

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: December 17, 2014

P. Ashwood-Smith
L. Xia
Huawei Technologies
R. Iyengar
T. Tsou
Huawei Technologies USA
A. Sajassi
Cisco Technologies
M. Boucadair
C. Jacquenet
France Telecom
M. Daikoku
KDDI corporation
A. Ghanwani
Dell
R. Krishnan
Brocade
June 15, 2014

NV03 Operations, Administration, and Maintenance Requirements
draft-ashwood-nvo3-oam-requirements-01

Abstract

This document provides framework and requirements for Network Virtualization Overlay (NV03) Operations, Administration, and Maintenance (OAM). This document for the most part gathers requirements from existing IETF drafts and RFCs which have already extensively studied this subject for different data planes and layering. As a result this draft is high level and broad. We begin to ask which are truly required for NV03 and expect the list to be narrowed by the working group as subsequent versions of this draft are created.

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 17, 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	3
1.1.	OSI Definitions of OAM	3
1.2.	Requirements Language	5
1.3.	Relationship with Other OAM Work	6
2.	Terminology	6
3.	NV03 Reference Model	6
4.	OAM Framework for NV03	8
4.1.	OAM Layering	8
4.2.	OAM Domains	9
5.	NV03 OAM Requirements	10
5.1.	Discovery	10
5.2.	Connectivity Fault Management	10
5.2.1.	Connectivity Fault Detection	10
5.2.2.	Connectivity Fault Verification	11
5.2.3.	Connectivity Fault localization	11
5.2.4.	Connectivity Fault Notification and Alarm Suppression	11
5.3.	Frame Loss	11
5.4.	Frame Delay	11
5.5.	Frame Delay Variation	11
5.6.	Frame Throughput	12
5.7.	Frame Discard	12
5.8.	Availability	12
5.9.	Data Path Forwarding	12
5.10.	Scalability	12
5.11.	Extensibility	12
5.12.	Security	13

5.13.	Transport Independence	13
5.14.	Application Independence	13
5.15.	Prioritization	13
5.16.	Logging and Traceability Requirements	14
6.	Items for Further Discussion	15
7.	IANA Considerations	18
8.	Security Considerations	18
9.	Acknowledgements	18
10.	References	18
10.1.	Normative References	18
10.2.	Informative References	18
	Authors' Addresses	19

[1.](#) Introduction

This document provides framework and requirements for Network Virtualization Overlay (NV03) Operations, Administration, and Maintenance (OAM). Given that this OAM subject is far from new and has been under extensive investigation by various IETF working groups (and several other standards bodies) for many years, this document draws from existing work, starting with [[RFC6136](#)]. As a result, sections of [[RFC6136](#)] have been reused with minor changes with the permission of the authors.

NV03 OAM requirements are expected to be a subset of IETF/IEEE etc. work done so far; however, we begin with a full set of requirements and expect to prune them through several iterations of this document.

[1.1.](#) OSI Definitions of OAM

The scope of OAM for any service and/or transport/network infrastructure technologies can be very broad in nature. OSI has defined the following five generic functional areas commonly abbreviated as "FCAPS" [[NM-Standards](#)]:

- o Fault Management,
- o Configuration Management,
- o Accounting Management,
- o Performance Management, and
- o Security Management.

This document focuses on the Fault, Performance and to a limited extent the Configuration Management aspects. Other functional

aspects of FCAPS and their relevance (or not) to NV03 are for further study.

Fault Management can typically be viewed in terms of the following categories:

- o Fault Detection;
- o Fault Verification;
- o Fault Isolation;
- o Fault Notification and Alarm Suppression;
- o Fault Recovery.

Fault detection deals with mechanism(s) that can detect both hard failures such as link and device failures, and soft failures, such as software failure, memory corruption, misconfiguration, etc. Fault detection relies upon a set of mechanisms that first allow the observation of an event, then the use of a protocol to dynamically notify a network/system operator (or management system) about the event occurrence, then the use of diagnostic tools to assess the nature and severity of the fault.

