Plane	9
5.1.3.	NVE-NVE Data Plane	10
5.2.	Control/Data Plane between NVEs and Hypervisors	11
5.2.1.	Distributed Deployment of NVE and Hypervisor	12
6.	Candidate Techniques	14
6.1.	Entity Authentication	14
6.2.	Packet Level Security	15
6.3.	Authorization	15
7.	IANA Considerations	15
8.	Security Considerations	15
8.1.	Automated Key Management in NV03	15
8.2.	Issues not Discussed	16
9.	Acknowledgements	16
10.	References	16
10.1.	Normative References	16
10.2.	Informative References	17
	Authors' Addresses	18

[1.](#) Introduction

Security is a key issue which needs to be considered during the design of a data center network. This document discusses the security risks that a NV03 network may encounter and tries to provide a list of essential security requirements that a NV03 network needs to fulfill. In addition, this draft introduces the candidate

techniques which could be potentially used to construct a security solution fulfilling the 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 work. [Section 5](#) provides a list of security requirements as well as the associated justifications. In [Section 6](#), the candidate techniques are introduced.

2. Terminology

This document uses the same terminology as found in the NV03 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.

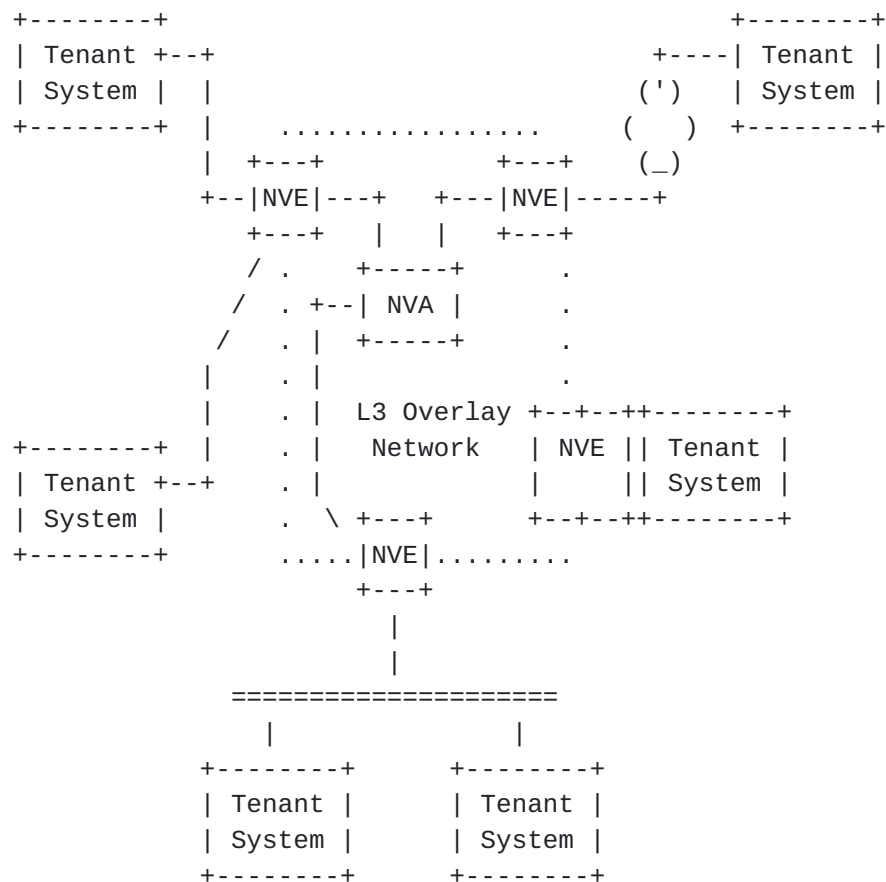
Network Virtualization Edge (NVE): An NVE implements network virtualization functions that allow for L2/L3 tenant separation and tenant-related control plane activity. An NVE contains one or more tenant service instances whereby a TS interfaces with its associated instance. The NVE also provides tunneling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

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

3. NV03 Overlay Architecture



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 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 advertisements and tunnel mappings are distributed to the NVEs by a logically centralized server (i.e., NVA).

