

Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: September 15, 2013

D. Rao
J. Mullooly
R. Fernando
Cisco
March 14, 2013

Layer-3 virtual network overlays based on BGP Layer-3 VPNs
draft-drao-bgp-l3vpn-virtual-network-overlays-01

Abstract

Virtual network overlays are being designed and deployed in various types of networks, including data centers. These network overlays address several requirements including flexible network virtualization and multi-tenancy, increased scale, and support for mobility of virtual machines. Such overlay networks can be used to provide both Layer-2 and Layer-3 network services to hosts at the network edge. New packet encapsulations are being defined and standardized to support these virtual networks. These encapsulations, such as VXLAN and NVGRE, are primarily based on IP and are currently defined to support a Layer-2 forwarding service.

BGP based Layer-3 VPNs, as specified in [RFC 4364](#), provide an industry proven and well-defined solution for supporting Layer-3 virtual network services. However, [RFC 4364](#) procedures use MPLS labels to provide the network virtualization capability in the data plane. With the increasing support for IP overlay encapsulations in data center devices, there is a strong preference to utilize this support to deploy Layer-3 virtual networks using the familiar policy and operational primitives of Layer-3 VPNs.

This document describes the use of BGP Layer-3 VPNs along with the new IP-based virtual network overlay encapsulations to provide a Layer-3 virtualization solution for all IP traffic, and specifies mechanisms to use the new encapsulations while continuing to leverage the BGP Layer-3 VPN control plane techniques and extensions. This mechanism provides an efficient incremental solution to support forwarding for IP traffic, irrespective of whether it is destined within or across an IP subnet boundary.

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 September 15, 2013.

Copyright Notice

Copyright (c) 2013 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	4
1.1.	Requirements Language	5
1.2.	Terminology	5
1.3.	Control plane signaling requirements	5
1.4.	Control plane model	5
1.5.	Overlay Encapsulations	6
2.	Virtual Network Identifier	6
2.1.	Virtual Network Identifier Scope	7
2.1.1.	Domain-scoped provisioned virtual network identifiers	7
2.1.2.	Per-device scoped allocated virtual network identifiers	7
2.1.3.	Global unicast table	8
2.1.4.	Virtual Network Identifier Specification	8
2.2.	Signaling Virtual Network Identifiers	8
2.2.1.	Signaling Requirements	8
2.2.2.	Signaling Specification	9
3.	Overlay Encapsulation specification	10
3.1.	Encapsulation for VXLAN and NVGRE	11
3.2.	Encapsulation for MPLS-in-GRE	11
3.3.	Multiple encapsulations	11
3.4.	Gateway device encapsulation handling	11
4.	Forwarding behavior	12
5.	Overlay Interconnection and Interworking Scenarios	12
5.1.	End-to-end overlay	12
5.2.	Virtual-network overlay VPN interworking	13
5.2.1.	Normalized interworking via VRF	13
5.2.2.	Seamless VPN interworking	13
6.	Virtual-Network Overlay Encapsulation Capability	13
7.	Acknowledgements	14
8.	IANA Considerations	14
9.	Security Considerations	14
10.	References	15
10.1.	Normative References	15
10.2.	Informative References	15
	Authors' Addresses	16

1. Introduction

Virtual network overlays are increasingly being designed and deployed in various types of networks, including data center networks. These virtual network overlays can be used to provide both Layer-2 and Layer-3 network services to hosts at the network edge. New encapsulations are being defined and standardized to support these virtual networks. These encapsulations are primarily based on IP transport, such as VXLAN and NVGRE. A significant characteristic of these encapsulations is the presence of an embedded virtual network identifier field that is part of the encapsulation header. The use of these encapsulations is defined in [VXLAN] and [NVGRE] and is being currently worked on as part of the NV03 architecture [NV03].

BGP based Layer-3 VPNs, as specified in [RFC 4364](#), provide an industry proven and well-defined solution for supporting Layer-3 virtual network services. The Layer-3 VPN BGP control plane is eminently suitable to provide a Layer-3 network virtualization solution in the data center.