After verifying that a fault has occurred along the data path, it is important to be able to isolate the fault to the level of a given device or link. Therefore, a fault isolation mechanism is needed in Fault Management. A fault notification mechanism should be used in conjunction with a fault detection mechanism to notify the devices upstream and downstream to the fault detection point. The fault notification mechanism should also notify NMS systems.

The terms "upstream" and "backward" are used here to denote the direction(s) from which data traffic is flowing. The terms "downstream" and "forward" denote the direction(s) to which data traffic is forwarded.

For example, when there is a client/server relationship between two layered networks (e.g., the NV03 layer is a client of the outer IP server layer, while the inner IP layer is a client of the NV03 server layer 2), fault detection at the server layer may result in the following fault notifications:

- o Sending a forward fault notification from the server layer to the client layer network(s) using the fault notification format appropriate to the client layer.

- o Sending a backward fault notification to the server layer, if applicable, in the reverse direction.
- o Sending a backward fault notification to the client layer, if applicable, in the reverse direction.

Finally, fault recovery deals with recovering from the detected failure by switching to an alternate available data path (depending on the nature of the fault) using alternate devices or links. In fact, the controller can provision another virtual network, thus automatically resolving the reported problem.

The controller may also directly monitor the status of virtual network components such as Network Virtualization Edge elements (NVEs) [[NV03-framework](#)] in order to respond to their failures. In addition to forward and backward fault notifications, the controller may deliver notifications to a higher level orchestration component, e.g., one responsible for Virtual Machine (VM) provisioning and management.

Note, given that the IP network on which NV03 resides is usually self healing, it is expected that recovery by the NV03 layer would not normally be required, although there may be a requirement for that layer to log that the problem has been detected and resolved. The special cases of a static IP overlay network, or possibly of a centrally controlled IP overlay network, may, however, require NV03 involvement in fault recovery.

Performance Management deals with mechanism(s) that allow determining and measuring the performance of the network/services under consideration. Performance Management can be used to verify the compliance to both the service-level and network-level metric objectives/specifications. Performance Management typically consists of measuring performance metrics, e.g., Frame Loss, Frame Delay, Frame Delay Variation (aka Jitter), Frame Throughput, Frame Discard, etc., across managed entities when the managed entities are in available state. Performance Management is suspended across unavailable managed entities.

1.2. 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 [[RFC2119](#)].

1.3. Relationship with Other OAM Work

This document leverages requirements that originate with other OAM work, specifically the following:

- o [\[RFC6136\]](#) provides a template and some of the high level requirements and introductory wording.
- o [\[IEEE802.1Q-2011\]](#) is expected to provide a subset of the requirements for NV03 both at the Tenant level and also within the L3 Overlay network.
- o [\[Y.1731\]](#) is expected to provide a subset of the requirements for NV03 at the Tenant level.
- o Section 3.3.2.1 of [\[NV03-DP-Reqs\]](#) lists several requirements specifically concerning ECMP/LAG.

2. Terminology

The terminology defined in [\[NV03-framework\]](#) and [\[NV03-DP-Reqs\]](#) is used throughout this document. We introduce no new terminology.

3. NV03 Reference Model

Figure 1 below reproduces the generic NV03 reference model as per [\[NV03-framework\]](#).

