Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 29, 2013

J. Gersch
Secure64 SW Corp
D. Massey
C. Olschanowsky
Colorado State University
L. Zhang
UCLA

February 25, 2013

DNS Resource Records for Authorized Routing Information
draft-gersch-grow-revdns-bgp-02

Abstract

This draft discusses the use of two DNS record types for storing BGP routing information in the reverse DNS. The RLOCK record allows prefix owners to indicate whether the DNS is being used to publish routing data. The SRO record allows operators to indicate whether an IPv4 or IPv6 prefix ought to appear in global routing tables and identifies authorized origin Autonomous System Number(s) for that prefix. The resulting published data can be used in a variety of contexts from routing security to address ownership.

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 August 29, 2013.

Table of Contents

1. Introduction

1.1. Overview

This draft describes a method in which a prefix owner can exploit the existing reverse DNS tree structure, along with the authentication provided by DNSSEC [RFC4033], to publish information about whether a prefix can be announced and identify the origin Autonomous System(s) that may originate a route to that prefix. This data is complementary to a variety of other data sources ranging from existing databases to new directions.

Publishing route information in the Reverse DNS takes advantage of infrastructure that already exists and has been globally deployed. No new infrastructure deployment is required, in contrast with approaches that use purpose-built resource certification.

Other key advantages to using the Reverse DNS are that it 1) has been in successful operation for many years, 2) has an existing operational model where prefix owners currently manage their IP address space (through various models from local operation to hosting companies), 3) has an existing operational model where both registries and providers delegate authority to entities receiving address space, 4) the resulting reverse DNS data can be authenticated using DNSSEC [RFC4033], and 5) the data can be easily checked using simple tools ranging from DNS query tools such as DIG to more elaborate systems.

A prefix owner must OPT-IN to the approach. Prefix owners who do not take any action are not impacted, but also do not gain any advantages. Prefix owners that do choose to participate would thereby enable a number of tools to make use of the published data. The objective of this draft is to standardize the format for indicating participation and publishing data. A variety of potential uses for the data are discussed later in the document, but are provided only to illustrate the usefulness of the data and should not be taken as a comprehensive list of all possible applications.

1.2. Scope

We limit the scope of this internet draft to associating an origin AS number with a prefix. Armed with this information, one can address both origin hijacks and sub-prefix hijacks in a manner that can be deployed in a reasonable time frame. Future expansion is readily made possible: the SRO record is kept simple for now, but is designed to reference additional future records that enable additional capabilities.

2. Conventions Used In This Document

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].

3. Overview of Route Publishing

This document defines two new DNS resource records types (RRTypes) and describes how these record types can be associated with prefixes.

1. The RLOCK RRType (Route Lock)

    * Purpose: Indicates that the Reverse DNS zone has enabled BGP route publishing.

    * The presence of the RLOCK Record at the apex of a Reverse DNS zone indicates that a prefix owner has OPTED-IN to BGP Route Publishing. All route announcements that map to this zone will be denied as INVALID unless an SRO record exists that specifically authorizes the announcement.

2. The SRO RRType (Secure Route Origin)

    * Purpose: Declare an authorized route origin ASN for comparison against BGP route announcements.

    * Placed in the Reverse DNS at the domain name corresponding to the associated CIDR address block.

Organizations that have been assigned and/or allocated CIDR address blocks also have Reverse-DNS delegations assigned to them from either the Regional Internet Registries (RIPE, ARIN, APNIC, etc.) or from a sub-delegation.

Address-block owners may use these record types to declare authoritative data for route origins associated with that address block. This data may be declared statically, with a long TTL (Time To Live) if the routing data changes infrequently. Alternatively, dynamic DNS and short TTLs can be used to rapidly publish and disseminate the authoritative information on a world-wide basis in near real-time.

The RLOCK and SRO records are to be stored in the reverse-DNS in zones with domain names that correspond to the associated CIDR address block. These domain names are to be constructed per the naming specification described in [I-D.gersch-dnsop-revDNS-CIDR].

The RLOCK and SRO records MUST be signed with DNSSEC and have a valid DNSSEC chain-of-trust.