However, [RFC 4364](#) mechanisms use MPLS labels as the mechanism to provide the network virtualization capability in the data plane. An MPLS label is signaled by a device advertising a VPN-IP route. This label can identify the virtual network when the device processes packets received with that label. [RFC 4364](#) does allow an MPLS label to be carried in an IP transport encapsulation such as MPLS-in-GRE.

This document specifies procedures to use the new IP-based virtual network overlay encapsulations such as VXLAN and NVGRE, while continuing to leverage the BGP based Layer-3 VPN control plane techniques and extensions. It also describes how virtual network overlays based on these encapsulations can efficiently interconnect with one another and with existing MPLS based L3VPN networks.

This document describes the protocol extensions necessary to allow advertising a VPN-IP NLRI with an attached VN-ID as well as an encapsulation attribute indicating the type of encapsulation, for example, VXLAN or NVGRE.

There are aspects of the signaling of encapsulation and VN-ID that can be leveraged across different kinds of services. Hence, the generic overlay encapsulation signaling extensions are defined in [[remote-next-hop](#)]. The current document provides the necessary context of how these extensions are used with the BGP IP-VPN NRIs.

1.1. 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](#)].

1.2. Terminology

VM: Virtual Machine

Edge device: The edge device is where end-hosts (eg. application VMs) attach to the overlay. This is where the tunnel encapsulation starts. It's called NVE in the NV03 terminology. The NVE is equivalent to a VPN PE in the context of BGP L3VPNs.

1.3. Control plane signaling requirements

While considering the leverage of the BGP L3VPN control plane with the IP overlay technologies, the following requirements should be supported.

1. Signal VN-ID with VPN-IP routes, that can be used with IP based overlay encapsulations.
2. Support signaling of multiple encapsulations per edge device.
3. Have flexibility to support single and per-encapsulation VN-ID spaces if needed.
4. Support both device-local and domain-global VN-ID/label spaces.
5. Support per-prefix granularity for VN-ID/encapsulation.
6. Support interoperability with legacy IP-VPN PEs.
7. Be efficient in signaling to provide good scalability.
8. Minimize protocol and deployment overhead.

1.4. Control plane model

The virtual network overlay described in this document uses regular BGP peering on the edge devices for policy constrained route distribution. A typical deployment would use Route Reflectors.

It is also feasible for an alternative protocol or provisioning framework to be used to control the forwarding plane population and forwarding behavior on the edge devices, as described in [vpe-

framework].

The extensions specified here are compatible with both approaches.

1.5. Overlay Encapsulations

Different tunnel encapsulations may be used to realize an overlay virtual network. Based on the encapsulation type being used, the virtual network identifier is appropriately encoded in the encapsulation header.

An overlay network may use the IP based VN-ID encapsulations such as VXLAN and NVGRE. It may also use an MPLS based encapsulation such as MPLS-in-GRE.

When VXLAN encapsulation is used, the virtual network identifier is carried as the 24-bit VNI in the VXLAN header.

When NVGRE encapsulation is used, the virtual network identifier is carried as the 24-bit VSID in the NVGRE header.

When MPLS-in-GRE is used, the regular MPLS VPN label serves as the data plane identifier for the virtual network or a specific destination.

A given overlay edge device may support a single encapsulation type or it may support multiple encapsulation types. In the latter case, it may signal the multiple encapsulations so that a receiving device can potentially use the one most suitable to it. An edge device may use the same encapsulation(s) for all routes or for a subset of routes.

2. Virtual Network Identifier

In [RFC 4364](#) based Layer-3 VPNs, a 20-bit MPLS label is assigned to an VPN-IP route by the device that advertises the route, with itself as the BGP next-hop. This label determines the forwarding behavior in the data plane for traffic being switched as per that route. A device receiving this route will encode this label in the packet header when sending traffic to the advertiser. The advertiser will take a unique forwarding action for traffic received with this label when compared to traffic with other labels. The label may also be used at the granularity of a VPN table and drive an IP lookup in that table. This MPLS VPN label is independent of the transport encapsulation that is used to carry this traffic to this PE from other PEs across a core network. The transport encapsulation may be native MPLS or be IP (eg, MPLS-in-GRE).