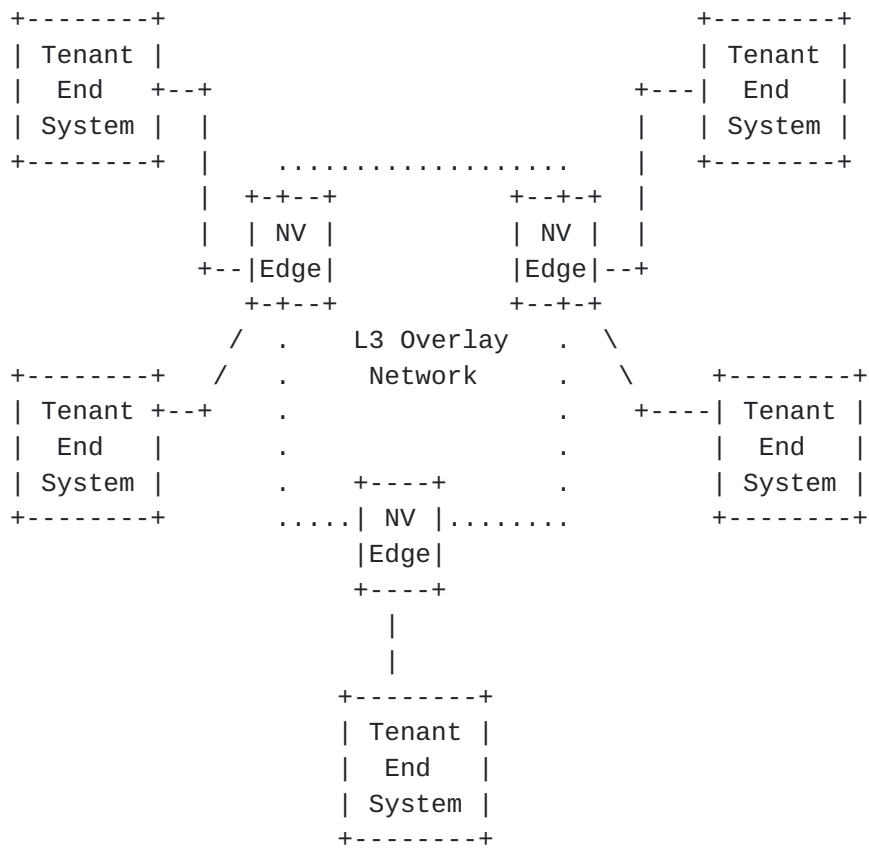


Figure 1: Generic reference model for DC network virtualization over a Layer3 infrastructure

Figure 2 below, reproduces the Generic reference model for the NV Edge (NVE) as per [[NV03-DP-Reqs](#)].

