Inter-Domain Routing
Internet-Draft
Intended status: Standards Track
Expires: April 10, 2016

H. Green
E. Zimmer
Blue Ridge Envisioneering, Inc.

October 8, 2015

DDoS-Alert Extensions
draft-green-idr-ddosae-00

Abstract

This document defines extensions to BGP-4 to enable the exchange of information about detected malicious traffic (e.g., Distributed Denial of Service Attacks) and provide options for coordinated, collaborative responses to mitigate such traffic. The extensions are backward compatible - a BGP speaker that supports the extensions can interoperate with speakers that do not support the extensions.

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 April 10, 2016.

Copyright Notice

Table of Contents

1. Introduction

Distributed Denial of Service (DDoS) attacks pose a significant risk to network operations. Mitigating these attacks requires a coordinated response, as many systems do not have the capacity to work through a large scale attack. BGP enabled devices are also likely to have the ability to filter and/or throttle traffic; they are also widely distributed throughout networks, making them ideal for mitigating DDoS attacks.

DDoS-AE provides an open, vendor agnostic, mechanism to enable network devices to rapidly disseminate information about detected attacks; thereby, enabling a distributed response to mitigate the detected attacks. A key advantage of DDoS-AE over other solutions [RFC5575] is that the DDoS Alert messages can traverse over BGP speakers that do not directly support the extension, allowing greater dissemination of information about ongoing network attacks. An optional feature in the DDoS-AE system is interfacing to a Central Service (CS) for bridging the gap between DDoS-AE BGP speakers that are not connected, and to receive tailored DDoS response cues to improve coordination and efficacy of the response to the detected attacks.

Participants in the DDoS-AE system do not have to implement traffic filtering or DDoS detection mechanisms to still benefit and contribute to the overall system. For example, if a device or policy limits the ability to perform filtering and/or throttling of identified malicious traffic, the device could still generate alert messages when it detects new attack traffic. Similarly, if a device does not have the capability to inspect traffic and detect attacks, it could still receive alerts and implement traffic policies to mitigate the reported attacks. Finally, if all a device does is forward the DDoS-AE alerts between DDoS-AE participants it still improves the ability of the system as a whole to detect and mitigate attacks.

Because some attacks may attempt various techniques for concealment in legitimate traffic, more advanced and complex descriptions/ signatures of the traffic may be required to ensure minimal impact to legitimate traffic. In these more complex cases, the DDoS-AE system offers the option to report detailed signatures through the web-based Central Service (CS), which will then coordinate responses with participants using a more rich set of traffic descriptors that would be too difficult and cumbersome to include in BGP messages. The BGP messages in these cases are still useful as a first response, however, as they can enable participants to begin throttling traffic matching a more course signature; reducing the effects of the attack and minimizing impacts to legitimate traffic matching the course signature. Participants interfacing with the CS then would receive verbose traffic signatures enabling them to setup targeted policies that take more severe actions to matching traffic, such as dropping the packets entirely.

To simplify the introduction of DDoS-AE a new optional, transitive, attribute is introduced into BGP-4 that will contain the information needed to identify and respond to malicious traffic. The DDoS-AE attribute (DDOSAE_ALERT) will specify information about identified attack traffic in a standardized, yet minimal manner, so that devices can implement traffic policies to help mitigate the attack. Guidelines are also defined for how devices should respond to received DDoS-AE alert messages, beyond the core protocol message exchange functions. Details about the interface to the CS are not included in this description as they are auxiliary to the functions of the described BGP extensions.

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 [RFC2119].

2. DDoS-AE Alert Attribute - DDOSAE_ALERT (Type Code TBD1)

This is an optional transitive attribute that can be used to distribute information about malicious traffic, i.e. Distributed Denial of Service (DDoS) attack traffic, called Alerts. The DDoS-AE Alert Attribute is included on UPDATE messages [RFC4271] where the advertised NLRI is the detected target of a network attack. By following the existing rules for BGP route processing, information regarding the attack to the specified network can be efficiently propagated to devices that may transport traffic destined to the network under attack.

