Network Working Group
Request for Comments: 1364

K. Varadhan
OARnet

September 1992

BGP OSPF Interaction

Status of this Memo

Abstract

This memo defines the various criteria to be used when designing Autonomous System Border Routers (ASBR) that will run BGP with other ASBRs external to the AS and OSPF as its IGP.

Table of Contents

1. Introduction

This document defines the various criteria to be used when designing Autonomous System Border Routers (ASBR) that will run BGP [RFC1267] with other ASBRs external to the AS, and OSPF [RFC1247] as its IGP.

This document defines how the following fields in OSPF and attributes in BGP are to be set when interfacing between BGP and OSPF at an ASBR:

```
OSPF cost and type      vs. BGP INTER-AS METRIC
OSPF tag                vs. BGP ORIGIN and AS_PATH
OSPF Forwarding Address vs. BGP NEXT_HOP
```

For a more general treatise on routing and route exchange problems, please refer to [ROUTE-LEAKING] and [NEXT-HOP] by Philip Almquist.

This document uses the two terms "Autonomous System" and "Routing Domain". The definitions for the two are below:

The term Autonomous System is the same as is used in the BGP-3 RFC [RFC1267], given below:

> "The use of the term Autonomous System here stresses the fact that, even when multiple IGPs and metrics are used, the administration of an AS appears to other ASs to have a single coherent interior routing plan and presents a consistent picture of what networks are reachable through it. From the standpoint of exterior routing, an AS can be viewed as monolithic: reachability to networks directly connected to the AS must be equivalent from all border gateways of the AS."

The term Routing Domain was first used in [ROUTE-LEAKING] and is given below:

> "A Routing Domain is a collection of routers which coordinate their routing knowledge using a single (instance of) a routing protocol."

2. Route Exchange

This section discusses the constraints that must be met to exchange routes between an external BGP session with a peer from another AS and internal OSPF routes.

BGP does not carry subnet information in routing updates. Therefore, when referring to a subnetted network in the OSPF routing domain, we consider the equivalent network route in the context of BGP.

Multiple subnet routes for a subnetted network in OSPF are collapsed into one network route when exported into BGP.

2.1. Exporting OSPF routes into BGP

1. The administrator must be able to selectively export routes into BGP via an appropriate filter mechanism.

   This filter mechanism must support such control with the granularity of a single network.

   Additionally, the administrator must be able to filter based on the OSPF tag and the various sub-fields of the OSPF tag. The settings of the tag and the sub-fields are defined in section 4 in more detail.

   - By default, no routes must be exported from OSPF into BGP. A single mechanism must permit all OSPF inter-area and intra-area routes to be exported into BGP.

   OSPF external routes of type 1 and type 2 must never be exported into BGP unless they are explicitly configured.

2. When configured to export a network, the ASBR must advertise a network route for a subnetted network, as long as at least one subnet in the subnetted network is reachable via OSPF.

3. The network administrator must be able to statically configure the BGP attribute INTER-AS METRIC to be used for any network route.

   - By default, the INTER_AS METRIC must default to 1.

   Explanatory text: The OSPF cost and the BGP INTER-AS METRIC are of different widths. The OSPF cost is a two level metric. The BGP INTER-AS METRIC is only an optional non-transitive attribute. Hence, a more complex BGP INTER-AS METRIC-OSPF cost mapping scheme is not necessary.

4. When an ASBR is advertising an OSPF route to network Y to external BGP neighbours and learns that the route has become unreachable, the ASBR must immediately propogate this information to the external BGP neighbours.

5. An implementation of BGP and OSPF on an ASBR must have a mechanism to set up a minimum amount of time that must elapse between the learning of a new route via OSPF and subsequent advertisement of the route via BGP to the external

neighbours.

- The default value for this setting must be 0, indicating that the route is to be advertised to the neighbour BGP peers instantly.

  Note that [RFC1267] mandates a mechanism to dampen the inbound advertisements from adjacent neighbours.

2.2. Importing BGP routes into OSPF

1. BGP implementations should allow an AS to control announcements of BGP-learned routes into OSPF. Implementations should support such control with the granularity of a single network. Implementations should also support such control with the granularity of an autonomous system, where the autonomous system may be either the autonomous system that originated the route or the autonomous system that advertised the route to the local system (adjacent autonomous system).

   - By default, no routes must be imported from BGP into OSPF. Administrators must configure every route they wish to import.