4. Overview of Route Verification

Various applications could be written to use BGP records published in the Reverse DNS. One example is an application to perform near-real-time route origin verification that alerts operators of hijacks or directly interacts with a router to prevent the hijack. Another application could perform a nightly analysis that generates router prefix filters. A third application could cross-check data in the Internet Routing Registries (IRR) against the data in the reverse DNS. This list is not intended to be comprehensive, but instead aims to illustrate the potential uses of the published data.

These applications analyze BGP announcements by performing DNS queries to classify route route announcements into one of the following three categories:

1. "VALID": a DNSSEC-validated SRO RRSET was received and one of the route origins in the RRSET matches the origin contained in the BGP route announcement.

2. "INVALID": a route hijack was detected.

   A. The DNSSEC-validated SRO responses received did NOT match the origin of the route announcement. This is indicative of an origin hijack.

   B. There was no SRO record at the domain name corresponding to this address block, but the authoritative zone did contain an RLOCK statement. This is indicative of a sub-prefix hijacks.

3. "NOTFOUND": there was no SRO record for this prefix and no RLOCK record to protect the zone, or the data did not properly validate with DNSSEC. In this case, the algorithm cannot authoritatively state that the prefix is valid or invalid, so it is simply marked as not found. Most routes today are in this category, as it takes a specific action to OPT-IN to this methodology.

This verification algorithm MUST "fail-safe". If a query for a DNS record fails, or if DNSSEC fails to validate the record, the algorithm MUST behave as if no DNS records were present in the first place. This results in marking a BGP announcement as "NOTFOUND". One could completely unplug a router verification application at any time and internet routing would continue to work just as it does today. The default state is always "not found".

Note that this implies the verification algorithm MUST use DNSSEC-enabled queries (set the DO bit) and MUST check for a validated response (the AD bit). A successful DNSSEC-downgrade attack would

result in classifying records as "NOTFOUND". However the redundancy in DNS would allow checking of multiple slave DNS servers should DNSSEC fail to validate.

The core of the verification algorithm can be summarized as follows:

1. Given a BGP announcement from a router feed, offline database, prefix filter, or other source, perform a DNSSEC-validated query for the SRO records at the domain name corresponding to the CIDR prefix in the BGP announcement.

2. Case 1: If no records exist (NXDOMAIN or NOERROR with number of answers=0), use the AUTHORITY section of the answer to determine the covering zone. Perform a query to that domain name (the zone apex) for an RLOCK record. There are two possible responses to the RLOCK query:

   A. NOERROR, answer=0: the RLOCK does not exist; the zone owner has not opted in. Mark the announcement as "NOTFOUND".

   B. RLOCK exists: the zone owner has OPTED-IN. Mark the announcement as "INVALID" since no SRO record exists to vouch for the announcement. This may be an example of a sub-prefix hijack.

3. Case 2: One or more SRO records were returned from the query. Loop through each SRO in the RRSET to compare the origin with the data in the route announcement. If a record with a matching set of data is found, mark the announcement as "VALID". If no match is found, mark the announcement as "INVALID".

5. The RLOCK Resource Record

The RLOCK resource record indicates "Route Lock". This record is placed at the apex of a reverse-DNS zone to indicate that the zone is being used to publish routing information. If this record is present, all route announcements for the CIDR address block covered by this zone MUST be marked as "INVALID" unless they are specifically authorized by a SRO record.

The main purpose of the RLOCK statement is to indicate participation (OPT-IN) and as a side-effect prevent sub-prefix route hijacks. Applications that query for an SRO record may get an NXDOMAIN or NOERROR with 0 answers. In this case, the application queries the domain name specified in the AUTHORITY section for an RLOCK record (this will be at the zone apex). If the RLOCK is present, the route announcement MUST be marked as "INVALID". Otherwise there is no SRO and no RLOCK, so the route announcement MUST be marked as "NOTFOUND".

The effective span of control for an RLOCK is dependent on the structure of the Reverse DNS zone. To be more specific, a Reverse DNS zone that has no delegations will have a span of control that covers all prefixes at or below the CIDR prefix specified by the domain name at the zone apex. Any zone delegation (also known as a "cut point") starts a new zone authority. Those prefixes in the delegated zone will not be covered by the parent zone's RLOCK. As an example, consider the zone at 129.82.0.0/16 and assume that it has only one delegation at 129.82.138.0/24. The /16 RLOCK covers all prefixes within the /16 to /32 range with the exception of prefixes within the 129.82.138.0/24 through /32 range. The child zone would need to have its own RLOCK.