On the other hand, the various IP overlay encapsulations support a virtual network identifier explicitly within their encapsulation header. A virtual network identifier is a value that at a minimum can identify a specific virtual network table in the forwarding plane, and may be used to perform an IP address lookup. It may also drive a specific forwarding action for packets destined to a particular destination address or prefix.

It is typically a 24-bit value which can support upto 16 million individual network segments or end-hosts. For instance, VXLAN defines a 24-bit VNI while NVGRE uses a 24-bit VSID that is carried in the GRE key field of the GRE header.

2.1. Virtual Network Identifier Scope

The scope of these virtual network identifiers fall into two broad categories. It is important to support both cases, and in doing so, ensure that the scope of the identifier be clear and the values not conflict with each other.

2.1.1. Domain-scoped provisioned virtual network identifiers

Based on the provisioning mechanism used, a virtual network identifier typically has a domain-wide scope within the network domain, where a unique value is assigned to a given virtual network or a given IP destination route at one or more edge devices.

This scope is useful in environments such as data centers where virtual networks and VMs are automatically provisioned by central orchestration systems. The system must support a very large number of VN-IDs given the scope.

2.1.2. Per-device scoped allocated virtual network identifiers

There are scenarios where it is also necessary for an identifier to have significance local to each network edge device that advertises the route. In this case, the same value may be used by different edge devices to represent different forwarding classes.

When it is locally scoped, the virtual network identifier may be dynamically allocated by the advertising device. This allocation follows the same semantics of an MPLS VPN label, and supports similar forwarding behaviors as specified in [RFC 4364](#). The device may, for example, be a DC-WAN edge device that supports L3VPN Inter-AS Option B and use this allocation for routes received from other ASBRs.

2.1.3. Global unicast table

The overlay encapsulation can also be used to support forwarding for routes in the global or default routing table. A virtual network identifier value can be allocated for the purpose as per the above options.

2.1.4. Virtual Network Identifier Specification

The above requirements can be achieved in a simple manner by splitting the virtual network ID number space to support both domain-wide and device-local scopes.

- o Values upto 1 million (or less than 20 bits) SHOULD be treated with the same semantics as MPLS VPN labels and have significance local to the advertiser.

For future expansion, this draft stipulates that the 16 numerical values in the end of the VN-ID range be reserved for future use. These special values could be used to indicate the presence of other types of IP payloads.

- o Values greater than 1 million (greater than 20 bits) SHOULD be treated as per their original definition, ie domain-wide scoped values.

These limits are not mandatory, but are recommended defaults. As long as the provisioning system can ensure conflict-free operation, the boundary between local and domain scoped ranges can be adjusted higher or lower by configuration.

- o A virtual network identifier value of zero SHOULD be used by default to indicate the global or routing table.

2.2. Signaling Virtual Network Identifiers

2.2.1. Signaling Requirements

The Introduction section listed the desirable characteristics of signaling of VN-IDs. This section elaborates on a couple of those requirements.

- o The device may support a single VN-ID space across all its supported encapsulations.

This is expected to be common deployed scenario. A given edge device will be provisioned by a single network orchestrator or controller. The device may support multiple encapsulations in order to

interoperate with remote edge devices that support a different encapsulation. However, the single network orchestrator will manage the VN-ID space that will be common across multiple encapsulations on this device.

- o A device may support an independent VN-ID space per-supported encapsulation.

This is expected to be applicable mostly at network gateway devices that interconnect two different overlay domains and support the same virtual network across these domains. These border devices are likely to be managed by two different orchestrators, and hence need to support different VN-ID spaces. In this case, they typically advertise routes of one domain into another.

- o An edge device may support an independent VN-ID space per-supported encapsulation.

This is assumed to not be a common scenario, where an edge device within the domain is being shared or managed by multiple orchestration systems. However, in case this scenario must be supported, the edge device must be able to support multiple distinct VN-ID spaces. An alternative scheme would be to divide the VN-ID range among the orchestration systems.

- o It is required to support prefix-level VN-ID assignment.

Supporting prefix-level granularity is useful in various scenarios, for example, at an interworking point between DC (VXLAN) and WAN (MPLS).