     A mechanism may allow an administrator to configure an ASBR to import all the BGP routes into the OSPF routing domain.

2. The administrator must be able to configure the OSPF cost and the OSPF metric type of every route imported into OSPF.

   - The OSPF cost must default to 1; the OSPF metric type must default to type 2.

3. Routes learned via IBGP must not be imported into OSPF.

4. The ASBR must never generate a default route into the OSPF routing domain unless explicitly configured to do so.

   A possible criterion for generating default into an IGP is to allow the administrator to specify a set of (network route, AS_PATH, default route cost, default route type) tuples. If the ASBR learns of the network route for an element of the set, with the corresponding AS_PATH, then it generates a default route into the OSPF routing domain, with cost "default route cost" and type, "default route type". The lowest cost default route will then be injected into the OSPF

routing domain.

This is the recommended method for originating default routes in the OSPF routing domain.

3. BGP Identifier and OSPF router ID

The BGP identifier must be the same as the OSPF router id at all times that the router is up.

This characteristic is required for two reasons.

i. Consider the scenario in which 3 routers, RT1, RT2, and RT3, belong to the same autonomous system.

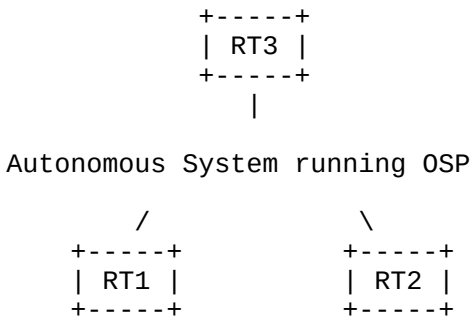

Both RT1 and RT2 have routes to an external network X and import it into the OSPF routing domain. RT3 is advertising the route to network X to other external BGP speakers. RT3 must use the OSPF router ID to determine whether it is using RT1 or RT2 to forward packets to network X and hence build the correct AS_PATH to advertise to other external speakers.

More precisely, RT3 must use the AS_PATH of the route announced by the ASBR, whose BGP Identifier is the same as the OSPF routerID corresponding to its route for network X.

ii. It will be convenient for the network administrator looking at an ASBR to correlate different BGP and OSPF routes based on the identifier.

4. Setting OSPF tags, BGP ORIGIN and AS_PATH attributes

The OSPF external route tag is a "32-bit field attached to each external route . . . It may be used to communicate information between AS boundary routers; the precise nature of such information is outside the scope of [the] specification." [RFC1247]

OSPF imports information from various routing protocols at all its ASBRs. In some instances, it is possible to use protocols other than EGP or BGP across autonomous systems. It is important, in BGP, to differentiate between routes that are external to the OSPF routing domain but must be considered internal to the AS, as opposed to routes that are external to the AS.

Routes that are internal to the AS and that may or may not be external to the OSPF routing domain will not come to the various BGP speakers via IBGP. Therefore, ASBRs running BGP must have knowledge of this class of routes so that they can advertise these routes to the various external AS without waiting for IBGP updates about these routes.

Additionally, in the specific instance of an AS intermixing routers running EGP and BGP as external gateway routing protocols, using OSPF as an IGP, the network administrator does not have to configure IBGP on every ASBR running EGP and not running BGP, if this information can be carried in the OSPF tag field.

We use the external route tag field in OSPF to intelligently set the ORIGIN and AS_PATH attributes in BGP. Both the ORIGIN and AS_PATH attributes are well-known, mandatory attributes in BGP. The exact mechanism for setting the tags is defined below.

The tag is broken up into sub-fields shown below. The various sub-fields specify the characteristics of the route imported into the OSPF routing domain.

The high bit of the OSPF tag is known as the "Automatic" bit. When this bit is set to 1, the following sub-fields apply:

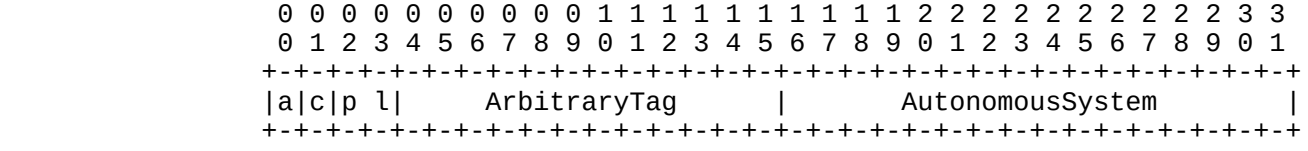

a is 1 bit called the Automatic bit, indicating that the Completeness and PathLength bits have been generated automatically by a router. The meaning of this characteristic and its setting are defined below.

c is 1 bit of Completeness information. The meaning of this characteristic and its settings are defined below.

pl are 2 bits of PathLength information. The meaning of this characteristic and its setting are defined below.