The RLOCK record MUST be signed with DNSSEC and have an associated RRSIG record. If a resolving DNS server cannot validate the DNSSEC signature, the SRO record should be ignored as if it were not even present in the zone.

The Type value for the RLOCK RR type is currently unassigned. We are temporarily using private RRTYPE TYPE65400 until a formal number is assigned by IANA.

The RLOCK RR is class independent.

The RLOCK RR has no special TTL requirements.

Example use of RLOCK records, taken directly from the current testbed, are included in the appendix.

5.1. RLOCK RDATA Wire Format

The SRO RDATA wire format consists of one optional field, a 4 octet Activation Time field. If the Activation Time field is not present, the RDATA length is 0 octets. If the Activation Time field is present, the RDATA Length is 4 octets.

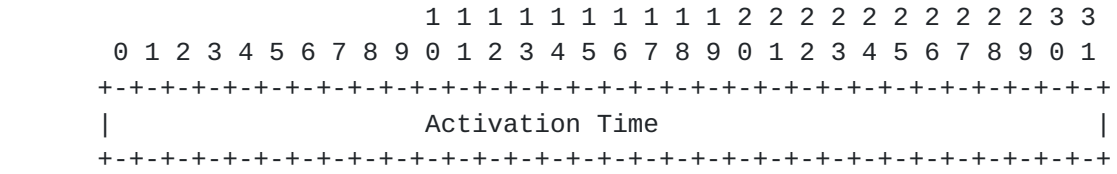

5.1.1. The Activation Time field

The optional Activation Time field specifies a starting date and time from which the RLOCK record is considered to be active and may be used for route origin validation. The RLOCK record MUST NOT be used for validation prior to this activation time.

The Activation Time field value specifies a date and time in the form of a 32-bit unsigned number of seconds elapsed since 1 January 1970 00:00:00 UTC, ignoring leap seconds, in network byte order. The longest interval that can be expressed by this format without wrapping is approximately 136 years.

If the Activation Time field is not present, or if it contains a value of 0, immediate activation for the RLOCK record is in effect.

The purpose of the Activation Time field is to permit publication of RLOCK records in the reverse DNS prior to its use in formal route validation. This enables two key capabilities: first,it allows for the testing of RLOCK records in a safe manner by informing an application to "don't validate, but tell me what you would have done". Second, it allows an ISP to publish an RLOCK and SRO record that defines a large covering prefix to give advance warning to that ISP's customers. Customers that have their own AS number and a delegated address block within the ISP's larger block may want to publish their own SRO record if they share a zone with the ISP, or they will risk having their announcements being marked INVALID because the ISP has claimed a larger covering prefix. Alternatively, the customer may be independent from the ISP's zone and manage their own reverse-DNS zone that has been delegated to them. In this case they may choose to publish an SRO record or to do nothing. The cut point of the zone delegation stops the ISP's covering prefix from extending into the new zone.

5.2. RLOCK Presentation Format

The presentation format of the RDATA portion is as follows:

```
[ ACTIVATION_TIME ]
```

where the ACTIVATION TIME is an optional field.

The optional Activation Time field value, if present, MUST be represented either as an unsigned decimal integer indicating seconds since 1 January 1970 00:00:00 UTC, or in the form YYYYMMDDHHmmSS in UTC, where:

YYYY is the year (0001-9999);
MM is the month number (01-12);
DD is the day of the month (01-31);
HH is the hour, in 24 hour notation (00-23);
mm is the minute (00-59); and
SS is the second (00-59).

Note that it is always possible to distinguish between these two formats because the YYYYMMDDHHmmSS format will always be exactly 14 digits, while the decimal representation of a 32-bit unsigned integer can never be longer than 10 digits.

5.3. RLOCK RR Examples

The following examples shows an RLOCK resource record that enables routing security for the zone covering 129.82.0.0/16. The first example specifies immediate activation. The second specifies activation starting on July 4, 2013 at 9:30 UTC.

```
82.129.in-addr.arpa.  86400  IN    RLOCK

120.15.in-addr.arpa.  86400  IN    RLOCK  20130704093000
```

6. The SRO Resource Record

Zones that participate in this approach to BGP route security use "Secure Route Origin" (SRO) resource records to identify authorized origin Autonomous System Number(s) for a prefix.