2.2.2. Signaling Specification

This document specifies two options for signaling VN-IDs.

1. The VN-ID is encoded within an Virtual-Network Overlay encapsulation attribute that also contains the encapsulation type and associated parameters.

This enables the device to signal a VN-ID per encapsulation that it may support. For example, the device may use VN-ID1 for VXLAN and VN-ID2 for NVGRE. The VN-ID is encoded as a 24-bit value in each encapsulation attribute.

When multiple VN-IDs need to be signaled, one per overlay encapsulation type, then the VN-ID MUST be included in the overlay encapsulation attribute as defined in [[remote-next-hop](#)].

When MPLS-in-GRE is one of the encapsulations, there is no change from current behavior. The VPN label is encoded in the label field in the IP-VPN NLRI.

2. The VN-ID is encoded in the label field in the IP-VPN NLRI.

This option is used when a device supports a single VN-ID space across all encapsulations. The benefit of this encoding is it's efficiency of packing, even when used for per-prefix VN-ID assignment. With this option, the 24-bit VN-ID for VXLAN and NVGRE is encoded as a 3-byte label field in the IP-VPN NLRI.

When a VN-ID or VPN label is to be signaled, the value MUST be encoded in the 3-octet label field in the IP or IP-VPN NLRI.

This offers the most efficient packing of prefixes in BGP update messages. The device may still advertise multiple encapsulation types with this route, but they will all use the same VN-ID value.

An advertising device will select the suitable option as per the requirements stated above, based on configuration.

3. Overlay Encapsulation specification

Signaling the VN-ID must be coupled with signaling the appropriate overlay encapsulation type. An overlay encapsulation attribute MUST be carried with each route.

The section above specified two options of signaling VN-ID. In both options, the accompanying encapsulation attribute indicates that a 24-bit VN-ID is specified with the NLRI and must be encoded in the signaled encapsulation header.

The encapsulation attribute also indicates the accompanying parameters to be used in the packet header.

[RFC 5512](#) defines a Tunnel Encapsulation Extended Community that can be used to signal different tunnel types. It defines an Encapsulation Sub-TLV that can be used to specify encapsulation parameters.

[remote-next-hop] specifies a Remote-Next-Hop attribute which reuses the Encapsulation Sub-TLV from [RFC 5512](#), but adds the flexibility to signal the encapsulation attribute and parameters along with each individual route. The address specified as the remote next-hop identifies the end-point or destination of the encapsulated packets that use the dependent routes as well as the tunnel encapsulation

parameters.

Hence, the Remote-Next-Hop attribute is used to signal VN-ID encapsulations. New tunnel types are defined for VXLAN and NVGRE. The format for the tunnel parameters are specified in [remote-next-hop].

3.1. Encapsulation for VXLAN and NVGRE

When VXLAN and NVGRE encapsulations are used, the header by definition contains an Ethernet MAC address within the overlay header. When these encapsulations are used for Layer-3 as specified in this document, the MAC addresses are not relevant. A single well-known MAC address may be specified for the purpose of deterministically driving a Layer-3 lookup based on the inner IP or IPv6 address.

Alternatively, an overlay egress edge device may choose to specify a MAC address as part of the encapsulation header in its route advertisement. In this case, any ingress edge device sending traffic as per this route MUST use the above specified MAC address as the destination MAC address in the header. The egress device may use this address to drive the Layer-3 table lookup or for other purposes.

3.2. Encapsulation for MPLS-in-GRE

When MPLS-in-GRE is one of the encapsulations, there is no change from current behavior. A tunnel type of [MPLS-in-GRE] as defined in [RFC 5512](#) is used in the Remote-Next-Hop attribute.

3.3. Multiple encapsulations

A given overlay edge device MAY advertise multiple Encapsulation Sub-TLVs, in order to signal multiple encapsulations. It MAY encode a different VN-ID in each Sub-TLV as per rules specified earlier.

A receiving edge device may support one or more encapsulations that are signaled by the advertising edge device. In that case, the receiving device can select any of these encapsulations for sending traffic to the advertiser. If a receiving device supports no encapsulations that were signaled by the advertiser, then it will not send any traffic for these routes to the advertiser.

