

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: May 3, 2012

T. Narten, Ed.
IBM
M. Sridharan
Microsoft
D. Dutt
Cisco
D. Black
EMC
L. Kreeger
Cisco
October 31, 2011

Problem Statement: Overlays for Network Virtualization
draft-narten-nvo3-overlay-problem-statement-01

Abstract

This document describes issues associated with providing multi-tenancy in large data center networks and an overlay-based network virtualization approach to addressing them. A key multi-tenancy requirement is traffic isolation, so that a tenant's traffic is not visible to any other tenant. This isolation can be achieved by assigning one or more virtual networks to each tenant such that traffic within a virtual network is isolated from traffic in other virtual networks. The primary functionality required is provisioning virtual networks, associating a virtual machine's NIC with the appropriate virtual network, and maintaining that association as the virtual machine is activated, migrated and/or deactivated. Use of an overlay-based approach enables scalable deployment on large network infrastructures.

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 May 3, 2012.

Internet-Draft

Overlays for Network Virtualization

October 2011

Copyright Notice

Copyright (c) 2011 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.

Internet-Draft Overlays for Network Virtualization October 2011

Table of Contents

1.	Introduction	4
2.	Problem Details	5
2.1.	Multi-tenant Environment Scale	5
2.2.	Virtual Machine Mobility Requirements	5
2.3.	Span of Virtual Networks	5
2.4.	Inadequate Forwarding Table Sizes in Switches	6
2.5.	Decoupling Logical and Physical Configuration	6
2.6.	Support Communication Between VMs and Non-virtualized Devices	6
2.7.	Overlay Design Characteristics	6
3.	Defining Virtual Networks and Tenants	7
3.1.	Limitations of Existing Virtual Network Models	8
3.2.	Virtual Network Instance	8
3.3.	Tenant	9
4.	Network Overlays	9
4.1.	Benefits of an Overlay Approach	10
4.2.	Standardization Issues for Overlay Networks	10
4.2.1.	Overlay Header Format	10
4.2.2.	Fragmentation	11
4.2.3.	Checksums and FCS	11
4.2.4.	Middlebox Traversal	12
4.2.5.	OAM	12
5.	Control Plane	12
5.1.	Populating the Forwarding Table of a Virtual Network Instance	12
5.2.	Handling Multi-destination Frames	13
5.3.	Associating a VNID With An Endpoint	13
5.4.	Disassociating a VNID on Termination or Move	13
6.	Related Work	13
6.1.	ARMD	13
6.2.	TRILL	14
6.3.	L2VPNs	14
6.4.	Proxy Mobile IP	14
6.5.	LISP	14

6.6.	Individual Submissions	15
7.	Further Work	15
8.	Summary	15
9.	Acknowledgments	15
10.	IANA Considerations	15
11.	Security Considerations	15
12.	Informative References	16
	Authors' Addresses	17

[1.](#) Introduction

Server virtualization is increasingly becoming the norm in data centers. With server virtualization, each physical server supports multiple virtual machines (VMs), each running its own operating system, middleware and applications. Virtualization is a key enabler of workload agility, i.e., allowing any server to host any application and providing the flexibility of adding, shrinking, or moving services within the physical infrastructure. Server virtualization provides numerous benefits, including higher utilization, increased data security, reduced user downtime, reduced power usage, etc.

Large scale multi-tenant data centers are taking advantage of the benefits of server virtualization to provide a new kind of hosting, a virtual hosted data center. Multi-tenant data centers are ones in which each tenant could belong to a different company (in the case of a public provider) or a different department (in the case of a internal company data center). Each tenant has the expectation of a level of security and privacy separating their resources from those of other tenants. Each virtual data center looks similar to its physical counterpart, consisting of end stations connected by a network, complete with services such as load balancers and firewalls. The network within each virtual data center can be a pure routed network, a pure bridged network or a combination of bridged and routed network. The key requirement is that each such virtual network is isolated from the others, whether the networks belong to the same tenant or different tenants.