The SRO record contains one mandatory field, the ORIGIN ASN field. The SRO record MAY contain also contain one to three optional fields in the following order: a FLAGS field, a PREFIX LIMIT field, and an ACTIVATION TIME field. These fields MUST be appended to the RDATA in that specific order.

The SRO record MUST be signed with DNSSEC [RFC4033] and have an associated RRSIG record. If a resolving DNS server cannot validate the DNSSEC signature, the SRO record should be ignored and an attempt should be made to query an alternate DNS server. If all servers fail, the route prefix should be classified as "NOTFOUND".

The Type value for the SRO RR type is currently unassigned. We are temporarily using TYPE65401 until a formal number is assigned by IANA.

The SRO RR is class independent.

The SRO RR has no special TTL requirements.

6.1. SRO RDATA Wire Format

The SRO RDATA wire format consists of four fields: a 4 octet ORIGIN AS NUMBER field, a 1 octet FLAGS field, a 1 octet PREFIX LIMIT field, and a 4 octet SRO ACTIVATION TIME field. The RDATA Length is a constant 10 octets. Missing fields in the presentation format are filled with 0's in the wire data format.

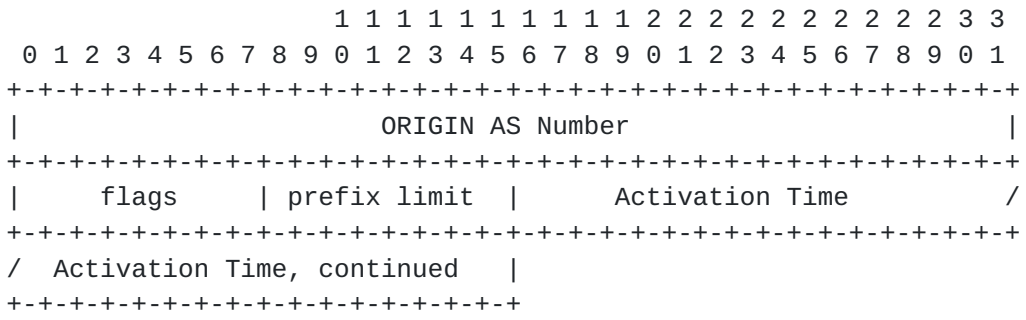

6.1.1. The Origin AS Number field

The Origin AS Number field is used to specify an AS Number that is authorized to originate a route announcement for the CIDR address block corresponding to the SRO record's reverse DNS domain name. If a CIDR address block can be announced from multiple AS numbers, then multiple SRO records should be defined at the domain name corresponding to that CIDR address block.

The Origin AS Number field accommodates both 2 octet and 4 octet AS Numbers. 2 octet AS numbers MUST be encoded with leading zeroes to construct a complete 4 octet field.

6.1.2. The Flags field

The Flags field is reserved for future expansion; currently there are no flags defined and the Flags field, if present, MUST be set to 0. In the future various bits of the field will be used to define extensions and additional records related to the SRO record.

6.1.3. The Prefix Limit field

The Prefix Limit field specifies the maximum length of a prefix that the SRO statement is authorizing. If the field contains a value of 0 then the origin AS is only authorized for the exact prefix corresponding to that domain name.

Any other integer value from 1 to 32 for IPv4 and 1 to 128 for IPv6 specifies the length of the most specific IP prefix that the SRO is authorizing. For example, if an IP address prefix is 10.0/16 and the prefix limit is 22, the AS is authorized to advertise any prefix under 10.0/16, as long as it is no more specific than /22. In this example, the AS would be authorized to advertise 10.0/16, 10.0.128/20 or 10.0.255/22, but not 10.0.255.0/24. An SRO statement must be present at each domain name corresponding to these prefixes (see Appendix A for a discussion on how to create multiple SRO records using wildcard DNS names). The Prefix Limit field is used to inform a route validation algorithm that if an SRO response is returned for a domain name past the prefix limit, the SRO statement MUST be ignored as if it were not present.