Because there may be multiple types of attacks targeting the same destination at any given time, this attribute may contain multiple Alert entries. The Attribute Length field for the Path Attribute and the Alert Length fields in the individual entries are used to determine the individual Alert entry boundaries.

The attribute is encoded as one or more entries of the following fields shown below:

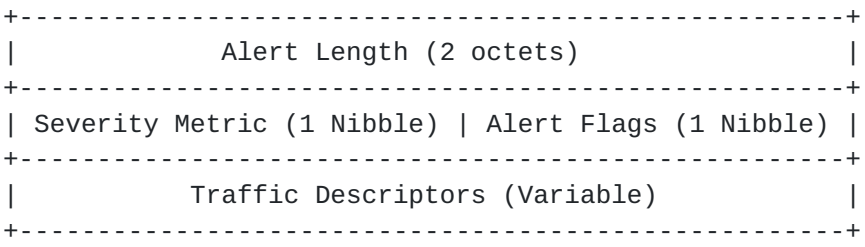

Figure 1: DDoS-AE Alert Attribute

2.1. Attribute Field Definitions

Alert Length
: This field is used to differentiate between multiple Alert entries for a given target prefix. It is a 2 octet field describing the length in octets of the current Alert entry. The length count includes the 2 octets of the Alert Length field. It can be used to completely skip over an Alert entry during processing if an unrecognized Traffic Descriptor or error is found.

Severity Metric (SM)
: This field is used to report the measured severity of the reported attack traffic. This field is also used by the Alert Distribution Process.

  The value is calculated by the node generating the alert based on the measured rate of the described attack traffic observed by that

node in relation to the total amount of all measured traffic at the observing node. The ratio is then normalized so that it ranges between 0 and 15, where a value of 15 indicates the attack traffic has saturated the observing node.

A value of 0 SHOULD not be used because it means there is no longer an attack detected. If that was the case, then the entire attribute for the target should be removed, either by sending another UPDATE for the same target, with the DDoS-AE Alert attribute removed, or by sending an UPDATE removing the specific route entirely.

Alert Flags

This field is used to provide additional information about the processing state of the information included in the Alert message. It is a 4 bit field consisting of the following flags:

Reported to CS Flag (CS)

High order bit (0) that when set (1) indicates that the Alert message has been reported to the Central Service (CS). This allows nodes that do not interact with the CS to report Alerts and have other nodes that do interact with the CS ensure the Alert is reported.

Drop Safe Flag (DS)

Second high order bit (1) that when set (1) indicates that the description in this Alert contains sufficient detail that nodes are encouraged to completely drop all matching traffic. When not set (0), the implication is the description may match a significant amount of legitimate traffic and dropping that traffic would not be recommended, in this case bandwidth throttling policies would be the preferred response.

Reserved Flags

Bits 2 - 3 are currently reserved.

Traffic Descriptors

A variable length field that lists Traffic Descriptors that further describes the attack traffic being reported. Traffic Descriptors are encoded as the following triplet:

<Type (1 Octet), Length (1 Octet), Value (Variable)>

Descriptor Type is a one octet field that identifies the traffic descriptor being described. See Section 2.2 for a complete listing of available Traffic Descriptor Types and their associated Value encoding.

Descriptor Length is a one octet field that contains the length of the Descriptor Value field in octets. Descriptor Value is a variable length field that is interpreted according to the value of the Descriptor Type field.

Some Descriptor Types MAY appear multiple times in one Alert message. If a Descriptor Type entry conflicts with a previous entry in the same Alert message then the later entry SHOULD be ignored. If a node detects an unknown or unsupported Descriptor Type it MAY ignore the value in the Response Action, however, it MUST maintain the entry for distribution to other nodes.

2.2. Traffic Descriptor Types

Traffic Descriptors are used to further describe attack traffic so that it can be targeted more accurately, minimizing impact to legitimate traffic on a network. These Traffic Descriptors have been selected and designed to be high level, generic, and flexible to ensure compatibility with as many traffic filtering/policing implementations as possible. Specifically, the descriptors are such that they do not require a filter to maintain state of traffic streams, meaning these descriptors should be compatible with any stateless filter.