This document outlines the problems encountered in scaling the number of isolated networks in a data center, as well as the problems of managing the creation/deletion, membership and span of these networks and makes the case that an overlay based approach, where individual networks are implemented within individual virtual networks that are dynamically controlled by a standardized control plane provides a number of advantages over current approaches. The purpose of this document is to identify the set of problems that any solution has to address in building multi-tenant data centers. With this approach, the goal is to allow the construction of standardized, interoperable implementations to allow the construction of multi-tenant data centers.

[Section 2](#) describes the problem space details. [Section 3](#) defines virtual networks. [Section 4](#) provides a general discussion of overlays and standardization issues. [Section 5](#) discusses the control plane issues that require addressing for virtual networks. [Section 6](#) and [7](#) discuss related work and further work.

[2.](#) Problem Details

The following subsections describe aspects of multi-tenant networking that pose problems for large scale network infrastructure. Different problem aspects may arise based on the network architecture and scale.

[2.1.](#) Multi-tenant Environment Scale

Cloud computing involves on-demand elastic provisioning of resources for multi-tenant environments. A common example of cloud computing is the public cloud, where a cloud service provider offers these elastic services to multiple customers over the same infrastructure. This elastic on-demand nature in conjunction with trusted hypervisors to control network access by VMs calls for resilient distributed network control mechanisms.

[2.2.](#) Virtual Machine Mobility Requirements

A key benefit of server virtualization is virtual machine (VM) mobility. A VM can be migrated from one server to another, live i.e. as it continues to run and without shutting down the VM and

restarting it at a new location. A key requirement for live migration is that a VM retain its IP address(es) and MAC address(es) in its new location (to avoid tearing down existing communication). Today, servers are assigned IP addresses based on their physical location, typically based on the ToR (Top of Rack) switch for the server rack or the VLAN configured to the server. This works well for physical servers, which cannot move, but it restricts the placement and movement of the more mobile VMs within the data center (DC). Any solution for a scalable multi-tenant DC must allow a VM to be placed (or moved to) anywhere within the data center, without being constrained by the subnet boundary concerns of the host servers.

[2.3.](#) Span of Virtual Networks

Another use case is cross pod expansion. A pod typically consists of one or more racks of servers with its associated network and storage connectivity. Tenants may start off on a pod and, due to expansion, require servers/VMs on other pods, especially the case when tenants on the other pods are not fully utilizing all their resources. This use case requires that virtual networks span multiple pods in order to provide connectivity to all of the tenant's servers/VMs.

[2.4.](#) Inadequate Forwarding Table Sizes in Switches

Today's virtualized environments place additional demands on the forwarding tables of switches. Instead of just one link-layer address per server, the switching infrastructure has to learn addresses of the individual VMs (which could range in the 100s per server). This is a requirement since traffic from/to the VMs to the rest of the physical network will traverse the physical network infrastructure. This places a much larger demand on the switches' forwarding table capacity compared to non-virtualized environments, causing more traffic to be flooded or dropped when the addresses in use exceeds the forwarding table capacity.

[2.5.](#) Decoupling Logical and Physical Configuration

Data center operators must be able to achieve high utilization of server and network capacity. For efficient and flexible allocation, operators should be able to spread a virtual network instance across servers in any rack in the data center. It should also be possible to migrate compute workloads to any server anywhere in the network while retaining the workload's addresses. This can be achieved today by stretching VLANs (e.g., by using TRILL or OTV).

However, in order to limit the broadcast domain of each VLAN, multi-destination frames within a VLAN should optimally flow only to those devices that have that VLAN configured. When workloads migrate, the physical network (e.g., access lists) may need to be reconfigured which is typically time consuming and error prone.

[2.6.](#) Support Communication Between VMs and Non-virtualized Devices

Within data centers, not all communication will be between VMs. Network operators will continue to use non-virtualized servers for various reasons, traditional routers to provide L2VPN and L3VPN services, traditional load balancers, firewalls, intrusion detection engines and so on. Any virtual network solution should be capable of working with these existing systems.

