DNS Operations
Internet-Draft
Expires: March 11, 2006

S. Krishnaswamy
SPARTA Inc.

September 7, 2005

Split-View DNSSEC Operational Practices
draft-krishnaswamy-dnsop-dnssec-split-view-01.txt

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with Section 6 of BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/1id-abstracts.txt.

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html.

This Internet-Draft will expire on March 11, 2006.

Copyright Notice

Copyright (C) The Internet Society (2005).

Abstract

The security extensions to the Domain Name System (DNSSEC) allow for integrity protection, whereby it is possible to make a determination of the verity of data returned from the Domain Name System in response to a query. Current operation of the Domain Name System also allows for the creation of multiple views of data, where the answer returned in response to a query is dependent on the origin of the query. Data integrity and the ability to return possibly conflicting values as in split-views may be construed to be mutually conflicting goals; but this apparent dichotomy is resolvable in

practice through proper configuration. This document provides recommendations for correctly configuring the split-view DNSSEC environment in a typical enterprise network.

Table of Contents

1. Introduction 3
2. Split-View DNS 4
2.1. Background 4
2.2. Query Channeling 5
2.3. Controlling Errant Queries 7
2.4. Name Server Requirements 7
2.4.1. Internal Recursive Forwarder 7
2.4.2. Second-Level Recursive Name Server 7
2.4.3. Authoritative Internal and External-View Name Servers 7
3. Split-View DNSSEC 8
3.1. No Internal Validation 8
3.2. Same Key Signing 9
3.3. Partial Decoupling of Chains-Of-Trust 10
3.4. Complete Decoupling of Chains-Of-Trust 11
3.5. Multiple DS Records 11
3.6. Name Server Requirements 13
4. Packet Filtering Considerations 13
4.1. Inner Packet Filter 13
4.2. Outer Packet Filter 14
5. Summary 15
6. IANA Considerations 16
7. Security Considerations 16
8. Acknowledgements 16
9. References 17
9.1. Normative References 17
9.2. Informative References 17
Author's Address 18
Intellectual Property and Copyright Statements 19

1. Introduction

Split-view DNS is the term used to describe multiple views of DNS information for a domain based on where and by whom the query is sent. Split-views help contain DNS names to only those portions of the network that need to see these names. Although primarily meant to be a network management technique, the tailoring of the DNS to create an internal view of information hidden from the outside is also seen by some as improving their organization's security posture, by preventing the exposure of internal host names, knowledge of whose existence is deemed to be sensitive.

Relying solely on split-view DNS to protect sensitive hosts from attacks has proven to be less than adequate in the past. Attack vectors in recent Internet exploits have been able to successfully infect hosts with or without their IP addresses being published in the DNS. Conversely, publishing the IP addresses of hosts that are otherwise secured does not necessarily increase their vulnerability to these attacks. Name hiding through split-view DNS is primarily useful as part of a more comprehensive defense-in-depth strategy to provide one line of defense against name-based attacks.

The security extensions to DNS [1] provide for origin authenticity and data integrity. These properties are determined by validating the chain-of-trust from the signed record to some trusted key configured at the end resolver. In the case of split-view DNS every chain-of-trust in every view must validate properly. Some names may be common between multiple views but contain different data. Cache pollution is a possibility when data from the wrong view is returned in response to a query. Building a chain-of-trust from a trusted key above the zone that has split views, to data in the internal view of a zone can be especially problematic, caching problems notwithstanding.

The objective of this document is to describe approaches for configuring split-view DNSSEC environments with the additional requirement that no server be both authoritative and recursive at the same time. Separation of authoritative and recursive name servers not only provides simple role separation, but is also an important security measure in DNS for protecting authoritative name servers against compromised caches.

In cases where the different views of DNS information correspond to different physical networks, the name servers authoritative for the internal and external views of data are often separated by a firewall. Among some of the frequently observed DNS resolution misbehaviour [3] is the problem of resolvers aggressively retransmitting queries from behind misconfigured firewalls that allow

queries out, but drop all returned responses. This problem is exacerbated by a handful of errant queries that are sent by only a subset of internal resolvers, which makes problem isolation extremely difficult. This document provides recommendations for reducing the impact of errant queries in the split-view DNS setup and also makes recommedations for DNS-related packet filtering rules required to support the proper operation of the suggested configuration.