3.4. Gateway device encapsulation handling

When an intermediate gateway device changes the BGP next-hop to self before propagating a received route, it may need to advertise a new overlay encapsulation attribute with the local address as the

endpoint. While doing so, it MAY use an overlay encapsulation type that is different from the received route. It MAY also signal a different VN-ID or VPN label than what it received, as described in the various VN-ID requirements and rules earlier.

4. Forwarding behavior

- o Locally assigned virtual network identifiers

When the virtual network identifier is locally assigned, forwarding based on the identifier at the advertising device follows the semantics of an MPLS VPN label. The VN-ID may either drive an IP table lookup or provide a seamless binding to an output VN-ID or label.

- o Domain-scoped provisioned virtual network identifiers

With a provisioned VN-ID, forwarding behavior at a device which is provisioned with this value is governed by the forwarding action that has been provisioned. As one example, the VN-ID may be set up to represent a specific IP VRF table on all relevant edge devices, causing incoming packets with this VN-ID to undergo an IP lookup. Alternatively, the VN-ID may be configured on only one or two edge or border devices to directly forward incoming packets to an attached end-host, without undergoing an IP lookup.

In either case, the forwarding behavior at any ingress edge device (physical or virtual) remains the same. The ingress edge device adds an encapsulation as signaled by the advertising device, and includes the VN-ID in that encapsulation header.

5. Overlay Interconnection and Interworking Scenarios

Multiple virtual network overlay domains may be inter-connected using a couple of approaches. Both these approaches may co-exist in the same data center, and be used for connectivity to different kinds of external networks.

5.1. End-to-end overlay

The IP overlay encapsulation or tunnel extends end-to-end between edge devices in different data centers.

IP routes for hosts attached to each edge device are exchanged between the overlay domains either via route exchange between BGP speakers in each overlay domain, or via an orchestration framework

that manages the domains. The two networks may be located within the same ASN or may extend across ASes.

The routes are propagated to various edge device via the control plane mechanism used in the DC, along with the encapsulation and VN-ID or label to be used for sending traffic to a given destination edge device. All intermediate devices in the forwarding path between the two edge devices simply transport the IP encapsulated overlay traffic.

The tunnel endpoints, ie the edge devices need to be reachable across the ASes. This reachability may be provided by BGP peering.

5.2. Virtual-network overlay VPN interworking

The overlay encapsulation terminates at a border router such as the DC-WAN gateway device. The gateway device may re-encapsulate packets in another header when sending it onwards. This requires an interworking function which can be of multiple types.

5.2.1. Normalized interworking via VRF

The overlay based virtual network terminates into an L3VPN VRF at the DC-WAN gateway device. Internal routes of the DC as well as the external routes received from the WAN router are installed in the VRF forwarding table at the DC gateway router. The DC gateway will perform an IP lookup and forward traffic after doing the appropriate output MPLS or IP overlay/VPN encapsulation.

5.2.2. Seamless VPN interworking

In this case, the DC Gateway router performs a direct translation of VN-IDs/labels while switching packets between the DC and WAN interfaces without doing an IP lookup. The forwarding table at the DC Gateway router is set up to do a VN-ID or VPN label lookup and derive the output VN-ID or VPN label. The DC Gateway Router can act as an Inter-AS Option-B ASBR/ABR peering with other ASBRs/ABRs.

6. Virtual-Network Overlay Encapsulation Capability

Need for Capability Negotiation:

When the IP-VPN NLRIs are used along with a VN-ID based encapsulation, the label field will either represent a VN-ID or not be used. A device that supports vanilla IP-VPNs but does not understand the extensions specified in this document may not interpret the received IP-VPN NLRI as intended, leading to incorrect

forwarding plane behavior. The presence of a Capability protect against this issue and ensures interoperability with vanilla IP-VPN peers.

With the Virtual-Network Overlay Encapsulation Capability, a VN-capable BGP speaker will detect peers that are not capable of processing VN-ID encapsulation information received in BGP updates. The speaker **MUST** not send any VPN-IP routes that contain only a VN-ID based encapsulation to such peers. When the route contains both a VN-ID encapsulation and an MPLS-in-IP encapsulation, the speaker **MAY** send the route to the legacy peer with only the MPLS encapsulation information.