6.1.4. The Activation Time field

The Activation Time field specifies a starting date and time from which the RLOCK record is considered to be active and may be used for route origin validation. The RLOCK record MUST NOT be used for validation prior to this activation time.

The Activation Time field value specifies a date and time in the form of a 32-bit unsigned number of seconds elapsed since 1 January 1970 00:00:00 UTC, ignoring leap seconds, in network byte order. The longest interval that can be expressed by this format without wrapping is approximately 136 years.

If the Activation Time field contains a value of 0, immediate activation for the RLOCK record is in effect.

The purpose of the Activation Time field is to permit publication of SRO records in the reverse DNS prior to its use in formal route validation. This enables two key capabilities: first,it allows for the testing of RLOCK records in a safe manner by informing an application to "don't validate, but tell me what you would have done". Second, it allows an ISP to publish an RLOCK and SRO record that defines a large covering prefix to give advance warning to that ISP's customers. Customers that have their own AS number and a delegated address block within the ISP's larger block may want to publish their own SRO record if they share a zone with the ISP, or they will risk having their announcements being marked INVALID because the ISP has claimed a larger covering prefix. Alternatively, the customer may be independent from the ISP's zone and manage their own reverse-DNS zone that has been delegated to them. In this case they may choose to publish an SRO record or to do nothing. The cut point of the zone delegation stops the ISP's covering prefix from extending into the new zone.

6.2. SRO RRDATA Presentation Format

The presentation format of the RDATA portion is as follows:

```
ORIGIN_ASN [ FLAGS [ PREFIX_LIMIT [ [ ACTIVATION_TIME ] ] ]
```

where ORIGIN_ASN is the mandatory Origin AS Number field, and the remaining fields are optional. If not present, the wire data format will be filled with zeroes.

The Flags field is represented by an unsigned integer. The current accepted value MUST be 0.

The Prefix Limit field is represented by an unsigned integer in the range 0 to 128.

The ORIGIN AS NUMBER field is represented in asdot notation which is a combination of asplain and asdot+ notation. That is, any ASN in the 2-octet range is represented in asplain (simple decimal representation of the ASN). Any ASN above the 2-octet range is represented in asdot+ notation which breaks an ASN into two 16-bit

values separated by a dot. For example, AS65535 will be represented by the decimal number "65535" while AS65536 will be represented as "1.0".

The Activation Time field values MUST be represented either as an unsigned decimal integer indicating seconds since 1 January 1970 00:00:00 UTC, or in the form YYYYMMDDHHmmSS in UTC, where:

YYYY is the year (0001-9999);
MM is the month number (01-12);
DD is the day of the month (01-31);
HH is the hour, in 24 hour notation (00-23);
mm is the minute (00-59); and
SS is the second (00-59).

Note that it is always possible to distinguish between these two formats because the YYYYMMDDHHmmSS format will always be exactly 14 digits, while the decimal representation of a 32-bit unsigned integer can never be longer than 10 digits.

6.3. SRO RR Examples

The following example shows an SRO RR authorizing AS12145 as the origin for CIDR address block 129.82.0.0/16 and only for that prefix. Its activation time is immediate.

```
m.82.129.in-addr.arpa.  86400  IN  SRO 12145
```

The second example shows the use of a prefix limit field to authorize AS 12145 to cover the range from 129.82/16 through and including announcements to the /24 limit. Its activation is set to June 1, 2013 at 00:00 UTC.

```
m.82.129.in-addr.arpa.  86400  IN  SRO 12145 0 24 20130601000000
```

The third example shows two separate origins to be authorized for a prefix. This example also illustrates the use of the asdot notation. Two different prefix limits are used. Activation is immediate for the first SRO; the second is to be used after July 15, 2013 at noon UTC.

```
m.82.129.in-addr.arpa.  86400  IN  SRO 12145 0 24
                        86400  IN  SRO 3.421 0 18 20130715120000
```

7. Discussion and Related Work

This work is not the first to propose entering routing data in the Reverse DNS and there are also many other proposed approaches for publishing routing data. We first review some of the past work and then discusses the differences presented in this approach.