ArbitraryTag (or "at")
is 12 bits of tag information, which defaults to 0 but can be configured to anything else.

AutonomousSystem (or "as")
is 16 bits, indicating the AS number corresponding to the route, 0 if the route is to be considered as part of the local AS.

When the Automatic bit is set to 0, the following sub-fields apply:

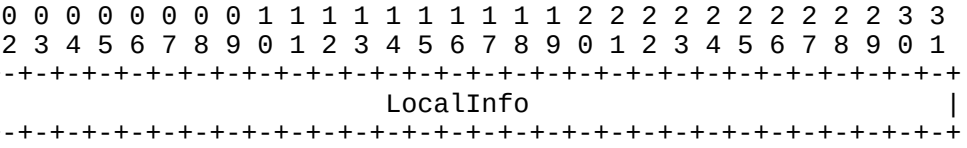

a is 1 bit called the Automatic bit, set to 0.

LocalInfo (or "li")
is 31 bits of an arbitrary value, manually configured by the network administrator.

The format of the tag for various values of the characteristics bits is defined below.

4.1. Semantics of the characteristics bits

The Completeness and PathLength characteristics bits define the characteristic of the route imported into OSPF from other ASBRs in the autonomous system. This setting is then used to set the ORIGIN and NEXT_HOP attributes when re-exporting these routes to an external BGP speaker.

- The "a" bit or the Automatic characteristic bit is set when the Completeness and PathLength characteristics bits are automatically set by a border router.

  For backward compatibility, the Automatic bit must default to 0 and the network administrator must have a mechanism to enable automatic tag generation. Nothing must be inferred about the characteristics of the OSPF route from the tag bits, unless the tag has been automatically generated.

- The "c" bit of the Completeness characteristic bit is set when the source of the incoming route is known precisely, for instance, from an IGP within the local autonomous system or EGP at one of the autonomous system's boundaries. It refers

to the status of the path information carried by the routing protocol.

- The "pl" or the PathLength characteristic sub-field is set depending on the length of the AS_PATH that the protocol could have carried when importing the route into the OSPF routing domain. The length bits will indicate whether the AS_PATH attribute for the length is zero, one, or greater than one.

  Routes imported from an IGP will usually have an AS_PATH of length of 0, routes imported from an EGP will have an AS_PATH of length 1, BGP and routing protocols that support complete path information, either as AS_PATHs or routing domain paths, will indicate a path greater than 1.

  The OSPF tag is not wide enough to carry path information about routes that have an associated PathLength greater than one. Path information about these routes will have to be carried via IBGP. Such routes must not be exported from OSPF into BGP.

For brevity in the following sections, the keywords O and P refer to the BGP ORIGIN and AS_PATH attributes respectively. Likewise, we use the abbreviations , "l" and "nh" for the local_AS and next_hop_AS respectively in the following sections.

4.2. Configuration parameters for setting the OSPF tag

- There must be a mechanism to enable automatic generation of the tag characteristic bits.

- Configuration of an ASBR running OSPF must include the capability to associate a tag value, for the ArbitraryTag, or LocalInfo sub-field of the OSPF tag, with each instance of a routing protocol.

- Configuration of an ASBR running OSPF must include the capability to associate an AS number with each instance of a routing protocol.

  Associating an AS number with an instance of an IGP is equivalent to flagging those set of routes imported from the IGP to be external routes outside the local autonomous system.

  Specifically, when the IGP is RIP [RFC1058], it should be possible to associate a tag and/or an AS number with every

interface running RIP on the ASBR.

4.3. Manually configured tags

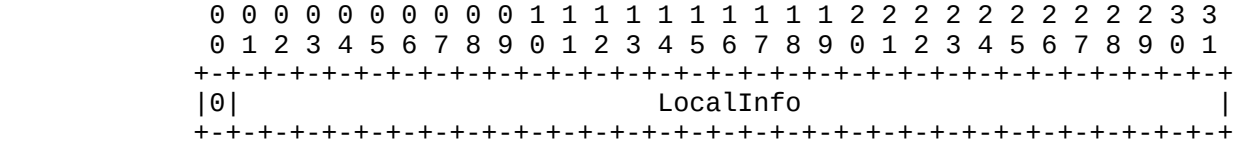

This tag setting corresponds to the administrator manually setting the tag bits. Nothing shall be inferred about the characteristics of the route corresponding to this tag setting.