Section 2 describes the general approach for configuring split-view DNS, which by itself, is independent of DNSSEC. Considerations for DNSSEC appear in Section 3 .

2. Split-View DNS

2.1. Background

Different views of the DNS can be created by a process of "query channeling". Here, different servers are made authoritative for the different views of the DNS information and queries are channeled to these name servers based upon their origination address.

It is also possible to use a single machine as the authoritative name server for both views of data by running multiple instances of the name server process on a machine with multiple network interfaces, and answering differently based on the query source. Some name server implementations also directly support split-view DNS. Variants include the view-based approach and the data tagging approach. In the former, the name server loads multiple zone databases and makes available answers from a particular zone based on the origin of the query. The second approach tags the data in the database itself as either being internally, externally or globally available.

Single name server approaches are susceptible to leakage of DNS information if the host on which they operate is compromised. Confidentiality of the namespace is directly tied to how resilient the name server is against such attacks. A much better alternative to protect a namespace of sensitive hosts is to have that entire namespace reside within a private delegation. By doing so, it is possible to have the protection given to the name server that serves these names commensurate with the protection given to the hosts themselves. Since hosts in the private branch are explicitly marked as such by virtue of their domain name, this method also allows the network administrator to better classify hosts as being public or private and lessens the opportunity for sensitive hosts to be inadvertantly placed in public domains. Private delegations are useful when name hiding is the only reason for namespace separation.

They have the drawback that they do not allow for transparency during name resolution; queries have to be made for specific names in specific views.

This document describes a generic configuration for split-view DNS using multiple nameservers without relying on any special special capablities from any machine or name server implementation. The architecture is modeled around a typical enterprise structure: the two views are for the internal and external portions of the network, with the external portion residing within the boundary network. The two networks are separated by a packet filtering firewall. A packet filtering firewall also separates the boundary network from the external Internet. Name hiding is not an objective of this split-view setup, but avoiding cache pollution is. Although the two concepts are related, this configuration is not recommended for hiding sensitive names because of the ease with which names can be leaked out due to trivial configuration errors. Again, if name hiding is the main objective for providing split-views, the approach of using a private delegation for sensitive names is strongly encouraged.

The suggested configuration uses a combination of multiple name servers and query forwarding. One name server answers queries for the internal view and forwards all requests for external data to a second name server. The second name server recursively answers queries but only if asked by the first. Other name servers are configured in such a way so as to decouple the roles of the authoritative and recursive name servers.

2.2. Query Channeling

Query channeling is the process of carefully controlling how the queries are sent to different name servers so as to avoid cache pollution.

Resolving outer data is straightforward since queries follow their normative paths. For the internal view, a two-level recursive server scheme is recommended. One server functions as a recursive forwarder and is responsible for answering all internal queries. This server forwards all queries for internal data to their respective authoritative name servers while recursively obtaining external answers from a second-level name server. The authoritative name servers do not perform any recursion themselves. The second-level name server is a simple caching name server that asks questions from the outside, but only if asked by the recursive forwarder. The recursive forwarder and the name servers authoritative for the internal data reside in the internal network; the second-level recursive name server that is used for returning external answers

resides in the boundary network.

The two-level recursive scheme controls where queries are directed to. Since queries for internal data are sent to authoritative name servers which are not also recursive, this scheme also controls where data is received from. In this way internal data is kept totally separate from external data, thus preventing cache pollution. Figure 1 illustrates the above setup.