[2.7.](#) Overlay Design Characteristics

There are existing layer 2 overlay protocols in existence, but they were not necessarily designed to solve the problem in the environment of a highly virtualized data center. Below are some of the characteristics of environments that must be taken into account by the overlay technology:

1. Highly distributed systems. The overlay should work in an environment where there could be many thousands of access switches (e.g. residing within the hypervisors) and many more end systems (e.g. VMs) connected to them. This leads to a distributed mapping system that puts a low overhead on the overlay tunnel endpoints.
2. Many highly distributed virtual networks with sparse

connectivity. Each virtual network could be highly dispersed inside the data center. Also, along with expectation of many virtual networks, the number of end systems connected to any one virtual network is expected to be relatively low; Therefore, the percentage of access switches participating in any given virtual network would also be expected to be low. For this reason, efficient pruning of multi-destination traffic should be taken into consideration.

3. Highly dynamic end systems. End systems connected to virtual networks can be very dynamic, both in terms of creation/deletion/power-on/off and in terms of mobility across the access switches.
4. Work with existing, widely deployed network Ethernet switches and IP routers without requiring wholesale replacement. The first hop switch that adds and removes the overlay header will require new equipment and/or new software.
5. Network infrastructure administered by a single administrative domain. This is consistent with operation within a data center, and not across the Internet.

3. Defining Virtual Networks and Tenants

Virtual Networks are used to isolate a tenant's traffic from other tenants (or even traffic within the same tenant that requires isolation). There are two main characteristics of virtual networks:

1. Providing network address space that is isolated from other virtual networks. The same network addresses may be used in different virtual networks on the same underlying network infrastructure.
2. Limiting the scope of frames to not exit a virtual network except through controlled exit points or "gateways".

3.1. Limitations of Existing Virtual Network Models

Virtual networks are not new to networking. VLANs are a well known construct in the networking industry. VLAN is a bridging construct which provides the semantics of virtual networks mentioned above: a MAC address is unique within a VLAN, but not necessarily across VLANs and broadcast traffic is limited to the VLAN it originates from. In the case of IP networks, routers have the concept of a Virtual Routing and Forwarding (VRF). The same router can run multiple instances of routing protocols, each with their own forwarding table. Each instance is referred to as a VRF, which is a mechanism that provides address isolation. Since broadcasts are never forwarded across IP subnets, limiting broadcasts are not applicable to VRFs. In the case of both VLAN and VRF, the forwarding table is looked up using the tuple {VLAN, MAC address} or {VRF, IP address}.

But there are two problems with these constructs. VLANs are a pure bridging construct while VRF is a pure routing construct. VLANs are carried along with a frame to allow each forwarding point to know what VLAN the frame belongs to. VLAN today is defined as a 12 bit number, limiting the total number of VLANs to 4096 (though typically, this number is 4094 since 0 and 4095 are reserved). Due to the large number of tenants that a cloud provider might service, the 4094 VLAN limit is often inadequate. In addition, there is often a need for multiple VLANs per tenant, which exacerbates the issue.

There is no VRF indicator carried in frames. The VRF is derived at each hop using a combination of incoming interface and some information in the frame. Furthermore, the VRF model has typically assumed that a separate control plane governs the population of the forwarding table within that VRF. Thus, a traditional VRF model assumes multiple, independent control planes and has no specific tag within a frame to identify the VRF of the frame.

[3.2.](#) Virtual Network Instance

To overcome the limitations of a traditional VLAN or VRF model, we define a new mechanism for virtual networks called a virtual network instance. Each virtual network is assigned a virtual network instance ID, shortened to VNID for convenience. A virtual network instance provides the semantics of a virtual network: address disambiguation and multi-destination frame scoping. A virtual network can be either routed or bridged. So, a VNID can be used for both bridged networks and routed networks and so is unlike a VLAN or a VRF. To build large multi-tenant data centers, a larger number space than the 12b VLAN is required. 24 bits is the most common value identified by multiple solutions that attempt to address this problem space (or similar problem spaces). To simplify the building and

administration of these large data centers, we require that the VNID be carried with each frame (similar to a VLAN, but unlike a VRF). Finally, because of the nature of a virtual data center and to allow scaling virtual networks to massive scales, we don't require a separate control plane to run for each virtual network. We'll identify other possible mechanisms to populate the forwarding tables for virtual networks in [section 5.1](#).

