ISC SOFTWARE GUILD

BIND Forum Newsletter

DNS Lookaside Validation

Issue 2

May 2006

Inside this issue:

For a long time, ISC has been supporting the development of the DNS Security extensions by tracking and contributing to the standards development process as well as providing a reference implementation as part of the BIND code.

In that line, this article describes the process called Domain Lookaside Validation (DLV) as implemented in BIND 9.3 and later, and the additions to the DNSSECbis validation model that it introduces, designed to be an aid in the early operational deployment of the DNS Security extensions (DNSSECbis).

The article also describes the operation of the infrastructure necessary to make DLV operational on the Internet.

Funding for this development work was provided by Keio University.

Review of the DNSSEC validation process

DNSSECbis defines several new DNS resource records (RRs) that enable cryptographic validation of DNS information. One of the most crucial new resource record types is the DS (Delegation Signer) RR.

This record is generated by a parent zone administrator from DNSKEY RRs provided by the child zone and asserts the fact that the keys used by the child are indeed provided by the child zone administrator and hence that the DNS delegation tree is being followed correctly.

This record occurs at zone cuts and differs from most other records at this point in that the authoritative RR resides at the parent zone and not at the child zone's apex, unlike for instance the NS RR set.

Ideally, for this verification process to be automatic and easy to operate on the Internet, the root zone would be signed and it would sign the DS records of the TLDs, and these would in turn sign their child zones, and so forth.

In practice, getting the root zone to be signed is proving to be a slow process, not only because of technical constraints, such as the not-quite-solved problem of key rollover, the process by which a key that is in use is replaced by a new one, but also because of political issues such as that of key ownership. Given that this key would be configured in all resolvers that wish to make use of DNSSEC, this distribution process is not a small problem and it needs to be stable and well defined.

In the absence of a signed root zone, it would still be manageable to a certain extent to have a collection of keys for each of the TLDs, though the key management concerns above

also apply here and they would be somewhat of a burden on system administrators and a potential source of accidents.

Having a greater number of hand-configured keys would be too costly.

In order to help kick-start the operational deployment of DNSSECbis on the Internet, an addition to the DNSSEC validation algorithm, DLV, is being introduced as part of the BIND 9.3 and 9.4.0 servers and resolver code.

Domain Lookaside Validation (DLV)

DLV is a local policy mechanism that changes a validator's behaviour in the event that no DNSSECbis metadata is available, or if the available metadata makes use of a key that does not exist in the parent zone. It is not a protocol extension per se, and is not the result of the IETF standards process. Rather, it is a local extension used by co-operating zone administrators and validator operators who wish to publish and consume keys before the normal DNSSECbis mechanisms are usable due to an unsecured parent zone. DLV is expected to become irrelevant, and to die, after the root zone and most TLD zones are secured.

By definition, if a DNSSECbis validator would have found a secure result without DLV, then this result is unchanged by the DLV logic. However, if DNSSECbis would have resulted in an insecure result, then DLV processing can produce a secure result if the zone administrator has registered their keys in a DLV namespace subscribed to by that DLV-capable validator.

How DLV is implemented - DLV RRs

A new DLV RR is used, which is structurally identical to the DS RR but has none of the normal DS RR metadata semantics - in other words it is never returned automatically as a side effect of queries for other RR types. The type code number for the DLV RR is in "specification required space" and has been allocated by the Internet Assigned Numbers Authority (IANA).

DLV RRs are stored in one or more secured zones in normal DNS (for example, DLV.ISC.ORG). The name of the DNSKEY RR described by a DLV RR is just the owner name of the DLV RR, minus the name of the DLV namespace. So a DLV RR with owner name VIX.COM.DLV.ISC.ORG stored in DLV namespace DLV.ISC.ORG would describe a DNSKEY RR whose owner name is VIX.COM.

A DLV namespace can consist of more than one zone, if zone cuts are necessary to create interior administrative boundaries. For example, if ten million registrations all occurred for child zones of COM, the DLV namespace operator might place an interior zone cut at COM.DLV.ISC.ORG in order to balance the resulting nameserver load.

Given that registration in a DLV namespace is optional and is meant to support DNSSEC deployment, all zones in a DLV namespace must be secure, and the apex DNSKEY must be trusted by all DLV validators. Therefore, interior zone cuts in the DLV namespace will have normal DS and DNSKEY RRs, and all RRsets in a DLV namespace will have RRSIGs.

Using DLV

A DLV-capable validator will be configured to know the name of one or more DLV namespaces, and to know the trusted key used to sign each DLV namespace. When DLV has been enabled in a validator, several exceptions that would normally yield non-secure results are handled specially.