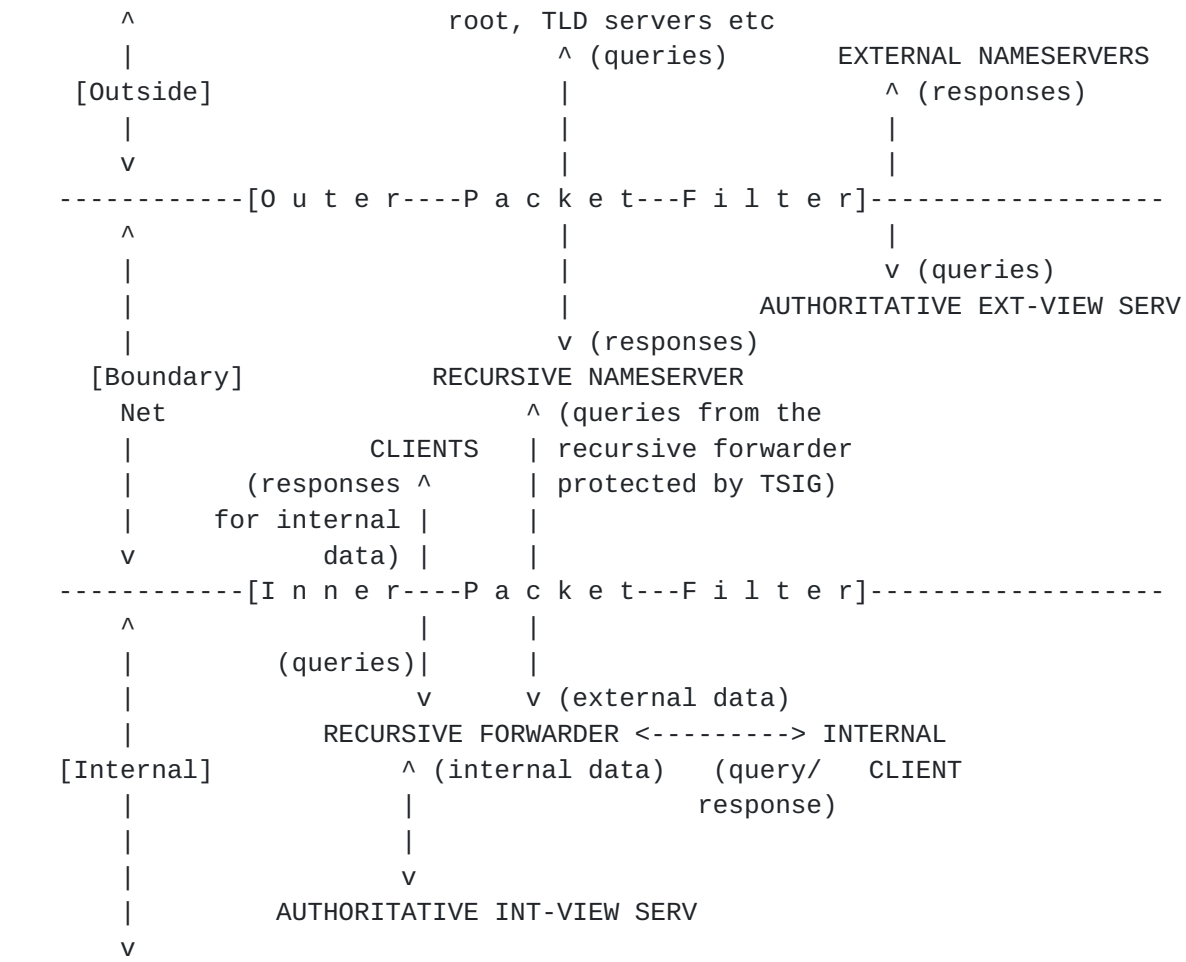

Figure 1

It is useful to note that the internal recursive forwarder must not attempt to recursively answer queries if the authoritative name server for internal-view data fails to respond. If it did so, external data could be returned in such circumstances and lead to cache pollution. Since neither the server authoritative for a forwarded zone nor the server doing the forwarding can recursively answer queries for delegations from that zone, the internal recursive forwarder must explicitly forward queries for every internal delegation to its respective authoritative name server. This rule

can be relaxed while forwarding queries to name servers that are simultaneously authoritative for the child as well as the parent zone.

2.3. Controlling Errant Queries

DNS queries that are sent from the internal recursive forwarder to the outside should only be directed towards the second-level recursive name server. Since the second-level name server has no knowledge of internal-view data, internal resolvers must not use it directly for resolving queries. Only properly configured internal recursive forwarders that are approved to send queries to this name server must do so, and that too solely for the purpose of resolving external answers. TSIG is the recommended method for controlling which recursive forwarders are approved for sending queries to the second-level name server. Having these rules alternatively configured in the packet filter is also possible, but using TSIG for performing this authorization eases packet filter administration for DNS.

2.4. Name Server Requirements

This section summarizes the list of requirements for the various name servers involved in the split-view configuration.

2.4.1. Internal Recursive Forwarder

- Ability to forward queries to specific name servers.
- Ability to control forwarding behaviour such that the recursive option is not tried, even if the name server that queries are normally forwarded to fails to respond.
- Ability to recursively answer queries.
- Ability to protect the integrity of messages using TSIG for selected destinations.