7.1. Prior Work on CIDR names for Routing

Over a decade ago, [I-D.bates-bgp4-nlri-orig-verif] proposed to use the reverse DNS to verify the origin AS associated with a prefix. This requires both a naming convention for converting the name into a prefix and additional resource record types for storing origin information, along with recommendations on their use. More recently [I-D.donnerhacke-sidr-bgp-verification-dnssec] including links to IRR data and also includes the notion of policy in adjacency, but this approach also introduces a new reverse DNS tree under "BGP.ARPA." CNAME and DNAME records must be used in publishing the data.

Our approach differs in several respects. We rely on the existing reverse DNS tree without creating a new hierarchy such as "BGP.ARPA.". We exploit the naming convention in [I-D.gersch-dnsop-revDNS-CIDR] so one does not need to introduce CNAME or DNAME records (though an operator could choose to do so if so desired). We assume optional participation and introduce the concept of an RLOCK resource record to indicate participation. We currently limit our approach to detecting false sub-prefix and false origin route announcements. Extensions to include links to other databases such as IRR can be achieved in combination with or in lieu of an SRO record and further path validation can be included, but the scope of this document is intentionally limited, both for clarity and to match actual implementation. Finally, we separate the publishing technique which is specified in this document from the variety of ways in which one may make use of the data, recognizing that different operators will make different choices on how to make use of the data.

7.2. RPKI

A great deal of work has been done in the sidr working group on Resource Public Key Infrastructure [RFC6480][RFC6481][RFC6482][RFC6483].

RPKI, also known as Resource Certification, is a specialized public key infrastructure (PKI) framework designed to secure Border Gateway Protocol (BGP). RPKI provides a way to connect Internet number resource information (such as Autonomous System numbers and IP Addresses) to a trust anchor. The certificate structure mirrors the

way in which Internet number resources are distributed. That is, resources are initially distributed by the IANA to the Regional Internet Registries (RIRs), who in turn distribute them to Local Internet Registries (LIRs), who then distribute the resources to their customers. RPKI can be used by the legitimate holders of the resources to control the operation of Internet routing protocols to prevent route hijacking and other attacks. [cited from Wikipedia].

The publication of BGP route origin information in the reverse-DNS is a complementary technique to RPKI. While there is some overlap in the techniques, there are also different goals for the reverse-DNS.

The Reverse-DNS publication method uses DNSSEC as its base trust model, not a chain of certificates. If an organization has a DNSSEC-signed delegation for a reverse-DNS address block, that organization is the legitimate owner and may place SRO and RLOCK statements in their zone without the interaction of any other organization. If an address block is sold or transferred, either the RIR (Regional Internet Registry) will change its signed delegation records to reflect the change, or the organization itself may independently implement a signed sub-delegation.

8. Security Considerations

Applications that query the DNS for SRO and RLOCK records MUST request them from DNSSEC-enabled servers and have the DO bit set. Responses that are returned MUST be checked to verify that the D bit is set indicating that the responses have been validated. Otherwise the response should be ignored.

The absence of DNSSEC or the inability to contact any nameservers MUST indicate the route is viable (NOTFOUND).

9. IANA Considerations

RRType numbers need to be assigned for the SRO and RLOCK records. The current testbed temporarily substitutes TYPE65400 for the RLOCK record and TYPE65401 for the SRO record.

10. Acknowledgments

We would like to thank Danny McPherson for his comments and suggestions. In addition, this document was aided via numerous discussions at NANOG, IETF and private meetings with ISPs, telecomm carriers, and research organizations too numerous to mention by name. Thanks to all for your comments and advice.

11. Change History

Changes from version 01 to 02

Removed the last_hop field from the SRO record

Changed the structure of the RLOCK to accommodate an activation date.

Changed the structure of the SRO record to accommodate activation date, prefix limit and flag fields. These changes were suggested from a year's practical experience in running a reverse-DNS testbed.

Added appendix A discussing the use of DNS wildcards with SRO prefix limits.

Modified the examples in Appendix B to reflect changes to the SRO and RLOCK records.

Changes from version 00 to 01

Removed all discussion of an "alternate naming scheme". A related draft on prefix naming had proposed both a standard naming scheme and an alternate naming scheme, but the alternate naming scheme was removed. Since the alternate naming scheme no longer exists, it was no longer necessary to discuss how this draft would deal with the alternate alternate naming scheme.