Since a DNSSECbis validator always signals its desire to receive DNSSECbis metadata when it sends or forwards a query, a response that lacks such metadata is usually treated as either evidence of an insecure zone (if the parent zone's delegation did not include a DS RRset), or as evidence of an attack or misconfiguration (if a DS RRset indicated that the zone was secure). Because of "the grandparent problem" where a name server who is authoritative for both a zone and its ancestor (for example, LAB.ISC.ORG and ORG), every DNSSECbis validator has to do some extra work to distinguish these cases.

When using DLV, a validator who has determined that there really is no matching parent DS RR for a zone, will perform a lookup for the corresponding DLV RR in the locally configured DLV namespace that best matches the zone's apex. Therefore if a validator knows a DLV namespace for the root (.) zone and another DLV namespace for MIL, missing DS RRs for the NAVY.MIL zone will cause lookups in the known MIL DLV namespace, whereas missing DS RRs anywhere else will cause lookups in the known root (.) DLV namespace.

Last, in order to keep the number of additional queries necessary when using DLV to a minimum, the caching algorithm of the resolvers is altered to exercise a more aggressive negative caching.

Deploying DLV

DLV operation requires the existence of a DLV registry, whose job is to accept DLV RRs, as a zone operator would accept DS RRs for its child zones, verifying the legitimacy of the registrations. These records must be published in a secure zone with high availability.

The DLV registry should be a public benefit corporation with strong ties to the research and protocol development communities, such that deployment statistics will be tracked and publicly disclosed, and query names seen at the DLV namespace's nameservers will not be used commercially, and the registry function can be altered or ended according to the needs of the DNSSECbis community.

Zone administrators who generate new DNSKEY key pairs and associated DS RRs will have to form a relationship to the DLV Registry and submit DLV RRs to the DLV Registry for publication. This is analogous to submitting DS RRs to a parent zone registrar (for example, in the normal DNSSECbis data model, when a new DNSKEY is created for VIX.COM, the corresponding DS RR would be sent to a COM registrar for ultimate inclusion in the COM zone.

Note that after a zone's parent zone is secured, and if the parent zone's DNSKEY is submitted to the DLV Registry, then the zone's own DNSKEY will no longer need to be published in the DLV Registry. Also note that once all ancestor zones from a given zone up to the root (.) zone are secured, then there is no need to continue publishing keys in the DLV Registry.

Finally, operators of validating full resolvers and caching forwarders will have to install software having DLV capabilities, and enable those capabilities, and configure one or more

trusted keys for the chosen DLV namespace. It is also advisable to monitor the DLV Registry to become aware of any changes to the DLV technology or to the configured trusted keys for the DLZ zone. (The same can be said for changes to the DNSSECbis technology or to the root (.) zone keys or to methods of automatically rolling over the root (.) zone keys.)

ISC is running a DLV registry for use by anyone, free of charge, as a means to bootstrap deployment of DNSSEC in the real world. As part of this service, ISC is working together with Domain Registrar's so domain holders can use already established relationships to provide their DNSSEC information for use in the DLV tree. See http://www.isc.org/ops/dlv for more information.

Minutes of third BIND Forum member's meeting

The third BIND Forum members' meeting took place on Sunday 31 July, 2005 at the Palais de Congres in Paris, France, next to IETF 63.

The agenda for the meeting was as follows.

- Agenda bashing
- Review of current releases
- Review of new features for BIND 9.4
- Plans for BIND 9.5
- Enabling more and better communication between ISC and membership
- DNSSEC deployment (DLV)
- Deprecating BIND 8 as a caching resolver (Presentation by a member: JPRS)

Review of current release

- BIND 9.3.x: A review of new features as posted on ISC's website was presented.
- BIND 8: only bug fixes. Called attention to the use of BIND 8 in a forwarder configuration, as per recent security advisory (See http://www.isc.org)

Review of new features for BIND 9.4

New documentation

ISC has improved reference documentation generation from the code itself (doxygen)

There is also a better generation system. It is now easier to regenerate documentation and to produce more formats both by ISC and users.

ISC is proceeding with review of the ARM's content (initial phase only)

This work will not be complete by the time of 9.4.0 release but in some later minor release.

DLZ

Integration of the generic API and non-database specific parts has been done.

Drivers that are specific to each database backend (e.g. bdb, mysql, postgresql, etc) will be shipped as contributed code with each BIND 9 release but not developed by ISC itself.

Caveats of DLZ include:

- does not support incoming axfr
- does not support all RR types
- does not support dynamic updates

This feature is turned on/off a per zone basis via the configuration file and must be compiled in using configure script options

Queries are answered via the DB back-end and a memory tree is not built inside the server for these zones

Provides instant visibility for any changes done to the database

gss tsig

This feature has been promised earlier but not delivered.

Reasons for this are non-conformity of the main target platform OS (Windows server) with the standard available in ways that have been impossible to work around so far.