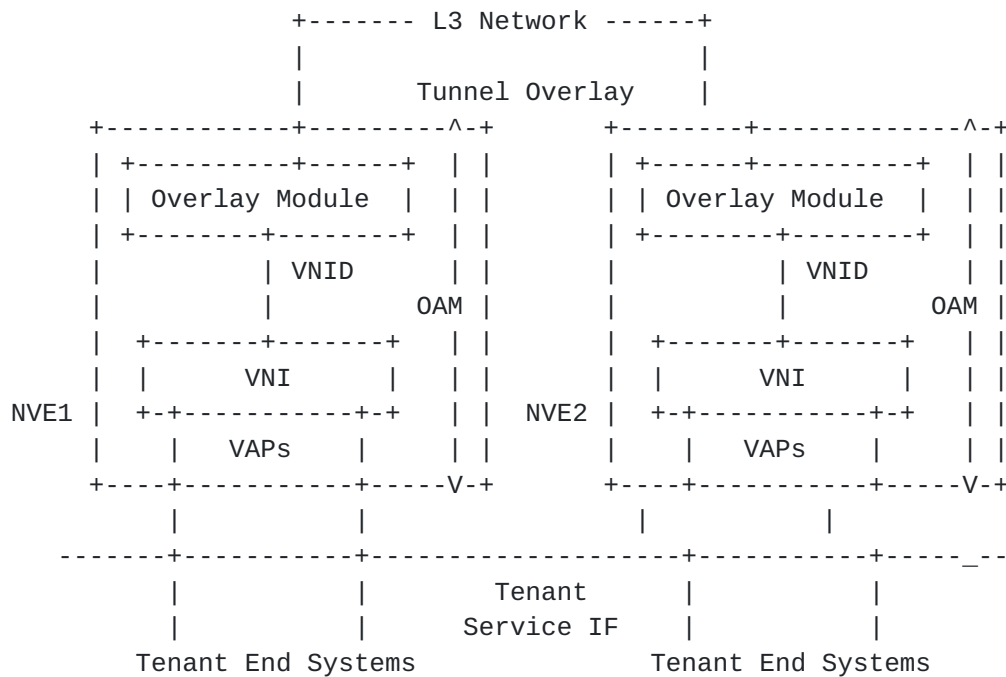


Figure 2: Generic reference model for NV Edge

4. OAM Framework for NV03

Figure 1 showed the generic reference model for a DC network virtualization over an L3 (or L3VPN) infrastructure while Figure 2 showed the generic reference model for the Network Virtualization (NV) Edge.

L3 network(s) or L3 VPN networks (either IPv6 or IPv4, or a combination thereof), provide transport for an emulated layer 2 created by NV Edge devices. Unicast and multicast tunneling methods (de-multiplexed by Virtual Network Identifier (VNID)) are used to provide connectivity between the NV Edge devices. The NV Edge devices then present an emulated layer 2 network to the Tenant End Systems at a Virtual Network Interface (VNI) through Virtual Access Points (VAPs). The NV Edge devices map layer 2 unicast to layer 3 unicast point-to-point tunnels and may either map layer 2 multicast to layer 3 multicast tunnels or may replicate packets onto multiple layer 3 unicast tunnels.

4.1. OAM Layering

The emulated layer 2 network is provided by the NV Edge devices to which the Tenant End Systems are connected. This network of NV Edges can be operated by a single service provider or can span across multiple administrative domains. Likewise, the L3 Overlay Network

can be operated by a single service provider or span across multiple administrative domains.

While each of the layers is responsible for its own OAM, each layer may consist of several different administrative domains. Figure 3 shows an example.

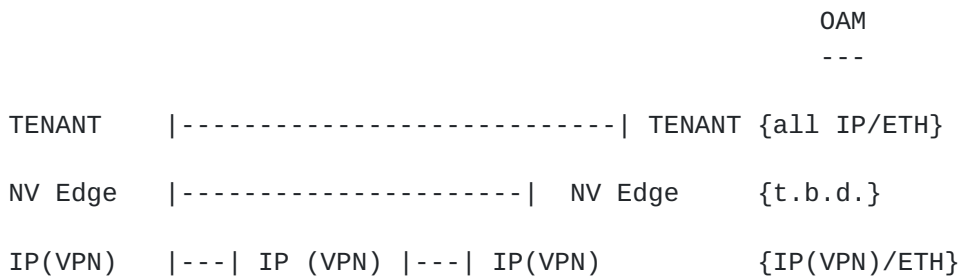


Figure 3: OAM layers in an NV03 network

For example, at the bottom, at the L3 IP overlay network layer IP(VPN) and/or Ethernet OAM mechanisms are used to probe link by link, node to node etc. OAM addressing here means physical node loopback or interface addresses.

Further up, at the NV Edge layer, NV03 OAM messages are used to probe the NV Edge to NV Edge tunnels and NV Edge entity status. OAM addressing here likely means the physical node loopback together with the VNI (to de-multiplex the tunnels).

Finally, at the Tenant layer, the IP and/or Ethernet OAM mechanisms are again used but here they are operating over the logical L2/L3 provided by the NV-Edge through the VAP. OAM addressing at this layer deals with the logical interfaces on Vswitches and Virtual Machines.

[4.2.](#) OAM Domains

Complex OAM relationships exist as a result of the hierarchical layering of responsibility and of breaking up of end-to-end responsibility.

The OAM domain above NV03, is expected to be supported by existing IP and L2 OAM methods and tools.