To minimize complexity in the Alerts and ease interpretation by traffic filtering/policing implementations all Traffic Descriptor entries in an Alert SHOULD be considered to be the minimum criteria for matching described traffic. In other words, ALL supported Traffic Descriptor entries in an Alert SHOULD be satisfied by traffic in question in order to be considered a match. If attack traffic cannot be completely distinguished from legitimate traffic using the provided Traffic Descriptors then the Drop Safe flag SHOULD be set to 0.

This document defines the following values for Traffic Descriptor Types:

0 - IP Protocol / Next Header

Value Encoding: 1 Octet Integer

Value of the IPv4 Protocol field or IPv6 Next Header field. An entry of this type MUST be specified if any of the protocol independent convenience Descriptor Types are present in the Alert. Valid values are those found in the IANA Assigned Internet Protocol Numbers [RFC5237][RFC7045].

1 - IP Protocol / Next Header Compare

Value Encoding: Compare Triplet (Section 2.3)

Compare the value of the IPv4 Protocol field or IPv6 Next Header field.

2 - Source Port Compare

Value Encoding: Compare Triplet (Section 2.3)

Protocol independent way to compare the Source Port of the Transport Layer protocol of the described traffic. How this field is applied depends on the value of the IP Protocol / Next Header (Type 0) entry.

3 - Destination Port Compare

Value Encoding: Compare Triplet (Section 2.3)

Protocol independent way to compare the Destination Port of the Transport Layer protocol of the described traffic. How this field is applied depends on the value of the IP Protocol / Next Header (Type 0) entry.

4 - Network Header Offset Compare*

Value Encoding: Offset Compare Quadlet (Section 2.4)

Used to compare a value at a specific offset from the start of the Network Layer (IPv4/IPv6) header.

5 - Transport Header Offset Compare*

Value Encoding: Offset Compare Quadlet (Section 2.4)

Similar to Network Header Offset Compare, except the start of the offset begins at the beginning of the first Transport Layer Protocol Header. This allows for variable length options in the Network Layer Protocol Header.

6 - ANY IP Options Compare*

Value Encoding: Compare Triplet (Section 2.3)

Specify comparisons to perform over the IP Options present in the subject packet. A match is valid if ANY of the IP Options present in the subject packet evaluate to true for the specified comparison.

7 - ALL IP Options Compare*

Value Encoding: Compare Triplet (Section 2.3)

Like Any IP Options, but ALL present IP Options in subject packet must evaluate to true for the specified comparison.

8 - NO IP Options Compare*

Value Encoding: Compare Triplet (Section 2.3)

The opposite of All IP Options, in that NONE of the present IP Options must evaluate to true for the specified comparison.

9 - First Fragment

Value Encoding: No Value Needed

Match packets that are the first of a fragmented packet series.

10 - Is Fragment

Value Encoding: No Value Needed

Match packets that are not the first of a fragmented packet series, but are trailing fragments.

11 - Not Fragment

Value Encoding: No Value Needed

Match packets that are not fragmented.

12 - TTL/Hop Limit Compare

Value Encoding: Compare Triplet (Section 2.3)

Protocol independent way to compare value of the TTL/Hop Limit. How this field is applied depends on the value of the IP Protocol / Next Header (Type 0) entry.

13 - TCP Initial

Value Encoding: No Value Needed

Match packets that are the initial packet in a TCP connection. Essentially looking for TCP packets with ACK flag set to 0 and SYN flag set to 1. Should only have an effect if the

IP Protocol / Next Header (Type 0) is present with a value of TCP (6).

14 - TCP Established

Value Encoding: No Value Needed

Match packets that are not the initial packet in a TCP connection. Essentially looking for TCP packets with the ACK or RST flags set. Should only have an effect if the
IP Protocol / Next Header (Type 0) is present with a value of TCP (6).

15 - TCP Flags Compare*

Value Encoding: Compare Triplet (Section 2.3)

Compare values of TCP flags. Should only have an effect if the IP Protocol / Next Header (Type 0) is present with a value of TCP (6).

16 - ICMP Type Compare

Value Encoding: Compare Triplet (Section 2.3)

Compare value of ICMP Type field. Should only have an effect if the IP Protocol / Next Header (Type 0) is present with a value of ICMP (1) or ICMPv6 (58).