[3.3](#). Tenant

Tenant is the administrative entity that that is responsible for and manages a specific virtual network and its associated services (whether virtual or physical). In a cloud environment, a tenant would correspond to the customer that has defined and is using a particular virtual network. However, there is a one-to-many mapping between tenants and virtual network instances. A single tenant may operate multiple individual virtual networks, each associated with a different service.

[4](#). Network Overlays

To address the problems of decoupling physical and logical configuration and allowing VM mobility without exploding the forwarding table sizes in the switches and routers, a network overlay model can be used.

The idea behind an overlay is quite straightforward. The original frame is encapsulated by the first hop network device. The encapsulation identifies the destination as the device that will perform the decapsulation before delivering the frame to the endpoint. The rest of the network forwards the frame based on the encapsulation header and can be oblivious to the payload that is carried inside. To avoid belaboring the point each time, the first hop network device can be a traditional switch or router or the virtual switch residing inside a hypervisor. Furthermore, the endpoint can be a VM or it can be a physical server. Some examples of network overlays are tunnels such as IP GRE [[RFC2784](#)], LISP[I-D.ietf-lisp] or TRILL [[RFC6325](#)].

With an overlay, the VNID can be carried within the overlay header so that every frame has its VNID explicitly identified in the frame. Since both routed and bridged semantics can be supported by a virtual data center, the original frame carried within the overlay header can be an Ethernet frame complete with MAC addresses or just the IP packet.

[4.1.](#) Benefits of an Overlay Approach

The use of a large (e.g., 24-bit) VNID would allow 16 million distinct virtual networks within a single data center, eliminating current VLAN size limitations. This VNID needs to be carried in the data plane along with the packet. Adding an overlay header provides a place to carry this VNID.

A key aspect of overlays is the decoupling of the "virtual" MAC and IP addresses used by VMs from the physical network infrastructure and the infrastructure IP addresses used by the data center. If a VM changes location, the switches at the edge of the overlay simply update their mapping tables to reflect the new location of the VM within the data center's infrastructure space. Because an overlay network is used, a VM can now be located anywhere in the data center that the overlay reaches without regards to traditional constraints implied by L2 properties such as VLAN numbering, or the span of an L2 broadcast domain scoped to a single pod or access switch.

Multi-tenancy is supported by isolating the traffic of one virtual network instance from traffic of another. Traffic from one virtual network instance cannot be delivered to another instance without (conceptually) exiting the instance and entering the other instance via an entity that has connectivity to both virtual network instances. Without the existence of this entity, tenant traffic remains isolated within each individual virtual network instance. External communications (from a VM within a virtual network instance to a machine outside of any virtual network instance, e.g. on the Internet) is handled by having an ingress switch forward traffic to an external router, where an egress switch decapsulates a tunneled packet and delivers it to the router for normal processing. This router is external to the overlay, and behaves much like existing external facing routers in data centers today.

Overlays are designed to allow a set of VMs to be placed within a single virtual network instance, whether that virtual network provides the bridged network or a routed network.

[4.2.](#) Standardization Issues for Overlay Networks

[4.2.1.](#) Overlay Header Format

Different overlay header formats are possible as are different possible encodings of the VNID. Existing overlay headers maybe extended or new ones defined. This document does not address the exact header format or VNID encoding except to state that any solution MUST:

1. Carry the VNID in each frame
2. Allow the payload to be either a complete Ethernet frame or only an IP packet

[4.2.2.](#) Fragmentation

Whenever tunneling is used, one faces the potential problem that the packet plus the encapsulation overhead will exceed the MTU of the path to the egress router. If the outer encapsulation is IP, fragmentation could be left to the IP layer, or it could be done at the overlay level in a more optimized fashion that is independent of the overlay encapsulation header, or it could be left out altogether, if it is believed that data center networks can be engineered to prevent MTU issues from arising.