The OAM domain below NV03, is expected to be supported by existing IP/L2 and MPLS OAM methods and tools. Where this layer is actually multiple domains spliced together, the existing methods to deal with these boundaries are unchanged. Note however that exposing LAG/ECMP detailed behavior may result in additional requirements to this

domain, the details of which will be specified in the future versions of this draft.

When we refer to an OAM domain in this document, or just 'domain', we therefore refer to a closed set of NV Edges or MEPs and the tunnels which interconnect them.

Note, whether for the scenario of inter-domain or multi-layer, each domain (or layer) is responsible for its own OAM, no correlation of OAM function exists between each domain (or layer). When an E2E connection in Tenant layer spans across multiple domains and has multiple underlay layers of NV Edge layer and L3 IP (VPN) layer, current OAM implementation for the E2E connection of Tenant layer such as Fault or Performance Management can only be performed per domain and layer manually and more manual labor is needed. An automatic coordination process among OAM functions of each domain or layer may be useful here for improving efficiency and intelligence.

In the case where a gateway device is used to connect two different domains (whether for changing the encapsulation or other reasons), it is necessary to provide mechanisms to monitor the path through the gateway which involves the removal of one overlay header and the creation of a new one.

5. NV03 OAM Requirements

The following numbered requirements originate from [\[RFC6136\]](#). All are included however where they seem obviously not relevant (to the present authors) an explanation as to why is included.

5.1. Discovery

R1) NV03 OAM MUST allow an NV Edge device to dynamically discover other NV Edge devices that share the same VNI within a given NV03 domain. This may be based on a discovery mechanism used to set up data path forwarding between NVEs.

5.2. Connectivity Fault Management

5.2.1. Connectivity Fault Detection

R2) NV03 OAM MUST allow proactive connectivity monitoring between two or more NV Edge devices that support the same VNIs within a given NV03 domain. NV03 OAM MAY act as a protection trigger. That is, automatic recovery from transmission facility failure by switchover to a redundant replacement facility may be triggered by notifications from NV03 OAM.

R3) NV03 OAM MUST allow monitoring/tracing of all possible paths in the underlay network between a specified set of two or more NV Edge devices. Using this feature, equal cost paths that traverse LAG and/or ECMP may be differentiated.

5.2.2. Connectivity Fault Verification

R4) NV03 OAM MUST allow connectivity fault verification between two or more NV Edge devices that support the same VNI within a given NV03 domain.

5.2.3. Connectivity Fault localization

R5) NV03 OAM MUST allow connectivity fault localization between two or more NV Edge devices that support the same VNI within a given NV03 domain.

5.2.4. Connectivity Fault Notification and Alarm Suppression

R6) NV03 OAM MUST support fault notification to be triggered as a result of the faults occurring in the underneath network infrastructure. This fault notification SHOULD be used for the suppression of redundant service-level alarms.

5.3. Frame Loss

R7) NV03 OAM MUST support measurement of per VNI frame loss between two NV Edge devices that support the same VNI within a given NV03 domain.

5.4. Frame Delay

R8) NV03 OAM MUST support measurement of per VNI two-way frame delay between two NV edge devices that support the same VNI within a given NV03 domain.

R9) NV03 OAM MUST support measurement of per VNI one-way frame delay between two NV Edge devices that support the same VNI within a given NV03 domain.

5.5. Frame Delay Variation

R10) NV03 OAM MUST support measurement of per VNI frame delay variation between two NV Edge devices that support the same VNI within a given NV03 domain.

5.6. Frame Throughput

R11) NV03 OAM MAY [*** Should this be stronger? ***] support measurement of per VNI frame throughput (in frames and bytes) between two NV Edge devices that support the same VNI within a given NV03 domain. This feature could be an effective way to confirm whether or not assigned path bandwidth conforms to service level agreement before providing the path between two NV Edge devices.

5.7. Frame Discard

R12) NV03 OAM MAY support measurement of per VNI frame discard between two NV Edge devices that support the same VNI within a given NV03 domain. This feature MAY be effective to monitor bursty traffic between two NV Edge devices.