17 - ICMP Code Compare

Value Encoding: Compare Triplet (Section 2.3)

Compare value of ICMP Code field. Should only have an effect if the IP Protocol / Next Header (Type 0) is present with a value of ICMP (1) or ICMPv6 (58).

* - Indicates Traffic Descriptor Type may be present more than once per Alert. Unless otherwise specified there SHOULD be no more than one entry per Traffic Descriptor Type per Alert.

2.3. Compare Triplet Encoding

The Compare Triplet is used by several Traffic Descriptor types to specify a comparison operator, and comparator value. The Compare Triplet is encoded as the following triplet:

<Compare Operator (1 Octet), Length (1 Octet), Value (Variable)>

Compare Operator is a 1 octet field specifying the comparison operator/match behavior. Compare operators are defined in Section 2.5.

Comparator Value Length is a 1 octet field containing the length in octets of the Comparator Value field.

Comparator Value is a variable length field containing the value to use in the comparison operation.

2.4. Offset Compare Quadlet Encoding

The Offset Compare Quadlet is similar to the Compare Triplet (Section 2.3), but adds a 2 octet Offset Amount field to the beginning of the Triplet.The Compare Quadlet is encoded as the following quadlet:

<Offset (2 Octets), Compare Operator (1 Octet), Length (1 Octet), Value (Variable)>

The Offset Amount field is a 2 octet value specifying the offset in bytes. The starting point for offset calculation is dependent on the context in which the type is used. The other fields have the same definition as in the Compare Triplet Encoding (Section 2.3).

2.5. Compare Operator Definitions

This document defines the following values for Compare Operators:

0 - Match

Match the exact value.

1 - Mask

Perform bit-wise AND operation then match result to the mask value.

2 - Less Than (<)

Determine if the value at the specified offset is < the Comparator Value.

3 - Greater Than (>)

Determine if the value at the specified offset is > the Comparator Value.

4 - Not Equal (!=)

Determine if the value at the specified offset is != to the Comparator Value.

5-255 - Reserved

Reserved for future use.

3. Alert Processing

An Alert is used to describe detected malicious traffic so that participants in the DDoS-AE system can coordinate a response to mitigate the attack. Alerts leverage existing BGP processes for exchanging NLRI and therefore the same rules for NLRI announcements are followed. This helps ensure that Alerts are generated by speakers about network segments with which they have a legitimate interest, and ensures the Alerts are propagated only to other speakers that also have concern with the network under attack.

Alerts are target centric, meaning they focus on malicious traffic streams destined to the same target. The target could be a single host or an entire subnet. While it is possible that one party could direct a single attack against multiple targets, for the purposes of DDoS-AE each distinct subnet target would be considered a unique attack for Alert generation purposes. Due to the nature of DDoS attacks, there will likely be multiple sources generating the malicious traffic destined to the identified target.

3.1. Alert Generation/Updating

Alerts are generated when a participating node detects a new attack or malicious traffic stream. The details of how malicious traffic streams are detected are outside the scope of this document and left up to the discretion of the node implementing this extension. It is recommended that system designers allow for flexibility in the generation of alerts so they may be generated in both an automated and manual fashion.

When a new malicious traffic stream is detected at a DDoS-AE node, an Alert is generated by sending an UPDATE message advertising an updated NLRI message for the detected traffic stream destination. The UPDATE should follow the existing BGP rules for propagation to peers as if any other optional transitive attribute regarding the route had been updated.

The content of the Alert attribute SHOULD be minimal, with sufficient detail to accurately describe the malicious traffic, while avoiding

legitimate traffic. If an organization detects an attack that is targeting multiple addresses in their network block, then it would be recommended to generate the Alert for the smallest possible subnet capturing the addresses under attack. However, if there is the possibility that portions of the advertised subnet are not under attack and there is the potential that another sub-organization is using portions of that address space, then it is RECOMMENDED to generate multiple Alerts for each minimal address block, rather than one Alert for a larger block that encompasses more addresses than are really under attack.