Related to fragmentation is the question of how best to handle Path MTU issues, should they occur. Ideally, the original source of any packet (i.e, the sending VM) would be notified of the optimal MTU to use. Path MTU problems occurring within an overlay network would result in ICMP MTU exceeded messages being sent back to the egress tunnel switch at the entry point of the overlay. If the switch is embedded within a hypervisor, the hypervisor could notify the VM of a more appropriate MTU to use. It may be appropriate to specify a set of best practices for implementers related to the handling of Path MTU issues.

[4.2.3.](#) Checksums and FCS

When tunneling packets, both the inner and outer headers could have their own checksum, duplicating effort and impacting performance. Therefore, we strongly recommend that any solution carry only one set

of checksum or frame FCS.

When the inner packet is TCP or UDP, they already include their own checksum, and adding a second outer checksum (using the same 1's complement algorithm) provides little value. Similarly, if the inner packet is an Ethernet frame, the frame FCS protects the original frame and a new frame FCS over both the original frame and the overlay header protects the new encapsulated frame.

In IPv4, UDP checksums can be disabled on a per-packet basis simply by setting the checksum field to zero. IPv6, however, specifies that UDP checksums must always be included. But even for IPv6, the LISP protocol [I-D.ietf-lisp] already allows a zero checksum field. The 6man working group is also currently considering relaxing the IPv6 UDP checksum requirement [[I-D.ietf-6man-udpzero](#)].

For Ethernet frames, L2 overlays such as TRILL already mandate only a single frame FCS.

[4.2.4.](#) Middlebox Traversal

One issue to consider is to whether the overlay will need to run over networks that include middleboxes such as NAT. Middleboxes may have difficulty properly supporting multicast or other aspects of an overlay header. Inside a data center, it may well be the case that middlebox traversal is a non-issue. But if overlays are extended across the broader Internet, the presence of middleboxes may be of concern.

[4.2.5.](#) OAM

Successful deployment of an overlay approach will likely require appropriate Operations, Administration and Maintenance (OAM) facilities.

[5.](#) Control Plane

The control plane needs to address the following pieces, at least:

1. A mechanism to populate the forwarding table of a virtual network

instance.

2. A mechanism to handle multi-destination frames within a virtual network instance.
3. A mechanism to allow an endpoint to inform the access switch which virtual network instance it wishes to join on a virtual network interface.
4. A mechanism to allow an endpoint to inform the access switch about its leaving the network so that the access switch can clean up state.

[5.1.](#) Populating the Forwarding Table of a Virtual Network Instance

When an access switch has to forward a frame from one endpoint to another, across the network, it has to consult some form of a forwarding table. When we use network overlays, the problem boils down to deriving the mapping between the inner and outer addresses i.e. deriving the destination address in the overlay header based on the destination address sent by the endpoint. Two well known mechanisms for populating the forwarding table (or deriving the mapping table) of a switch are (i) via a routing control protocol and

(ii) learning from the data plane as Ethernet bridges do. Another mechanism is through a centralized mapping database. Any solution must avoid problems associated with scaling a virtual network instance across a large data center.

[5.2.](#) Handling Multi-destination Frames

Another aspect of address mapping concerns the handling of multi-destination frames, i.e. broadcast and multicast frames, or the delivery of unicast packets when no mapping exists. Associating a infrastructure multicast address is one possible way of connecting together all the machines belonging to the same VNID. However, existing multicast implementations do not scale to efficiently handle hundreds of thousands of multicast groups, as would be required if one multicast group were assigned to each VNID.