5.8. Availability

A service may be considered unavailable if the service frames/packets do not reach their intended destination (e.g., connectivity is down) or the service is degraded (e.g., frame loss and/or frame delay and/or delay variation threshold is exceeded). Entry and exit conditions may be defined for the unavailable state. Availability itself may be defined in the context of a service type. Since availability measurement may be associated with connectivity, frame loss, frame delay, and frame delay variation measurements, no additional requirements are specified currently.

5.9. Data Path Forwarding

R13) NV03 OAM frames MUST be forwarded along the same path (i.e., links (including LAG members) and nodes) as the NV03 data frames.

R14) NV03 OAM frames MUST provide a mechanism to exercise/trace all data paths that result due to ECMP/LAG hops in the underlay network.

5.10. Scalability

R15) NV03 OAM MUST be scalable such that an NV edge device can support proactive OAM for each VNI that is supported by the device. (Note - Likely very hard to achieve with hash based ECMP/LAG).

5.11. Extensibility

R16) NV03 OAM should be extensible such that new functionality and information elements related to this functionality can be introduced in the future.

R17) NV03 OAM MUST be defined such that devices not supporting the OAM are able to forward the OAM frames in a similar fashion as the regular NV03 data frames/packets.

5.12. Security

R18) NV03 OAM frames MUST be prevented from leaking outside their NV03 domain.

R19) NV03 OAM frames from outside an NV03 domain MUST be prevented from entering the said NV03 domain when such OAM frames belong to the same level or to a lower-level OAM. (Trivially met because hierarchical domains are independent technologies.)

R20) NV03 OAM frames from outside an NV03 domain MUST be transported transparently inside the NV03 domain when such OAM frames belong to a higher-level NV03 domain. (Trivially met because hierarchical domains are independent technologies.)

5.13. Transport Independence

Similar to transport requirement from [RFC6136], we expect NV03 OAM will leverage the OAM capabilities of the transport layer (e.g., IP underlay).

R21) NV03 OAM MAY allow adaptation/interworking with its IP underlay OAM functions. For example, this would be useful to allow fault notifications from the IP layer to be sent to the NV03 layer and likewise exposure of LAG / ECMP will require such non-independence.

5.14. Application Independence

R22) NV03 OAM MUST [*** discuss -- is this too strong? ***] be independent of the application technologies and specific application OAM capabilities.

[Comment -- ECM: Noticed Nicira implementation has a dedicated NVP manager node to play the role of FCAPS here. It is both application layer and OAM layer. May not meet this requirement. In reality, due to the nature of overlay network, very often, vendors are going to make everything all together to a dedicated manager node.]

5.15. Prioritization

R23) NV03 OAM messages MUST be preferentially treated in NVE and between NVEs, since NV03 OAM MAY be used to trigger protection switching. As noted above (R2), protection switching is the automatic replacement of a failed transmission facility with a

working one providing equal or greater capacity, typically within a few tens of milliseconds from fault detection.

[Comment -- ECM: giving NV03 OAM messages priority treatment may interfere with measurements of frame delay and jitter.]

5.16. Logging and Traceability Requirements

Logging is required at the Network Virtualization Authority (NVA) and the Network Virtualization Edge (NVE) [and the NV03 Gateway, but the framework does not mention such a beast] in support of fault management and configuration management.

R24) All logs MUST contain a timestamp of sufficient accuracy and precision to allow accurate determination of the sequence of events at the reporting functional instance (i.e., NVA, NVE). Clocks on different functional instances SHOULD be synchronized to allow similar accuracy when comparing logs from different devices.

R25) All logs MUST identify the reporting functional instance.

R26) Implementations MUST be capable of reporting the following fault- related events:

1. Loss and resumption of connectivity

These reports SHOULD identify the affected VNI(s), but when the loss affects a large number of VNIs simultaneously the report SHOULD identify the underlying entity (e.g., route) if available.