The VN-ID encapsulation capability is described in [[remote-next-hop](#)].

7. Acknowledgements

The authors would like to acknowledge and thank Nabil Bitar, Dave Smith, Maria Napierala, Robert Raszuk, Eric Rosen, Ashok Ganesan and Luyuan Fang for their input and feedback.

8. IANA Considerations

There are no requirements to IANA at this time. A specific VN-ID range may be reserved for future use as applications for carrying payloads different than regular IP/VPN packets emerge in future.

9. Security Considerations

This draft does not add any additional security implications to the BGP/L3VPN control plane. All existing authentication and security mechanisms for BGP apply here.

There are security considerations pertaining to IP based overlay or tunnel encapsulations which are described in the respective overlay encapsulation specifications as well as in [RFC 5512](#).

There are certain measures that may be taken by default to increase the level of security provided at the overlay edge devices.

When an IP-based overlay encapsulation is used within a domain such as a data center, the network edge devices can enforce a default forwarding access rule to restrict the acceptance of such overlay encapsulated packets on their access interfaces attached to servers or VMs.

For example, when VXLAN is being used, an edge device may be directed to filter any VXLAN encapsulated packets (identified by the UDP port number) on their access interfaces. This rule can be further augmented by checking that the destination IP address of such VXLAN packets does not fall in the prefix range allocated to edge device addresses. Similarly, a DC edge device may be directed to not accept any VXLAN encapsulated packets on its interfaces connected to external routers, depending on the interconnectivity option being used.

10. References

10.1. Normative References

- [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-03](#) (work in progress), February 2013.
- [I-D.sridharan-virtualization-nvgre]
Sridharan, M., Greenberg, A., Venkataramaiah, N., Wang, Y., Duda, K., Ganga, I., Lin, G., Pearson, M., Thaler, P., and C. Tumuluri, "NVGRE: Network Virtualization using Generic Routing Encapsulation", [draft-sridharan-virtualization-nvgre-02](#) (work in progress), February 2013.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [min_ref] authSurName, authInitials., "Minimal Reference", 2006.

10.2. Informative References

- [I-D.fang-l3vpn-virtual-pe]
Fang, L., Ward, D., Fernando, R., Napierala, M., Bitar, N., Rao, D., Rijsman, B., and S. Ning, "BGP IP VPN Virtual PE", [draft-fang-l3vpn-virtual-pe-00](#) (work in progress), February 2013.
- [I-D.narten-iana-considerations-rfc2434bis]
Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [draft-narten-iana-considerations-rfc2434bis-09](#) (work in progress), March 2008.

[I-D.vandeveldde-idr-remote-next-hop]

Velde, G., Patel, K., Rao, D., Raszuk, R., and R. Bush,
"BGP Remote-Next-Hop",

[draft-vandeveldde-idr-remote-next-hop-03](#) (work in
progress), October 2012.

[RFC2629] Rose, M., "Writing I-Ds and RFCs using XML", [RFC 2629](#),
June 1999.

[RFC3107] Rekhter, Y. and E. Rosen, "Carrying Label Information in
BGP-4", [RFC 3107](#), May 2001.

[RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC
Text on Security Considerations", [BCP 72](#), [RFC 3552](#),
July 2003.

[RFC4271] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway
Protocol 4 (BGP-4)", [RFC 4271](#), January 2006.

[RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private
Networks (VPNs)", [RFC 4364](#), February 2006.

[RFC5512] Mohapatra, P. and E. Rosen, "The BGP Encapsulation
Subsequent Address Family Identifier (SAFI) and the BGP
Tunnel Encapsulation Attribute", [RFC 5512](#), April 2009.

Authors' Addresses

Dhananjaya Rao
Cisco
San Jose,
USA

Email: dhrao@cisco.com

John Mullooly
Cisco
New Jersey,
USA

Email: jmullool@cisco.com

Rex Fernando
Cisco
San Jose,
USA

Email: rex@cisco.com