[5.3.](#) Associating a VNID With An Endpoint

When an endpoint, such as VM or physical server, connects to the infrastructure, we must define a mechanism to allow the endpoint to identify to the access switch the network instance that it wishes to join. Typically, it is a virtual NIC (the one connected to the VM) coming up that triggers this association. The access switch can then determine the VNID to be associated with this virtual NIC. A standard protocol that all types of overlay encapsulation points can use to identify the VNID associated with an endpoint will be beneficial for supporting multi-vendor implementations. This protocol could also be used to distribute any per virtual network information (e.g. a multicast group address). This signaling can provide the stimulus to trigger the overlay termination points to perform any actions needed within the infrastructure network (e.g. use IGMP to join a multicast group).

5.4. Disassociating a VNID on Termination or Move

To enable cleaning up state in the access switch, we must define a mechanism to allow an endpoint to signal its disconnection from the network.

6. Related Work

6.1. ARMD

ARMD is chartered to look at data center scaling issues with a focus on address resolution. ARMD is currently chartered to develop a problem statement and is not currently developing solutions. While an overlay-based approach may address some of the "pain points" that

have been raised in ARMD (e.g., better support for multi-tenancy), an overlay approach may also push some of the L2 scaling concerns (e.g., excessive flooding) to the IP level (flooding via IP multicast). Analysis will be needed to understand the scaling trade offs of an overlay based approach compared with existing approaches. On the other hand, existing IP-based approaches such as proxy ARP may help mitigate some concerns.

6.2. TRILL

TRILL is an L2 based approach aimed at improving deficiencies and

limitations with current Ethernet networks. Approaches to extend TRILL to support more than 4094 VLANs are currently under investigation [[I-D.eastlake-trill-rbridge-fine-labeling](#)]

[6.3.](#) L2VPNs

The IETF has specified a number of approaches for connecting L2 domains together as part of the L2VPN Working Group. That group, however has historically been focused on Provider-provisioned L2 VPNs, where the service provider participates in management and provisioning of the VPN. In addition, much of the target environment for such deployments involves carrying L2 traffic over WANs. Overlay approaches are intended be used within data centers where the overlay network is managed by the data center operator, rather than by an outside party. While overlays can run across the Internet as well, they will extend well into the data center itself (e.g., up to and including hypervisors) and include large numbers of machines within the data center itself.

Other L2VPN approaches, such as L2TP [[RFC2661](#)] require significant tunnel state at the encapsulating and decapsulating end points. Overlays require less tunnel state than other approaches, which is important to allow overlays to scale to hundreds of thousands of end points. It is assumed that smaller switches (i.e., virtual switches in hypervisors or the physical switches to which VMs connect) will be part of the overlay network and be responsible for encapsulating and decapsulating packets.

[6.4.](#) Proxy Mobile IP

Proxy Mobile IP [[RFC5213](#)] [[RFC5844](#)] makes use of the GRE Key Field [[RFC5845](#)] [[RFC6245](#)], but not in a way that supports multi-tenancy.

[6.5.](#) LISP

LISP[I-D.ietf-lisp] essentially provides an IP over IP overlay where the internal addresses are end station Identifiers and the outer IP

addresses represent the location of the end station within the core IP network topology. The LISP overlay header uses a 24 bit Instance ID used to support overlapping inner IP addresses.

6.6. Individual Submissions

Many individual submissions also look to addressing some or all of the issues addressed in this draft. Examples of such drafts are VXLAN [[I-D.mahalingam-dutt-dcops-vxlan](#)], NVGRE [[I-D.sridharan-virtualization-nvgre](#)] and Virtual Machine Mobility in L3 networks [[I-D.wkumari-dcops-l3-vmmobility](#)].

7. Further Work

It is believed that overlay-based approaches may be able to reduce the overall amount of flooding and other multicast and broadcast related traffic (e.g, ARP and ND) currently experienced within current data centers with a large flat L2 network. Further analysis is needed to characterize expected improvements.

8. Summary

This document has argued that network virtualization using L3 overlays addresses a number of issues being faced as data centers scale in size. In addition, careful consideration of a number of issues would lead to the development of interoperable implementation of virtualization overlays.