An optional transit provider field was added the SRO record.

Examples were updated based on the SRO change and operational experience.

Added Cathie Olschanowsky as a co-author

12. References

12.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

12.2. Informative References

[I-D.bates-bgp4-nlri-orig-verif]
Bates, T., Bush, R., Li, T., and Y. Rekhter, "DNS-based NLRI origin AS verification in BGP", draft-bates-bgp4-nlri-orig-verif-00 (work in progress), January 1998.

[I-D.donnerhacke-sidr-bgp-verification-dnssec]
Donnerhacke, L. and W. Wijngaards, "DNSSEC protected routing announcements for BGP", draft-donnerhacke-sidr-bgp-verification-dnssec-04 (work in progress), May 2008.

[I-D.gersch-dnsop-revDNS-CIDR]
Gersch, J. and D. Massey, "Reverse DNS Naming Convention for CIDR Address Blocks", draft-gersch-dnsop-revDNS-CIDR-04 (work in progress), May 2012.

[RFC1876] Davis, C., Vixie, P., Goodwin, T., and I. Dickinson, "A Means for Expressing Location Information in the Domain Name System", RFC 1876, January 1996.

[RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, March 2005.

[RFC6480] Lepinski, M. and S. Kent, "An Infrastructure to Support Secure Internet Routing", RFC 6480, February 2012.

[RFC6481] Huston, G., Loomans, R., and G. Michaelson, "A Profile for Resource Certificate Repository Structure", RFC 6481, February 2012.

[RFC6482] Lepinski, M., Kent, S., and D. Kong, "A Profile for Route Origin Authorizations (ROAs)", RFC 6482, February 2012.

[RFC6483] Huston, G. and G. Michaelson, "Validation of Route Origination Using the Resource Certificate Public Key Infrastructure (PKI) and Route Origin Authorizations

(ROAs)", RFC 6483, February 2012.

Appendix A. Discussion of Prefix Limits use in conjunction with DNS wildcards

The combination of RLOCK and SRO records may be used to detect and/or prevent sub-prefix hijacks. For example, the zone file shown in appendix B has one SRO record for the /16 and four SRO records for the /18's. If a BGP announcement is made at any other prefix, say a /19, no SRO record will be found. In this case the verification algorithm must query for the RLOCK record. If present, a sub-prefix hijack is indicated.

This behavior is desirable in most situations. On the other hand, some organizations may want the ability to advertise an announcement at any arbitrary sub-prefix (e.g. traffic control). The semantics of the SRO statement dictates that if an organization intends to advertise any prefix, an SRO statement must be present to authorize it. This can be done by statically creating multiple SRO statements in a zone for each prefix that could be announced, or by using dynamic DNS to add and remove SRO statements on demand.

There are two ways to create a range of SRO statements within a zone file. The first method is to simply create all of the individual SRO statements required and put them in the zone. To cover the full octet from /16 to /23, for example, the administrator would need one SRO for the /16, two for the /17's, four for the /18's, eight for the /19's, and so on. This would be a total of 255 SRO records. One can see that this becomes impractical for IPv6 address ranges; the number of SRO records to authorize a /32 through /64 is enormous.

To manage this situation, administrators may use the Prefix Limit field in combination with a DNS wildcard domain name. Assume that the administrator of address block 2002:1488::/32 wants to pre-authorize any announcement from the /32 up to and including /64. Assume also that the zone apex is at the /32. In this case a single SRO record in the zone will suffice:

```
*.8.8.4.1.2.0.0.2.ip6.arpa.  IN SRO 12345 0 64 0
```

A query for an SRO record for any prefix, (e.g. /32, or /36 or /48 or even a /96) covered by this zone (i.e., up to any zone cut) will return this SRO record and the announcement can be verified. The prefix limit tells the verification algorithm when to stop using the SRO statement. If an announcement is made for a /96, for example, the SRO record will be returned. Since the prefix limit is 64, the validation algorithm MUST ignore this SRO record as if it did not exist, and perform a query for the RLOCK record. Furthermore, any zone cut also stops the extent of the SRO statement. Any new zone declared below the /32 can declare (or not declare) its own SRO and

RLOCK statements.