In many cases, due to attack traffic masquerading as legitimate traffic, it may be very difficult to distinguish legitimate traffic from malicious traffic. In these cases the Drop Safe flag should be cleared so that speakers implementing filters know to simply throttle matching target. In cases where the attack traffic can be perfectly described in the content of the Alert and virtually all legitimate traffic can be excluded, the Drop Flag SHOULD be set so that participating speakers implementing filters know it is safe to drop matching traffic completely.

The Severity Metric (SM) field SHOULD be set to a non-zero value based on the ratio of observed malicious traffic to legitimate traffic at the reporting node. A zero value would mean no traffic is observed, in which case, sending an Alert is meaningless and wasteful. See Alert Removal section for details about removing previous Alerts.

3.2. Alert Distribution

Alerts are distributed using the same mechanism as regular NLRI in BGP, through UPDATE messages. The same rules for processing UPDATE NLRI and distributing the NLRI should be followed. This is effective at distributing the Alert to speakers that may be in position to help mitigate the attack by following the reverse path of the incoming attack traffic. It also minimizes the Alerts that are sent to speakers that may not be able to assist in mitigating the detected attack. The DDoS-AE Alert attribute SHOULD NOT be used in the decision process for route selection.

3.3. Alert Aggregation

Alert aggregation is possible following the same rules as route aggregation in general. The DDoS-AE Alert attribute may be aggregated by combining the individual Alert entries within each of the aggregated DDoS-AE Alert Attributes, dropping duplicate entries.

Individual DDoS-AE Alert entries within a given DDoS-AE Alert Attribute may be further aggregated if the Traffic Descriptor entries all match. The Severity Metric value should contain the maximum value of the aggregated Alert entries. The Reported to CS Flag value is set if any of the aggregated Alerts have this flag set. The Drop Safe (DS) flag SHOULD be set to 0, unless all of the aggregated Alerts have this flag set.

3.4. Alert Removal

Alerts can be removed two ways:

1. Removing the advertised route using the Withdrawn Routes field in the UPDATE message (or the MP_UNREACH_NLRI attribute in [RFC4760]).

2. Sending an updated advertisement for the route but removing the DDoS-AE Alert attribute, or removing the specific Alert entry from the DDoS-AE Alert attribute in the updated advertisement.

4. BGP Capability Advertisement

A BGP speaker that uses DDoS-AE SHOULD use the Capability Advertisement procedures [RFC5492] to determine whether the speaker could use DDoS-AE with a particular peer and if any optional DDoS-AE features may be enabled. However, because DDoS-AE does not introduce new message types and the DDoS-AE path attributes are transitive optional, speakers MAY send Alert messages to peers in order to enable the possibility that the Alert values are passed on beyond the non-DDoS-AE peer and eventually make it to another indirectly connected DDoS-AE speaker.

To indicate support for DDoS-AE the Capability Optional Parameter Code field is set to TBD2 (requesting 74 in IANA Considerations (Section 7)). The Capability Length field is set to the value that minimally captures all the bits representing the supported optional DDoS-AE capabilities. Currently this length is 0.

5. Alert Refresh

Because DDoS-AE Alerts are distributed as attributes of existing NLRI, the ability to refresh information about active Alerts comes free with any BGP speaker that supports existing Route Refresh capabilities [RFC7313].

6. Acknowledgements

The authors would like to thank Dr. Dan Massey and the Cyber Security Divison (CSD) at the Department of Homeland Security (DHS) for their support of this effort. The authors would like to thank Nick Richard, David Fox, Andrew Krause, Keyur Patel, and Donald Sharp for support developing this concept. The authors also appreciate the support from the Quagga development community for the prototyping effort and the USC ISI DETER testbed team for providing the resources for evaluating the prototype system.

7. IANA Considerations

IANA is requested [RFC5226] to assign a BGP Path Attribute code through Standards Action [RFC4271]. The BGP Path Attribute code value requested is 30. The label for the requested BGP Path Attribute is requested to be DDOSAE_ALERT. It is referenced in this document as TBD1 (Section 2). The IANA registry for BGP Path Attributes is located at <http://www.iana.org/assignments/bgp-parameters/bgp-parameters.xhtml>.