2.4.2. Second-Level Recursive Name Server

- Ability to recursively answer queries.
- Ability to verify TSIG protection on messages.
- Ability to filter incoming queries based on the TSIG key used to protect the message.

2.4.3. Authoritative Internal and External-View Name Servers

- Ability to authoritatively answer queries for a zone.
- Ability to disable all recursive behaviour.

3. Split-View DNSSEC

Any data in any view that is likely to be spoofed has to be signed. The DNSSEC concern for split-view is ensuring that the internal and external chains-of-trust validate properly. This concern is addressed by making an appropriate choice of trusted and Secure Entry Point (SEP) keys.

While validating external data is relatively straightforward, there are multiple approaches that can be used for validating internal data. The method of choice depends on what the threat environment for the internal view is perceived to be, the amount of end-resolver configuration overhead that is needed, the ease of debugging and the ability to have administrative separation between the two split-views. The configuration overhead at end resolvers is mainly associated with the task of defining trust anchors at different validating resolvers. Having fewer keys is desirable in that it makes key management easier. It is also desirable to reduce the amount of reconfiguration required for clients that move between the two views of data, while still being able to tie an answer to a particular view. Often two views of a split zone are administered separately, so having different zone signing keys for the different views is also desirable. The different options for internal data validation are further outlined below.

3.1. No Internal Validation

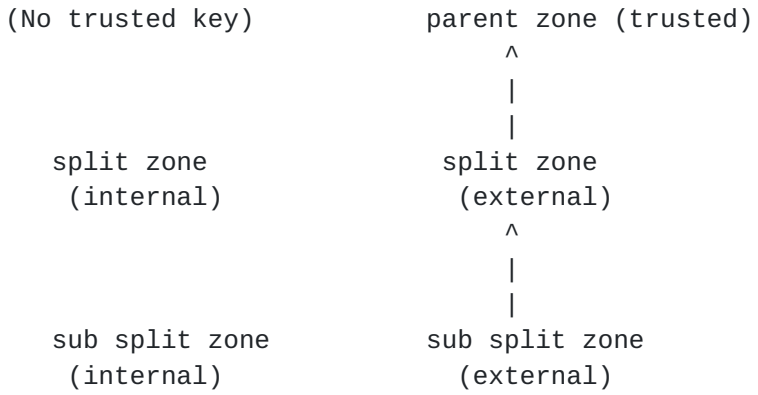

Figure 2

This is an option if the security requirements for the internal zone are more relaxed than the external zone. The threat environment for the internal zone in this scenario does not include DNS compromise and validation results returned from the internal recursive forwarder is not important. The internal recursive forwarder does not have any trusted key configured and does not perform any validation.

3.2. Same Key Signing

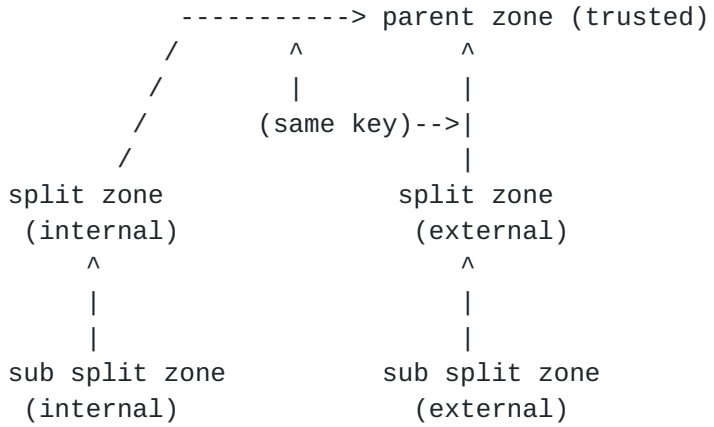

Figure 3

In this scenario, a single private key is used to sign both the internal and the external zone data. The glue DS and NS records at the delegation point of the split zone all correspond to the external view data. Validation proceeds by constructing two separate segments of the chain-of-trust. In the first segment, data at the level of the split and below is validated by constructing a chain-of-trust contained entirely within the internal view. If the trusted key is configured at or below the level of the split, validation stops at this point and queries are never sent to the outer view. If not, a second validation chain segment is constructed from the DS record covering the split to the trusted key. In forming the second validation segment all queries (including the query for the DS record of the split zone) are sent to the outer zone. Since the key referenced in the DS record is present in the apex key-set of both views, the chain-of-trust can be completed.