2. Loss and resumption of NVE responsiveness

These reports will be generated by adjacent NVAs or NVEs. They MUST identify the NVE concerned.

3. NVA or NVE change of operational state

These reports will be generated by the NVA or NVE concerned. They MUST indicate the old and new operational states and the cause.

4. Loss and resumption of a VAP

[What else?]

R27) Implementations MUST be capable of reporting the following events in support of configuration management and auditing. It MUST be possible to generate the reports at both the originating and

executing entities. The report generated at the originating entity MUST identify the executing entity and the report at the executing entity MUST identify the originating entity. Both reports MUST indicate the result of the transaction.

1. Virtual Access Point (VAP) creation or deletion

These reports MUST identify the VAP, the Tenant System, and the port supporting the VAP.

2. VNI creation or deletion

These reports MUST identify the VNI and the VAP.

3. VNI renumbering

These reports MUST identify the VAP and the old and new VNI numbers.

4. Reachability and forwarding information update

These reports MUST identify the previous and new file identifiers. (Assumption: reachability and forwarding information is passed as files, which are retained at the originating and executing entities for a fixed period for auditing purposes.)

[What else?]

R28) As a general requirement, implementations MUST provide a means whereby the operator can impose rate limits on the generation of specific reports. Implementations MUST further permit the operator to totally suppress reporting of specific events. However, if any report types have been suppressed, non-suppressible reports MUST be generated at regular intervals (e.g., once an hour) indicating what report types have been suppressed.

6. Items for Further Discussion

This section identifies a set of operational items which may be elaborated further if these items fall within the scope of the NV03.

o VNID renumbering support

- * Means to change the VNID assigned to a given instance MUST [*** discuss: is this too strong? ***] be supported.

- * System convergence subsequent to VNID renumbering MUST NOT take longer than a few seconds, to minimize impact on the tenant systems.
- * A NVE MUST be able to map a VNID with a virtual network context.
- o VNI migration and management operations
 - * Means to delete an existing VNI MUST be supported.
 - * Means to add a new VNI MUST be supported.
 - * Means to merge several VNIs MAY be supported.
 - * Means to retrieve reporting data per VNI MUST be supported.
 - * Means to monitor the network resources per VNI MUST be supported.
- o Support of planned maintenance operations on the NV03 infrastructure
 - * Graceful procedure to allow for planned maintenance operation on NVE MUST be supported. This includes undoing any configuration changes made for maintenance purposes after completion of the maintenance.
- o Support for communication among virtual networks
 - * For global reachability purposes, communication among virtual networks MUST be supported. This can be enforced using a NAT function.
- o Activation of new network-related services to the NV03
 - * Means to assist in activating new network services (e.g., multicast) without impacting running service SHOULD be supported.
- o Inter-operator NV03 considerations
 - * As NV03 may be deployed over inter-operator infrastructure, coordinating OAM actions in each individual domain are required to ensure an end-to-end OAM. In particular, this assumes existence of agreements on the measurement and monitoring methods, fault detection and repair actions, extending QoS classes (e.g., DSCP mapping policies), etc.