IANA is requested [RFC5226] to assign a BGP Capability Code from the First Come First Served range [RFC5492]. The BGP Capability Code value requested is 74. It is referenced in this document as TBD2 (Section 4). The IANA registry for BGP Capability Codes is located at <http://www.iana.org/assignments/capability-codes/capability-codes.xml>.

Value	Description	Reference
TBD1 (30)	BGP Path Attribute Type Code (DDOSAE_ALERT)	[RFC4271]
TBD2 (74)	BGP Capability Code (DDoS-AE Capability)	[RFC5492]

IANA Considerations Summary

8. Security Considerations

Exchanging information about detected malicious traffic, relies on the same trust relationship already present between BGP speakers. On its own, the exchange of traffic descriptors adds no additional security concerns to BGP. The trust and security levels are maintained because the Alerts are target centric, so the speaker that is announcing the Alert must also be advertising the network prefix associated with the Alert. Therefore existing policies and rules provide the assurance that the source of the Alert is the organization that is also the victim of the described attack(s). Scenarios where a false or malicous Alert might be issued are no different than what a poorly behaived BGP speaker might do, and can

be mitigated using the same techniques used to account for potentially bad BGP speakers.

Organizations that execute traffic shaping based on received Alerts should take care to ensure the source of the Alert is the same organization that they would expect to be advertising the NLRI on its own. This ensures the same degree of trust and security that is already inherent in BGP (for better or for worse).

Implementing traffic shaping in response to dynamic Alerts could make troubleshooting network issues more difficult. It is recommended that organizations generate detailed logs and human readable alerts whenever new traffic shaping policies are executed as a result of an Alert.

It is possible that malicious actors could specify traffic descriptors in an Alert to match NLRI destinations other than those in the associated NLRI announced by the BGP speaker. This could cause incautious routers to effect traffic destined to destinations other than the one in the associated NLRI update message. It is recommended that participants ensure the resulting traffic shaping policies only effect traffic destined to the addresses associated with the NLRI in the update message.

9. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <http://www.rfc-editor.org/info/rfc2119>.

[RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <http://www.rfc-editor.org/info/rfc4271>.

[RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 4760, DOI 10.17487/RFC4760, January 2007, <http://www.rfc-editor.org/info/rfc4760>.

[RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, DOI 10.17487/RFC5226, May 2008, <http://www.rfc-editor.org/info/rfc5226>.

[RFC5237] Arkko, J. and S. Bradner, "IANA Allocation Guidelines for the Protocol Field", BCP 37, RFC 5237, DOI 10.17487/RFC5237, February 2008, <http://www.rfc-editor.org/info/rfc5237>.

[RFC5492] Scudder, J. and R. Chandra, "Capabilities Advertisement with BGP-4", RFC 5492, DOI 10.17487/RFC5492, February 2009, <http://www.rfc-editor.org/info/rfc5492>.

[RFC5575] Marques, P., Sheth, N., Raszuk, R., Greene, B., Mauch, J., and D. McPherson, "Dissemination of Flow Specification Rules", RFC 5575, DOI 10.17487/RFC5575, August 2009, <http://www.rfc-editor.org/info/rfc5575>.

[RFC7045] Carpenter, B. and S. Jiang, "Transmission and Processing of IPv6 Extension Headers", RFC 7045, DOI 10.17487/RFC7045, December 2013, <http://www.rfc-editor.org/info/rfc7045>.

[RFC7313] Patel, K., Chen, E., and B. Venkatachalapathy, "Enhanced Route Refresh Capability for BGP-4", RFC 7313, DOI 10.17487/RFC7313, July 2014, <http://www.rfc-editor.org/info/rfc7313>.

Authors' Addresses

Harley Green
Blue Ridge Envisioneering, Inc.
5180 Parkstone Dr
Chantilly, Virginia 20151
USA

Email: harley@br-envision.com
URI: http://www.br-envision.com

Edward R. (Ned) Zimmer
Blue Ridge Envisioneering, Inc.
5180 Parkstone Dr
Chantilly, Virginia 20151
USA

Email: ned@br-envision.com
URI: http://www.br-envision.com