This approach allows flexibility in choosing the level at which the trusted key is configured, with the possibility of using the same trusted key for validating answers in both views.

Although easy to setup, this approach can be difficult to troubleshoot. There is no easy way to identify if the record obtained for a query corresponds to the internal view or the external view. Using the same key also makes administrative separation of the two views of data difficult.

3.3. Partial Decoupling of Chains-Of-Trust

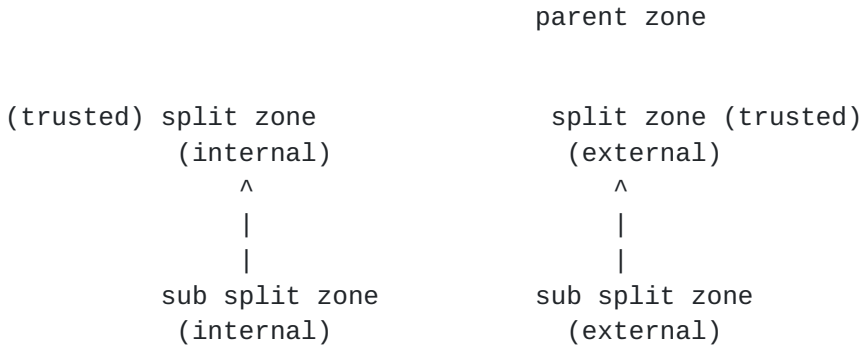

Figure 4

With the DS record in the parent always pointing to a key in the outer view, the construction of the chain-of-trust becomes problematic when the keys used to sign data in the two views of the split are different. The trusted key cannot be configured above the level of the split since there would be no way of linking the DS record in the outer zone to the apex DNSKEY set in the internal view of the split zone.

A simple solution is to configure the trusted key at the level of the split such that the chains-of-trust for the internal and external zones share no common records that might cause any ambiguity.

Having separate keys for the two views of data is useful for troubleshooting and in determining which view a given record belongs to. Cache pollution can be detected because such cases would lead to validation failures.

This configuration however involves more configuration overhead since trusted keys need to be configured for every zone that is split. This problem is more pronounced when dealing with validating stub resolvers on mobile nodes, where moving between the internal and external views would involve constant reconfiguration of its trusted keys.

3.4. Complete Decoupling of Chains-Of-Trust

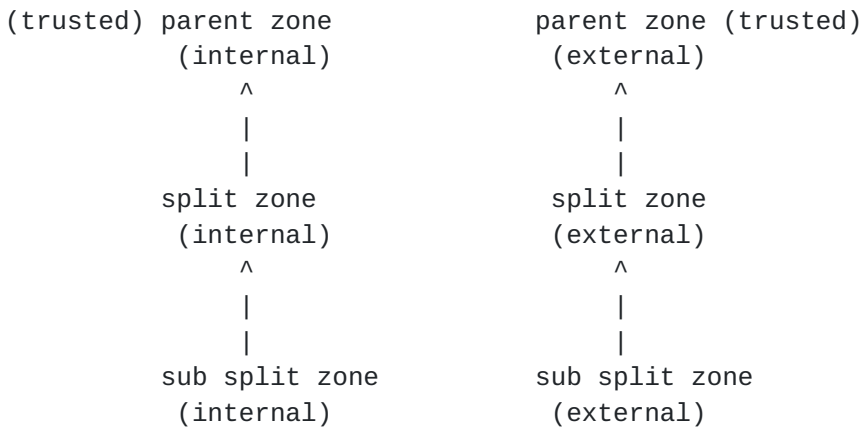

Figure 5

One problem with Section 3.3 is that trusted keys need to be configured for every zone that is split under the parent. An option to circumvent this while still retaining the advantages of the earlier setup is to split the parent also, and configure the trusted key at the level of the parent. An internal name server is configured as the authoritative server for the internal view of this split and the internal recursive forwarder is modified to forward all internal queries for the parent zone to it.

Although this option reduces the number of trusted keys at the end resolver, the trusted key still needs to change when moving between the two views. Since splitting the parent essentially creates two new zones, records in the parent that were previously common in both views would now need to be duplicated in the two split zones. The number of such records is typically not very large, but the overhead and complexity in maintaining duplicate records can still be a burden.