For backward compatibility with existing implementations of OSPF currently deployed in the field, this must be the default setting for importing routes into the OSPF routing domain. There must be a mechanism to enable automatic tag generation for imported routes.

The OSPF tag to BGP attribute mappings for these routes must be
a=0, li=Arbitrary_Value => O=<INCOMPLETE>, P=<l>

4.4. Automatically generated tags

4.4.1. Routes with incomplete path information, pl = 0.

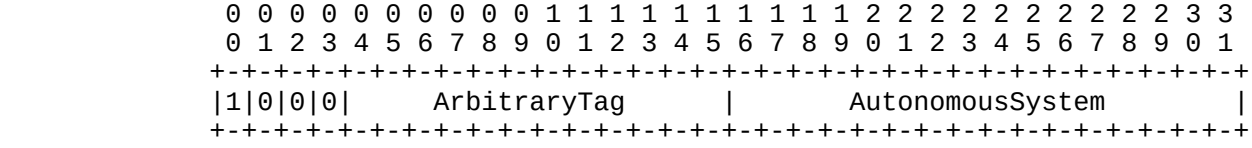

These are routes imported from routing protocols with incomplete path information and cannot or may not carry the neighbour AS or AS path as part of the routing information.

The OSPF tag to BGP attribute mappings for these routes must be
a=1,c=0,pl=00,as=0 => O=<EGP>, P=<l>

4.4.2 Routes with incomplete path information, pl = 1.

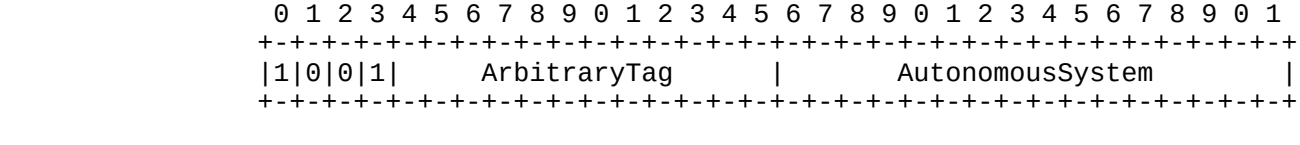

These are routes imported from routing protocols with incomplete

path information and carry the neighbour AS as part of the routing information.

The OSPF tag to BGP attribute mappings for these routes must be

a=1,c=0,pl=01,as=nh => O=<EGP>, P=<l, nh>

This setting should be used for importing EGP routes into the OSPF routing domain. This setting can also be used when importing BGP routes whose origin=<EGP> and AS_PATH=<nh>; if the BGP learned route has no other transitive attributes, then its propogation via IBGP can be suppressed.

4.4.3. Routes with incomplete path information, pl >= 1.

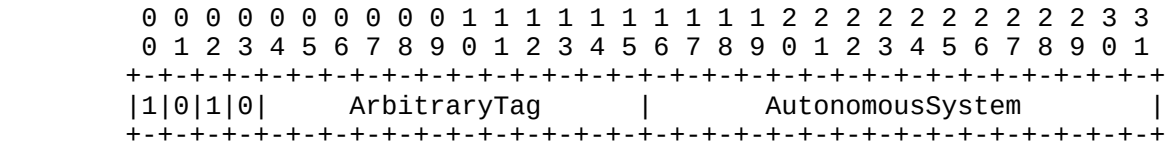

These are routes imported from routing protocols with truncated path information.

The OSPF tag to BGP attribute mappings for these routes must be

a=1,c=0,pl=10,as=don’t care

These are imported by a border router, which is running BGP to a stub domain, and not running IBGP to other ASBRs. This causes a truncation of the AS_PATH. These routes must not be re-exported into BGP at another ASBR.