- + [[DISCUSSION NOTE: Should inter-operator issues be declared out of scope?]]
- o An automatic coordination process among OAM functions of different domains or layers which an E2E connection in Tenant layer is tunnelled on
 - * NV03 OAM MAY support the automatic coordination of OAM functions among different domains or layers which belong to one Tenant layer E2E connection. The automatic coordination means OAM function in client layer or one domain triggers associated OAM functions in server layer or neighbouring domain. This triggered action performs at the domain boundaries, which is also the MEPs of the domain. Which OAM function in client layer or one domain can trigger which OAM functions in server layer or neighbouring domain depends on specific condition, and can be very flexible. But the basic rule is that the OAM functions performed simultaneously in different domains or layers can be synthesized together to get the final result.
 - * The OAM MEPs of domains MUST have the capability to know if it needs to perform the above automatic coordination process. This can be achieved by many ways, i.e., by configuration, by checking the flag field in OAM frames.
 - * When the OAM MEPs perform the automatic coordination, a specific global characteristic information MUST be carried and mapped between OAM frames used in different domains or layers, and be kept the same along the whole tenant layer E2E connection. The global characteristic information can be the tenant network identifier (e.g., VNID), ICMP sequence number, etc. It is used for identifying a set of correlated OAM results obtained from these domains or layers. This set of OAM results is then synthesized together to get the final diagnose result.
 - * NV03 OAM MUST support a Collection Point for collecting all the OAM results and synthesizing them. It can be the SDN controller, NVA, or NMS. An E2E OAM function in tenant network can trigger several OAM functions in different underlay networks, a Collection Point is needed to collect all the OAM results from different OAM MEPs of different domains or layers and synthesizes them.

7. IANA Considerations

This memo includes no request to IANA.

8. Security Considerations

TBD

9. Acknowledgements

The authors are grateful for the contributions of David Black, Dennis Qin, Erik Smith and Ziyue Yang to this latest version.

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

[I2RS-trace]

Clark, J., Salgueiro, G., and C. Pignataro, "Interface to the Routing System (I2RS) Traceability: Framework and Information Model", December 2013.

[IEEE802.1Q-2011]

IEEE, "IEEE Standard for Local and metropolitan area networks - Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks", 2011.

[NM-Standards]

ITU-T, "ITU-T Recommendation M.3400 (02/2000) - TMN Management Functions", February 2000.

[NV03-DP-Reqs]

Bitar, N., Lasserre, M., Balus, F., Morin, T., Jin, L., and B. Khasnabish, "NV03 Data Plane Requirements", April 2014.

[NV03-framework]

Lasserre, M., Balus, F., Morin, T., Bitar, N., and Y. Rekhter, "Framework for DC Network Virtualization", June 2014.

- [RFC6136] Sajassi, A. and D. Mohan, "Layer 2 Virtual Private Network (L2VPN) Operations, Administration, and Maintenance (OAM) Requirements and Framework", [RFC 6136](#), March 2011.
- [Y.1731] ITU-T, "ITU-T Recommendation Y.1731 (02/08) - OAM functions and mechanisms for Ethernet based networks", February 2008.

Authors' Addresses

Peter Ashwood-Smith
Huawei Technologies
303 Terry Fox Drive, Suite 400
Kanata, Ontario K2K 3J1
Canada

Phone: +1 613 595-1900
Email: Peter.AshwoodSmith@huawei.com

Liang Xia (Frank)
Huawei Technologies

Email: Frank.xialiang@huawei.com

Ranga Iyengar
Huawei Technologies USA
2330 Central Expy
Santa Clara, CA 95050
USA

Email: ranga.Iyengar@huawei.com

Tina Tsou
Huawei Technologies USA
2330 Central Expy
Santa Clara, CA 95050
USA

Email: Tina.Tsou.Zouting@huawei.com

Ali Sajassi
Cisco Technologies
170 West Tasman Drive
San Jose, CA 95134
USA

Email: sajassi@cisco.com

Mohamed Boucadair
France Telecom
Rennes 35000
France

Email: mohamed.boucadair@orange.com

Christian Jacquenet
France Telecom
Rennes 35000
France

Email: christian.jacquenet@orange.com

Masahiro Daikoku
KDDI corporation
3-10-10, Iidabashi, Chiyoda-ku
Tokyo 1028460
Japan

Email: ms-daikoku@kddi.com

Anoop Ghanwani
Dell
5450 Great America Pkwy
Santa Clara, CA
USA

Email: anoop@alumni.duke.edu

Ram Krishnan
Brocade
130 Holger Way
San Jose, CA 95134
USA

Email: ramk@brocade.com