4. Threat Model

To benefit describing the threats a NV03 network may have to face, in this work, attacks 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 has got certain privileges in changing the configuration or software of tenant systems (e.g., hypervisors or virtual machines) 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:

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 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,

1. Interfere with the normal operations of the overlay as a legal NV03 device, by sending packets containing invalid information or with improper frequencies,
2. Perform spoofing attacks and impersonate another legal NV03 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. 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:

1. Interfere with the normal operations of the overlay as a legal TS, by sending packets containing invalid information or with improper frequencies to NVEs,
2. 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
3. Access the contents of the data/control packets if they are encrypted with the keys held by the attacker.

4.4. Security Issues In Scope and Out of Scope

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

1. A underlying network connecting NOV3 devices (NVEs and NVAs) is relatively secure if it is located within a data center and cannot be directly accessed by any tenants or outsiders. However, a NV03 overlay for virtual data center may scatter across different geographically distributed sites which are connected through the public Internet. In this case, outside attacks may be raised from the underlying network connecting NV03 devices.
2. During the design of a security solution for a NV03 network, the attacks raised from compromised NVEs and hypervisors needs to be considered.
3. 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 consideration 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. 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.
2. 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. This type of attack is out of scope of this memo.

3. NVAs are centralized servers and play a critical role in NV03 overlays. 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. This work does not consider how to deal with this problem.

5. Security Requirements

5.1. Control/Data Plane of NV03 Overlay

In this section, the security requirements associated with the NVE-NVA control plane, the NVA-NVA control plane, and the NVE-NVE data plane are proposed.

5.1.1. NVE-NVA Control Plane

In a NVE-NVA control plane, it is assumed that a NVE only exchanges control traffics with its NVA using unicast.

REQ 1: The security solution for NV03 SHOULD enable two NV03 devices to mutually authenticate each other before they exchange any control packets.

Entity authentication can protect a network device against imposter attacks and then reduce the risk of DoS 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.

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

Unlike entity authentication, message authentication is performed on each incoming packet. 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. In addition, with the support of properly distributed keys, the packet level protection can also benefit the detection of spoofing attacks raised from insiders.

REQ 3: The security solution of a NV03 network MAY provide confidentiality protection for the control packets.

On many occasions, the control packets can be transported in plaintext. However, under the circumstances where some information contained within the control packets is considered to be sensitive or valuable, the information needs to be encrypted especially when the underlying network is not secure. Note that encryption will impose additional overhead in processing control packets and make NVAs more vulnerable to DoS/DDoS attacks.

REQ 4: Before accepting a control packet, a NOV3 device receiving the packet **MUST** be able to verify whether the packet comes from one who has the privilege to send that packet.

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 NVE tries to illegally elevate its privilege (e.g., using its credentials to communicate with other NVEs as a NVA, or attempting to access the mapping information of the VNs which it is not authorized to serve), it will be detected and rejected.

REQ 5: The security solution of NV03 **SHOULD** be able to provide distinct keys to protect the control traffics exchanged between a NVA and different NVEs 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 NVE, an attacker will not be able to use the keys it obtained to breach the security of the control traffics exchanged between the NVA and other NVEs.

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 NOV3. The following two requirements are used to benefit the migration of DoS/DDoS issue.

REQ 6: A NV03 device **MUST** send its control packets with limited frequencies.

Without this limitation, an attacker can attempt to perform DDoS attacks to exhaust the limited computing and memory resources of a NVA by manipulating the NVEs attached to the NVA to generate a significant member of mapping queries in a short period.

REQ 7: The amplification effect SHOULD be avoided

If in certain conditions the responses generated by a NVE are much longer than the received requests, the NVE may be taken advantage of by an attacker as a reflector to carry out DDoS attacks. Specifically, the attacker can concurrently send out a large amount of spoofed short requests to multiple NVEs with the source address of a victim (e.g., a NVA). The responses generated by the NVEs will be forwarded to the victim and overwhelm the victim's processing capability.