4.4.4. Routes with complete path information, pl = 0.

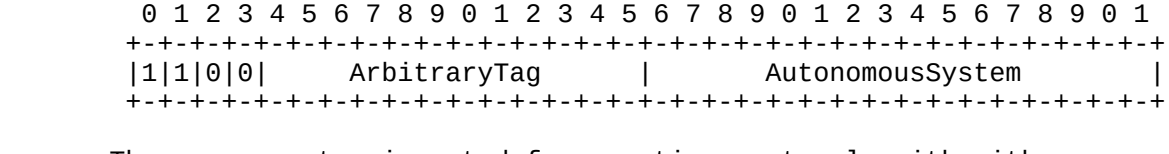

These are routes imported from routing protocols with either complete path information or are known to be complete through means other than that carried by the routing protocol.

The OSPF tag to BGP attribute mappings for these routes must be

a=1,c=1,pl=00,as=0 => O=<IGP>, P=<l>

This should be used for importing routes into OSPF from an IGP.

4.4.5. Routes with complete path information, pl = 1.

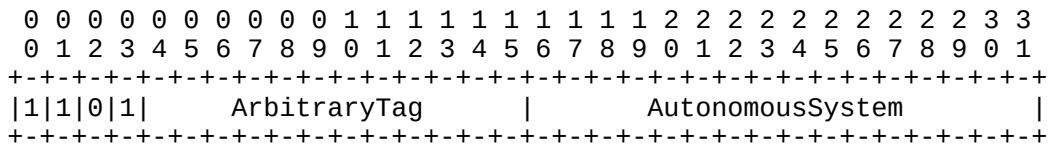

These are routes imported from routing protocols with either complete path information, or are known to be complete through means other than that carried by the routing protocol. The routing protocol also has additional information about the neighbour AS of the route.

The OSPF tag to BGP attribute mappings for these routes must be
a=1,c=1,pl=01,as=nh => O=<IGP>, P=<l, nh>

This setting should be used when the administrator explicitly associates an AS number with an instance of an IGP. This setting can also be used when importing BGP routes whose origin=<IGP> and AS_PATH=<nh>; if the BGP learned route has no other transitive attributes, then its propogation via IBGP can be suppressed.

4.4.6. Routes with complete path information, pl >= 1.

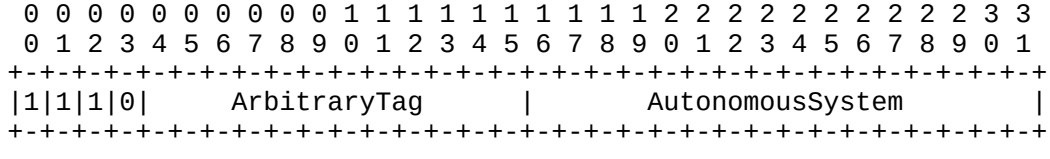

These are routes imported from routing protocols with complete path information and carry the AS path information as part of the routing information.

The OSPF tag must be set to
a=1,c=1,pl=10,as=don’t care

These routes must not be exported into BGP because these routes are already imported from BGP into the OSPF RD. Hence, it is assumed that the BGP speaker will convey this information to other BGP speakers via IBGP.

Note that an implementation may import BGP routes with a path length of 1 and no other transitive attributes directly into OSPF and not send these routes via IBGP. In this situation, it must use tag settings corresponding to 4.1.2.2, or 4.1.2.5.

4.5. Miscellaneous tag settings

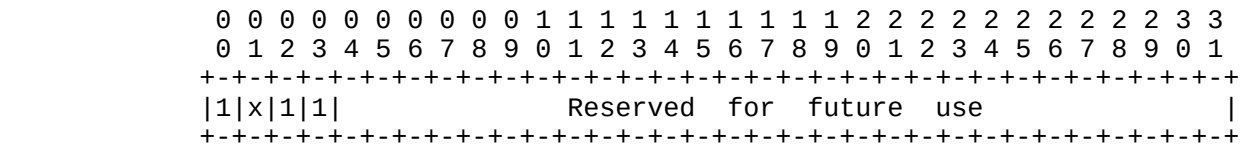

The value of pl = 3 is reserved during automatic tag generation. Routers must not generate such a tag when importing routes into the OSPF routing domain. ASBRs must ignore tags which indicate a pl = 3.

4.6. Summary of the tag sub-field setting

The following table summarises the various combinations of automatic tag settings for the Completeness and PathLength sub-field of the OSPF tag and the default behaviour permitted for each setting.