3.5. Multiple DS Records

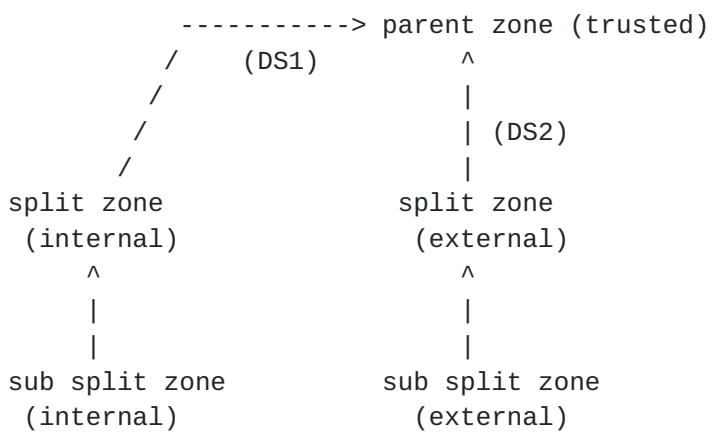

Figure 6

In this approach, the parent is not split; however, DS records corresponding to each of the two different views are published at the delegation point. The glue and NS records at the delegation point still corresponds to the external zone but this information is never used by the internal zone. As in option 4, the trusted key is configured at the level of the parent zone. The chain-of-trust from this trusted key to either zone is formed by using one of the two DS records, which ever is applicable at that view.

Since some coordination between the split zone and the parent is required to publish multiple DS records, this approach is most suitable when the split is made at a level lower than the organization apex (e.g. for example.com, the split is made at a level lower than example.com). This approach lends itself to using different keys in different views while still allowing for minimal configuration at the end resolvers; trusted keys need not be changed even if the nodes are mobile across the two views. This approach has the advantage that administrative separation of the two views of the split can be maintained while still having a single key configured at the end resolvers. Identifying which view a given record belongs to can be done by tracing the keys used to form the chain-of-trust.

While most of the internal zone contents can be kept private to the internal view, the DS record must still be exposed. Since data hiding is not the objective of the split-view setup this should not really be a problem in most cases. An attendent problem with multiple DS records is that since the validation algorithm iteratively looks for a DS record in the parent while completing the chain-of-trust there is some added computational overhead which increases as the number of DS records in the delegation point grows. In some circumstances there may not be sufficient flexibility to include all DS records in the parent, especially if the child is a part of a different organization. Lastly, the internal view is still

susceptible to an insider "attack", where data from the outside view injected in response to internal queries can corrupt the cache. This attack is common to all scenarios that use a common key for validating internal and external zone contents. Any cache pollution introduced due to administrator errors can also escape detection for the same reason.

3.6. Name Server Requirements

All name servers listed below must conform to the specifications given in [2]. Additionally, the Internal Recursive Forwarder must support the following:

- Ability to validate DNSSEC responses.
- Support for configurable DNSSEC trusted keys. It should be possible to configure more than one trusted key.

4. Packet Filtering Considerations

The following subsections define the rules that must be configured in the two packet filters depicted in Figure 1 in order to support the split-view configuration.

4.1. Inner Packet Filter

In order to allow the above configuration to work, any packet filtering system between the internal network and the boundary network must allow all of the following types of packets.

DNS queries from any internal address to the second-level recursive name server (finer level access control is done by TSIG):

Proto	SrcIP	SrcPort	DestIP	DstPort	AckBit
UDP	internal	>1023	rec.srv	53	N/A
UDP	internal	53	rec.srv	53	N/A
TCP	internal	>1023	rec.srv	53	Any

Responses to the above queries from the second-level recursive name server to any internal address:

Proto	SrcIP	SrcPort	DestIP	DstPort	AckBit
UDP	rec.srv	53	internal	>1023	N/A
UDP	rec.srv	53	internal	53	N/A
TCP	rec.srv	53	internal	>1023	Set

Queries from clients in the boundary network to any internal name server:

Proto	SrcIP	SrcPort	DestIP	DstPort	AckBit
UDP	client	>1023	internal	53	N/A
TCP	client	>1023	internal	53	Any

Responses to the above queries from the (any) recursive forwarder to clients in the boundary network:

Proto	SrcIP	SrcPort	DestIP	DstPort	AckBit
UDP	internal	53	client	>1023	N/A
TCP	internal	53	client	>1023	Set

4.2. Outer Packet Filter

Any packet filtering system configured between the boundary network and the external network needs to allow the following.