This method allows the creation of zone files for a sparsely populated IPv6 delegation. Individual zones can be created that manage their own sub-prefixes, and the top level zone can handle all the covering prefixes.

Appendix B. Examples

B.1. Example 1

This example shows data entered for the prefix 129.82.0.0/16. The prefix owner has authorized the announcement of 129.82.0.0/16 and the four /18's at 129.82.0.0/18, 129.82.64.0/18, 129.82.128.0/18, and 129.82.192.0/18. All the prefixes originate from AS12145.

Note: this data is directly cut and paste from actual deployment. TYPE 65400 is being used for RLOCK and TYPE 65401 for SRO records. This draft requests IANA to assign numbers for RLOCK and SRO, the values here are purely for illustrative purposes.

```
$TTL 3600
$ORIGIN 82.129.in-addr.arpa.

@    IN    SOA     rush.colostate.edu.  dnsadmin.colostate.edu. (
                        2012082800      ; serial number
                        900             ; refresh, 15 minutes
                        600             ; update retry, 10 minutes
                        86400           ; expiry, 1 day
                        3600            ; minimum, 1 hour
                       )

     IN    NS      dns1.colostate.edu.
     IN    NS      dns2.colostate.edu.

@                  IN   TYPE65400 \#0
;                  RLOCK
;                          OPT-IN; deny all route announcements
;                          except those authorized
;                          effective date = immediate

m                  IN   TYPE65401 \# 10 00002f71000000000000
; 129.82.0.0/16           SRO  12145

0.0.m              IN   TYPE65401 \# 10 00002f71000000000000
; 129.82.0.0/18           SRO 12145

1.0.m              IN   TYPE65401 \# 10 00002f71000000000000
; 129.82.64.0/18          SRO 12145

0.1.m              IN   TYPE65401 \# 10 00002f71000000000000
; 129.82.128.0/18         SRO 12145

1.1.m              IN   TYPE65401 \# 10 00002f71000000000000
; 129.82.192.0/18         SRO 12145

;  delegations required for 256 /24 zones which contain PTR records

1   IN  NS  dns1.colostate.edu.
    IN  NS  dns2.colostate.edu.
2   IN  NS  dns1.colostate.edu.
    IN  NS  dns2.colostate.edu.

;  continuation to 255 is left out for the sake of brevity
```

B.2. Example 2

This example shows data entered for the prefix 216.17.128.0/17. The prefix owner has authorized the announcement of 216.17.128.0/17. The prefix originates from AS6582.

```
        1.m.17.216.in-addr.arpa  NS   ns.frii.net
```

This delegation refers to the new /17 zone and the domain name is not in conflict with any of the pre-existing /24 zones at IN-ADDR.ARPA. This delegation is to be placed at the IN-ADDR.ARPA zone.

```
   $TTL 3600
   $ORIGIN 1.m.17.216.in-addr.arpa.

   @    IN   SOA     ns1.frii.net.  hostmaster.frii.net. (
                          2012021300      ; serial number
                          14400           ; refresh, 4 hours
                          3600            ; update retry, 1 hour
                          604800          ; expiry, 7 days
                          600             ; minimum, 10 minutes
                         )

        IN   NS      ns1.frii.net.
        IN   NS      ns2.frii.net.

   $ORIGIN 17.216.in-addr.arpa.

   1.m               IN   TYPE65400 \# 0
   ;                      RLOCK
   ;                           OPT-IN; deny all route announcements
   ;                           except those authorized
   ;                           activation date = immediate

   1.m               IN   TYPE65401 \# 10 000019b6000000000000
   ; 216.17.128.0/17      SRO 6582

   ; no other delegations or PTR records are needed in this zone file
   ; since the /24 delegations are at ARIN at xxx.17.216.IN-ADDR.ARPA
```

Authors' Addresses

Joe Gersch
Secure64 SW Corp
Fort Collins, CO
US

Email: joe.gersch@secure64.com

Dan Massey
Colorado State University
Fort Collins, CO
US

Email: massey@cs.colostate.edu

Cathie Olschanowsky
Colorado State University
Fort Collins, CO
US

Email: cathie@cs.colostate.edu

Lixia Zhang
UCLA
Los Angeles, CA
US

Email: lixia@cs.ucla.edu