Completeness := 0 | 1 ;
PathLength := 00 | 01 | 10 | 11;
ORIGIN := <INCOMPLETE> | <IGP> | <EGP>;

AS_PATH := valid AS path settings as defined in BGP.

	pl = 00	pl = 01	pl = 10	pl = 11
c = 0	<EGP><l>	<EGP><l,nh>	never export	reserved
c = 1	<IGP><l>	<IGP><l,nh>	out of band	reserved

The "out of band" in the table above implies that OSPF will not be able to carry everything that BGP needs in its routing information. Therefore, some other means must be found to carry this information. In BGP, this is done via IBGP.

5. Setting OSPF Forwarding Address and BGP NEXT_HOP attribute

Forwarding addresses are used to avoid extra hops between multiple routers that share a common network and that speak different routing protocols with each other.

Both BGP and OSPF have equivalents of forwarding addresses. In BGP, the NEXT_HOP attribute is a well-known, mandatory attribute. OSPF has a Forwarding address field. We will discuss how these are to be filled in various situations.

Consider the 4 router situation below:

RT1 and RT2 are in one autonomous system, RT3 and RT4 are in another.
RT1 and RT3 are talking BGP with each other.
RT3 and RT4 are talking OSPF with each other.

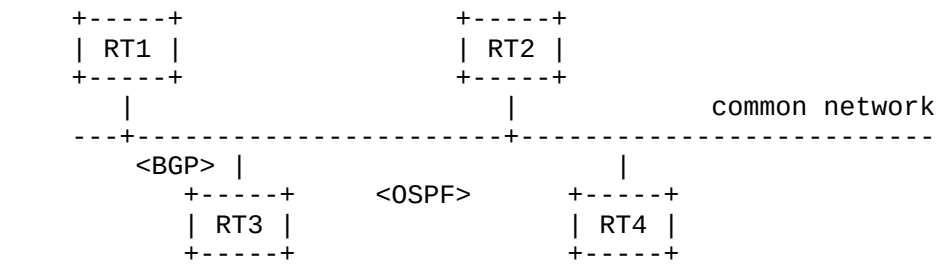

- Importing network X to OSPF:

  Consider an external network X, learnt via BGP from RT1.

  RT3 must always fill the OSPF Forwarding Address with the BGP

  NEXT_HOP attribute for the route to network X.

- Exporting network Y to BGP:

  Consider a network Y, internal to the OSPF routing domain, RT3's route to network Y is via RT4, and network Y is to be exported via BGP to RT1.

  If network Y is not a subnetted network, RT3 must fill the NEXT_HOP attribute for network Y with the address of RT4. This is to avoid requiring packets to take an extra hop through RT3 when traversing the AS boundary. This is similar to the concept of indirect neighbour support in EGP [RFC888, RFC827].

6. Security Considerations

Security considerations are not discussed in this memo.

7. Acknowledgements

I would like to thank Yakov Rekhter, Jeff Honig, John Moy, Tony Li, and Dennis Ferguson for their help and suggestions in writing this document, without which I could not have written this document. I would also like to thank them for giving me the opportunity to write this document, and putting up with my muddlements through various phases of this document.

I would also like to thank the countless number of people from the OSPF and BGP working groups who have offered numerous suggestions and comments on the different stages of this document.

8. Bibliography

[RFC827] Rosen, E., "Exterior Gateway Protocol (EGP)", RFC 827, BBN, October 1982.

[RFC888] Seamonson, L., and E. Rosen, "STUB Exterior Gateway Protocol", RFC 888, BBN, January 1984.

[RFC1058] Hedrick, C., "Routing Information Protocol", RFC 1058, Rutgers University, June 1988.

[RFC1267] Lougheed, K., and Y. Rekhter, "A Border Gateway Protocol 3 (BGP-3)", RFC 1267, cisco Systems, T.J. Watson Research Center, IBM Corp., October 1991.

[RFC1268] Rekhter, Y., and P. Gross, Editors, "Application of the Border Gateway Protocol in the Internet", RFC 1268, T.J. Watson Research Center, IBM Corp., ANS, October 1991.

[RFC1247] Moy, J., "The OSPF Specification - Version 2:", RFC 1247, Proteon, January 1991.

[ROUTE-LEAKING] Almquist, P., "Ruminations on Route Leaking", Work in Progress.

[NEXT-HOP] Almquist, P., "Ruminations on the Next Hop", Work in Progress.

9. Author's Address

Kannan Varadhan
Internet Engineer, OARnet
1224 Kinnear Road
Columbus, OH 43212-1136

EMail: kannan@oar.net