Queries from the recursive name server in the boundary network to the outside network:

Proto	SrcIP	SrcPort	DestIP	DstPort	AckBit
UDP	rec.srv	>1023	outside	53	N/A
UDP	rec.srv	53	outside	53	N/A
TCP	rec.srv	>1023	outside	53	Any

Responses to the above queries from the outside to the boundary network recursive name server:

Proto	SrcIP	SrcPort	DestIP	DstPort	AckBit
UDP	outside	53	rec.srv	>1023	N/A
UDP	outside	53	rec.srv	53	N/A
TCP	outside	53	rec.srv	>1023	Set

Queries from outside clients to the external-view authoritative servers:

Proto	SrcIP	SrcPort	DestIP	DstPort	AckBit
UDP	outside	>1023	auth.serv(ext view)	53	N/A
UDP	outside	53	auth.serv(ext view)	53	N/A
TCP	outside	>1023	auth.serv(ext view)	53	Any

Responses to the above queries from the external-view authoritative server to the outside:

Proto	SrcIP	SrcPort	DestIP	DstPort	AckBit
UDP	auth.serv(ext view)	53	outside	>1023	N/A
UDP	auth.serv(ext view)	53	outside	53	N/A
TCP	auth.serv(ext view)	53	outside	>1023	Set

Note that in this configuration, queries from all recursive name servers in the boundary network for any external view information may need to transit outward through the second-level packet filter and then back again into the boundary network. If the existing packet filter policy prevents such traffic patterns, all such recursive name servers would need additional forwarding statements to forward these queries directly to their respective authoritative name servers without going through the packet filter.

5. Summary

This document describes an approach for configuring split-view DNSSEC. The approach uses a two-level recursive scheme where an internal recursive forwarder resolves inside answers and marshalls all outside queries to a second-level recursive name server. TSIG between the internal and the second-level name servers protects against errant queries.

The recommended configuration has been shown to be adjustable for various needs and security consideration levels. Differences in these approaches make trade-offs between configuration overhead and validation overhead. Trading-off in favour of minimal operator overhead is useful for overall maintainability of the system, especially when split-view DNS is considered in the context of nodes that are mobile across the two views.

Although split-view DNSSEC is possible using the recommended setup,

it still involves significant effort: for configuring the various name servers, for setting up zone forwarding, for configuring and distributing shared keys for TSIG, and, depending on the configuration, for performing DS (or keyset) exchanges for every view of a split zone. Some configurations may also require multiple trusted keys in end resolvers which may change between views. Proper care must be taken to ensure that correct split-view behavior is consistently maintained.

6. IANA Considerations

This document has no actions for IANA.

7. Security Considerations

The configuration suggested in this document tries to minimize cache pollution, but misconfigurations are still easily possible. Any misconfigurations in the three different types of name servers or the two packet filters can either result in cache pollution and cause incorrect results to be returned, or impede the ability for end resolvers to validate data returned in response to queries. An improperly configured packet filter that allows errant DNS traffic through or denies legitimate responses can lead to aggressive retransmission of queries.

Each of the validation options outlined in Section 3 also introduce their own security considerations. Using a common key between both views of the split does not allow one to differentiate between internal and external data and troubleshooting is greatly encumbered. All approaches that use a common key for validating internal and external data are also susceptible to an insider attack where data from the outside view injected in response to internal queries can corrupt the cache. On the other hand, using a multitude of keys at end resolvers only increases the operator overhead and thus the chances for configuration errors.

8. Acknowledgements

The contributions, suggestions and remarks of the following persons to this draft are particularly acknowledged: Wesley Griffin, John Kelley, Russ Mundy and Sam Weiler. The two-level name server scheme described in this document builds upon work that was originally performed by Ed Lewis.

9. References

9.1. Normative References

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

[2] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Protocol Modifications for the DNS Security Extensions", RFC 4035, March 2005.

9.2. Informative References

[3] Larson, M. and P. Barber, "Observed DNS Resolution Misbehavior", Work in Progress ietf-dnsop-bad-dns-res, July 2005.

Author's Address

Suresh Krishnaswamy
SPARTA Inc.
7075 Samuel Morse Dr.
Columbia, MD 21046
US

Email: suresh@tislabs.com
URI: http://www.sparta.com

Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at http://www.ietf.org/ipr.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Disclaimer of Validity

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Copyright Statement

Copyright (C) The Internet Society (2005). This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.