5.1.2. 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.

Except the key deployment requirement (REQ 5), all the other requirements in the NVE-NVA control plane (REQs 1,2,3,4, 6, and 7) are applicable in the NVA-NVA control plane as well. Before two NVA communicate with each other, they should be able to mutually authenticated. In addition, message authentication can help a NV03 device to verify the authenticity of the received packets, and the sensitive information in the control packets need to be encrypted. Authorization is important to filter the invalid control packets and any un-privileged requests. Moreover, the approach to mitigating DoS/DDoS attacks needs to be considered in the control plane protocols.

The key deployment requirements for the NVA-NVA control plane are described as follows:

REQ 8: The security solution of NV03 SHOULD be able to provide different keys to protect the unicast control traffics exchanged between different NVAs respectively.

The purpose of this requirement is to provide a basic capability to confine the damage caused by compromised key. The compromise of a key will not affect the traffics protected by other keys.

REQ 9: 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 NVAs within different multicast groups.

In order to provide an essential packet level security protection specified in REQs 2 and 3, at least a group key may need to be

shared among the NVEs in a same mutlicast group. It is recommended to user different keys for different mutlicast groups.

5.1.3. NVE-NVE Data Plane

As specified in [[I-D.ietf-nvo3-framework](#)], a NV03 overlay needs to generate tunnels between NVEs for data packet transportation. When a data packet reaches the boundary of a overlay, the ingress NVE will encapsulate the packet and forward it to the destination egress NVE through a proper tunnel.

REQ 10: The security solution for NV03 SHOULD enable two NVEs to mutually authenticate each other before establishing a tunnel connecting them 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 11: The security solution of NV03 MUST 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 a underlying network devices on the path from replaying antique packets or injecting bogus data packets without being detected.

REQ 12: The security solution of NV03 MAY provide confidentiality protection for data traffics exchanged between NVEs.

If the data traffics from the TSeS are sensitive, they needs to be encrypted when being transported within the overlay. Otherwise, encryption will be unnecessary. In addition, in practice, tenants may also select to encrypt their sensitive data during transportation. Therefore this confidentiality requirement for data plane is then not as crucial as the integrity requirement.

REQ 13: 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 others. 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 14: 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 practice, a NVE may need to use the multicast capability provided by the underlying network to transfer multicast packets to other NVEs. In this case, in order to provide an essential packet level security protection specified in requirements 11 and 12, at least a group key may need to be shared among the NVEs in a same multicast group, in order to provide packet level authentication or optionally confidentiality protection for the multicast packets transferred within the group. It is recommended to deploy different keys for different multicast groups, in order to confine the insider attacks on NVEs.

REQ 15: Upon receiving a data packet, an egress NVE must be able to verify whether the packet is from a proper ingress NVE which is authorized to forward that packet.

In cooperation with authentication, authorization enables a egress NVE to detect the data packets which violate certain security policies, even when they are forwarded from a legal NVE. For instance, if a data packet belonging to a VN is forwarded from an ingress NVE which is not supposed to support that VN, the packet needs to be detected and discarded. Note that the detection of a 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.

5.2. Control/Data Plane between NVEs and Hypervisors

Apart from data traffics, the NVE and hypervisors may 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](#)].

A NVE and the hypervisors working with it can be deployed in a distributed way (e.g., the NVE is implemented in an individual

device, and the hypervisors are located on servers) or in a co-located way (e.g., the NVE and the hypervisors are located on the same server). In the former case, the data and control traffic between the NVE and the hypervisors are exchanged over network.

5.2.1. Distributed Deployment of NVE and Hypervisor

Five security requirements applicable for both control and data packets exchanged between NVEs and hypervisors are listed as follows:

REQ 16: The security solution for NV03 SHOULD enable the communicating NVE and hypervisor to mutually authenticate each other before exchanging any control/ data packets.

Mutual authentication is used to prevent an attacker from impersonating a legal NVE or a hypervisor without being detected and then reduce the risks of man-in-the-middle attacks. A successful authentication normally results in the distribution key materials to protect the security of subsequent communications.