ISC has working code for MIT and Heimdal Kerberos.

ISC has a couple of possible avenues to explore to get interoperability with MS

- When running on Windows platforms, use the native library as the API is almost identical and would possibly only require linking against the native library
- Trying to get access to contributed code that implements the necessary interoperability functionality

It is possible that full inter-operability will not be available in the code at the time of 9.4.0 release, but rather in one of the later point releases.

New RR types

Support for IPSECKEY and SPF RR types is being introduced in 9.4.0

Performance

BIND 9.4.0 will include some significant performance improvements:

- as an authoritative server in various configurations
- as a caching server
- zone loading when loading from binary zone representations

Plans for BIND 9.5

New resolver API

New requirements are being placed on the resolver API

- more need for abstraction.
- enabling various models of applications (event driven, multithreaded, etc.)
- ability to link against only subsets of the protocol as required by each specific application
- better communication with calling application with regards to outcome of DNSSEC validation processes, enabling the application to decide how to handle various failure modes.

XML configuration and process communication, internal architecture a more long term thought, not well defined currently.

Enabling more and better communication between ISC and membership

ISC renewed its constant offering of constant availability.

Proposal for more focused mailing lists, for instance.

A member commented on the absence of the promised quarterly newsletter, as a way to maintain a more continuous interaction. ISC committed to fixing this issue.

Other comments indicate a clear desire to have more means for ISC to report periodically to membership.

DNSSEC deployment (DLV)

A summary description of Domain Lookaside Validation was given. Open issues that have been heard about this feature were brought up, namely:

- it is not clear whether verification should look for the longest or the shortest label first (different members expressed different opinions)
- there were concerns about ISC running the registry as well as producing the code
- there were divided opinions about whether BIND should implement DLV at all or support more than one DLV registry
- Members expressed concerns regarding whether when searching the DLV tree the match with higher priority should be the one closer to the TLD (e.g. com.dlv.isc.org) or the more specific one (e.g. vix.com.dlv.isc.org). TLD operators were in favour of the first approach whereas other members were in favour of the second approach.
- There were several requests for a published specification that would enable members to better evaluate the consequences to their businesses of this feature
- There were expressions of concern on whether this is the way ISC was spending BIND Forum membership fees as the majority of members thought this was not a requested or welcome feature.

Deprecating BIND 8 as a caching resolver (Presentation by a member: JPRS)

JPRS has been suffering from certain incorrect behaviour of BIND 8 when used as a caching resolver, due to lack of EDNS(0), retrying policy and behaviour when coming across delegations where no glue is provided, resulting in performance impacts on JPRS's servers (TLD for .jp). JPRS issued a request to consider having ISC marking BIND 8 as deprecated so they could have some argument towards Japanese ISPs when asking them to migrate out of BIND 8.

Some members expressed concern about this, as BIND 8 is much faster as a caching resolver than BIND 9 currently is and it serves well for a large number of servers.

Proposals that were put forward for consideration were:

- Stop shipping binaries for BIND 8
- Move the BIND 8 releases on the ftp site to DEPRECATED or OLD directory
- Stop maintenance except for security vulnerabilities.

The proposal that got a more significant support was that for stopping the Windows binary distribution.

Commentary on events since the meeting

Since last year's BIND Forum members' meeting ISC has been hard at work developing BIND 9.4. Following are brief notes of the current state of affairs.

First, BIND 9.4 is now in public alpha release and we expect to have the first beta before the 2006 Member's meeting (see below). It has been harder than expected to get all the desired features implemented and in working condition.

Work has progressed consistently towards an implementation of gss-tsig and though this feature will not be part of the 9.4 release in order not to delay it further, ISC now has a beta version of a gss-tsig implementation that will interoperate with various versions of MS Windows servers. This feature will become available in BIND 9.5, which will be made available in a much shorter time frame than the gap between the 9.3 and 9.4 releases.

Performance is still the big highlight of the 9.4 release and ISC has now ironed out the quirks in the code that have been making it less stable than necessary for a release.

ISC has supported other developments in the DNS by implementing initially workshop-ready versions of DNS protocol extensions to enable real life testing and evolution of the protocol.

This is a fundamental part of BIND's role as a reference implementaion of the DNS protocol.

Joao Damas

In the next Issue

The next issue of the BIND Forum Newsletter will cover with detail the performance improvements of BIND 9.4, and which areas benefit most from it. We also be covering DNS anycast operations and how to use BIND to successfully deploy that technique.

Next BIND Forum Member meeting

The next BIND Forum annual member meeting will take place on the evening of Sunday July 9th in Montreal, to coincide with the 66th IETF meeting.
We will be reviewing BIND's roadmap and recent developments.
Details of the venue and time will be sent shortly to the forum_members@isc.org mailing list.