9. Acknowledgments

Helpful comments and improvements to this document have come from Ariel Hendel, Vinit Jain, and Benson Schliesser.

10. IANA Considerations

This memo includes no request to IANA.

11. Security Considerations

TBD

12. Informative References

[I-D.eastlake-trill-rbridge-fine-labeling]

Eastlake, D., Zhang, M., Agarwal, P., Dutt, D., and R. Perlman, "RBridges: Fine-Grained Labeling", [draft-eastlake-trill-rbridge-fine-labeling-02](#) (work in progress), October 2011.

[I-D.hasmit-otv]

Grover, H., Rao, D., Farinacci, D., and V. Moreno, "Overlay Transport Virtualization", [draft-hasmit-otv-03](#) (work in progress), July 2011.

[I-D.ietf-6man-udpzero]

Fairhurst, G. and M. Westerlund, "IPv6 UDP Checksum Considerations", [draft-ietf-6man-udpzero-04](#) (work in progress), October 2011.

[I-D.ietf-lisp]

Farinacci, D., Fuller, V., Meyer, D., and D. Lewis, "Locator/ID Separation Protocol (LISP)", [draft-ietf-lisp-15](#) (work in progress), July 2011.

[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-00](#) (work in progress), August 2011.

[I-D.sridharan-virtualization-nvgre]

Sridharan, M., Duda, K., Ganga, I., Greenberg, A., Lin, G., Pearson, M., Thaler, P., Tumuluri, C., Venkataramaiah, N., and Y. Wang, "NVGRE: Network Virtualization using Generic Routing Encapsulation", [draft-sridharan-virtualization-nvgre-00](#) (work in progress), September 2011.

[I-D.wkumari-dcops-l3-vm-mobility]

Kumari, W. and J. Halpern, "Virtual Machine mobility in L3 Networks.", [draft-wkumari-dcops-l3-vm-mobility-00](#) (work in progress), August 2011.

[RFC2661]

Townsley, W., Valencia, A., Rubens, A., Pall, G., Zorn, G., and B. Palter, "Layer Two Tunneling Protocol "L2TP"", [RFC 2661](#), August 1999.

Internet-Draft

Overlays for Network Virtualization

October 2011

Traina, "Generic Routing Encapsulation (GRE)", [RFC 2784](#), March 2000.

[RFC2890] Dommety, G., "Key and Sequence Number Extensions to GRE", [RFC 2890](#), September 2000.

[RFC5213] Gundavelli, S., Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, "Proxy Mobile IPv6", [RFC 5213](#), August 2008.

[RFC5844] Wakikawa, R. and S. Gundavelli, "IPv4 Support for Proxy Mobile IPv6", [RFC 5844](#), May 2010.

[RFC5845] Muhanna, A., Khalil, M., Gundavelli, S., and K. Leung, "Generic Routing Encapsulation (GRE) Key Option for Proxy Mobile IPv6", [RFC 5845](#), June 2010.

[RFC6245] Yegani, P., Leung, K., Lior, A., Chowdhury, K., and J. Navali, "Generic Routing Encapsulation (GRE) Key Extension for Mobile IPv4", [RFC 6245](#), May 2011.

[RFC6325] Perlman, R., Eastlake, D., Dutt, D., Gai, S., and A. Ghanwani, "Routing Bridges (RBridges): Base Protocol Specification", [RFC 6325](#), July 2011.

Authors' Addresses

Thomas Narten (editor)
IBM

Email: narten@us.ibm.com

Murari Sridharan
Microsoft

Email: muraris@microsoft.com

Dinesh Dutt

Cisco

Email: ddutt@cisco.com

Narten, et al.

Expires May 3, 2012

[Page 17]

Internet-Draft

Overlays for Network Virtualization

October 2011

David Black
EMC

Email: david.black@emc.com

Lawrence Kreeger
Cisco

Email: kreeger@cisco.com