REQ 17: The security solution of NV03 MUST be able to provide integrity protection, replay protection and origin authentication for the control/ 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 control packets into the network. In addition, because it is assumed 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 VN identification information in the packet headers so as to manipulate the NVE to transport the data packets within a VN to another.

REQ 18: If a NVE needs to communicate with multiple hypervisors, the security solution of a NV03 network SHOULD be able to provide different keys and ciphers to secure the control /data packets exchanged between different hypervisors and their NVEs respectively.

This requirement is used to benefit the damage confinement of inside attacks. For instance, the compromise of a hypervisor will not affect the security of control/data traffics exchanged between the NVE and other hypervisors.

REQ 19: Before accepting a control/data packet, a NVE or a hypervisor receiving the packet **MUST** verify that the device sending the packet is authorized to do so.

This is an authorization requirement. When receiving a control/data packet, besides authentication, authorization needs to be carried out by a NVE or a hypervisor to identify the role that the packet sender acts as and then assess the sender's privileges. Therefore, if a compromised hypervisor attempts to use its credentials to impersonate a NVE to communicate with other hypervisors, it will be detected.

REQ 20: The security solution of a NV03 network **SHOULD** be able to provide different security levels of protections for the control/data traffics exchanged between a NVE or a hypervisor.

The control and data traffics between a NVE and a hypervisor may be transported over the same path or even within the same security channel. However, when the control traffics and data traffics have different levels of sensitivity, the protection on them needs to be different. In this case, the security solution may need to use different security channels for control and data traffics respectively and so protect the data and control traffics exchanged between a hypervisor and a NVE with different keys and ciphers.

5.2.1.1. Control Plane

REQ 21: The security solution of a NV03 network **MAY** provide confidentiality protection for the control traffics exchanged between a NVE and a hypervisor.

The contents of the control/data packets need to be encrypted when they are considered to be sensitive.

Similar to REQs 6 and 7, the following two requirements are used to mitigate potential DDoS risks.

REQ 22: The frequency in forwarding control packets from a NVE or a hypervisors **MUST** be limited.

This is a common security requirement that can effectively avoid the capability of a device in processing control packets to be overwhelmed by the high frequent control packets generated by the devices attached to it.

REQ 23: Amplification effect **SHOULD** be Addressed.

If the responses generated by a NVE or a hypervisor are much longer than the received requests, an attacker may take advantage of the device as a reflector to perform DDoS attacks. Specifically, the attacker sends a large amount of spoofed short requests to NVEs or hypervisors with the source address of a victim. The responses will then be generated by the NVEs and forwarded to the victim and overwhelm its process capability. This issues should be considered in the design of the control protocols.

5.2.1.2. Data Plane

REQ 24: The security solution of a NV03 network MUST provide security gateways to control the data traffics across the boundaries of different VNs according to specified security policies.

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).

REQ 25: The security solution of a NV03 network MAY provide confidentiality protection for the data traffics exchanged between a NVE and a hypervisor.

When the contents of the data packets are sensitive to a tenant, 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 and the tenants cannot encrypt their sensitive data themselves.

6. Candidate Techniques

This section introduces the techniques which can potentially be used to fulfill the security requirements introduced in [Section 5](#).

6.1. 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.

The widely adopted protocols supporting entity authentication include: IKE[RFC2409], IKEv2[RFC4306], 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.

6.2. 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]). 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.

6.3. 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.

7. IANA Considerations

This document makes no request of IANA.

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

8. Security Considerations

8.1. Automated Key Management in NV03

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.

[8.2.](#) Issues not Discussed

Because this memo only tries to provide the most essential high level requirements, some important issues in designing concret security mechanisms are not covered in the requirements. Such issues include:

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

[9.](#) Acknowledgements

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

[10.](#) References

[10.1.](#) Normative References

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

10.2. 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-04](#) (work in progress), November 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-06](#) (work in progress), November 2013.
- [RFC2328] Moy, J., "OSPF Version 2", STD 54, [RFC 2328](#), April 1998.
- [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.

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>

