NLNETLABS

ANNUAL REPORT 2021

Table of Contents

About NLnet Labs

NLnet Labs is a not-for-profit foundation, founded in 1999. Over the past 20 years our mission has been to develop open source software and open standards for the benefit of the Internet, and to perform applied research on Internet protocols. We focus our efforts specifically on the Domain Name System and inter-domain routing. NLnet Labs' work supports the robustness, security and reliability of the Internet and safeguards the privacy of its users.

To accomplish our mission we collaborate with key Internet players around the world. Organisations we work with include the Internet Engineering Task Force (IETF), the Regional Internet Registries (RIRs), the Internet Corporation for Assigned Names and Numbers (ICANN), leading Top Level Domain (TLD) operators, the International Standards Organisation (ISO), the Internet Society (ISOC), and a wide variety of others in the field, ranging from individual researchers to major industry actors.

NLnet Labs plays a leading role in promoting technologies that stimulate trust, security, privacy, scalability and the global nature of the Internet. Our peers see us as a major stakeholder in the creation and use of open standards and open software. We are leading experts on core internet technologies, specifically DNS and routing.

We are a lean organisation with a team of around 15 people, consisting almost exclusively of developers and researchers, with minimal overhead. We attract talented people who want to make a difference in the well-being of the Internet, with a profound belief in open source and open standards.

We develop open-source software that is used across the Internet industry, ranging from the DNS root servers at the core of the Internet to small embedded devices running a secure recursive resolver, and routing security software that helps protect the network of large network operators.

Our researchers pioneer new technologies, help define future standards and build prototypes of technologies that promise to improve the Internet. We increase understanding of the Internet by studying its fundamental building blocks. By actively participating in both worlds – development and research – we bridge the gap between academia and industry, and introduce solutions that are practical as well as innovative.

We also contribute to policy and governance organisations. Our technical expertise and advice is widely recognised by policy-making bodies. We advise on public policy decisions that affect the security and privacy of Internet users across the globe, as well as the stability of the Internet itself.

Software Development

At a glance

In 2021 we continued to develop and extend our existing DNS and RPKI software. For the DNS products we have published several releases of Unbound, NSD and OpenDNSSEC. Next to bug and security fixes and stability improvements, we have been working to implement Extended DNS Errors, ZONEMD Zone Verification and support for SVCB/HTTPS records. Unbound also now supports the full set of RPZ triggers and actions. The documentation for Unbound and NSD was moved to the Read the Docs platform. A similar move for OpenDNSSEC is planned for next year.

Our routing security software continued to mature and evolve. This year we extended our portfolio from pure RPKI to include more generic BGP tooling. All in all, we published more than a dozen releases of Routinator, Krill and RTRTR. New is JDR, a service to explore, inspect and troubleshoot the RPKI.

This year we received substantial financial and infrastructural support for our RPKI projects from internet organisations as well as industry partners, allowing us to put in a lot of work to further develop the software. We discuss these additions in more detail below.

In 2021, we incorporated a long list of bug fixes in our DNS library LDNS. Even though we now limit our work on LDNS mainly to maintenance, we took the opportunity to implement support for ZONEMD and the SVCB and HTTPS record types.

We made several bug fixes and some improvements to the user experience for the getdns library and Stubby resolver.

We also did some limited refactoring and made a few bug fixes for Domain Crate, our DNS library for the Rust programming language, and implemented support for Extended DNS Errors.

The rest of this section gives an overview of our software projects and discusses this year's achievements in more detail.

DNS(SEC) Software Projects

Unbound

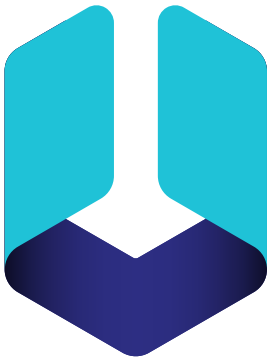

In 2021 we published two significant releases of Unbound (1.13.1 and 1.14.0) and several maintenance updates (including bug fixes). New features include support for Extended DNS Errors, ZONEMD Zone Verification, support for SVCB/HTTPS records, and the implementation of remaining RPZ triggers/actions.

Extended DNS Errors are specified in RFC 8914. Although primarily created to extend SERVFAIL to provide additional information about the cause of DNS and DNSSEC failures, this mechanism allows all response types to contain extended error information.

ZONEMD is defined in RFC 8976 and allows for a new record type containing a cryptographic message digest over the full zone. The ZONEMD record itself is protected by a DNSSEC signature. Since the record is part of the zone file, the message digest can be used by a recipient to verify that the zone is correct and complete.

Unbound now checks the ZONEMD records of zones loaded as auth-zone, also validating DNSSEC if available.

An HTTPS record allows you to specify full information about a specific HTTPS service (typically a website). It serves as an alias, its main contents being the target address, the HTTP versions supported, and optionally a set of IPv4 and/or IPv6 addresses (hints). The SVCB record is a generalised version of the HTTPS record, to be used with a service name instead of a host name. These relatively new record types solve various problems with the traditional CNAME and DNAME aliases.

Response Policy Zones (RPZs) allows network operators to include response policies in their recursive resolvers. That way, requests from their networks to the global Internet can be blocked or redirected (a "DNS firewall"), for example to prevent users from ending up at a known malicious site. Since RPZ policies are specified in zone file format, they can be maintained and distributed using existing zone transfer mechanisms.

With the addition of nsdname, nsip and clientip, Unbound now supports the full set of RPZ triggers. The same is true for RPZ actions, now also including the tcp-only action.

The Unbound documentation has been moved to the Read the Docs platform and is available here: https://unbound.docs.nlnetlabs.nl/en/latest/.

About Unbound: Unbound is a DNSSEC-validating, recursive, caching DNS resolver. It is designed to be fast and lean, and incorporates modern features based on open standards. The software runs on FreeBSD, OpenBSD, NetBSD, MacOS, Linux and Microsoft Windows, with pre-built packages available for most platforms. It is included in the standard repositories of most Linux distributions. Installation and configuration are designed to be easy: just a few lines of configuration is enough to set up a resolver for your machine or network.

NSD

In a series of minor updates (versions 4.3.5-4.3.9), we carried out various maintenance tasks (including several bug fixes) to ensure that NSD remains a dependable and performant authoritative name server. Along the way, we also added several new features.

NSD now supports Extended DNS Errors as specified in RFC 8914. Although primarily created to extend SERVFAIL to provide additional information about the cause of DNS and DNSSEC failures, this mechanism allows all response types to contain extended error information.

Also added was support for the SVCB and HTTPS record types. (we explained the mechanism and value of SVCB and HTTPS in the Unbound section above).

Finally, NSD now supports XoT, allowing zone transfers (AXFR/IXFR) over TLS, thereby protecting the security and privacy of the transfer.

Just like we did for Unbound, we have moved the documentation of NSD to Read the Docs: https://nsd.docs.nlnetlabs.nl/en/latest/.

About NSD: Name Server Daemon (NSD) is an authoritative DNS name server. It has been developed for operations in environments where speed, reliability, stability and security are essential. The software is designed with a pure philosophy that prioritises raw performance. This means that if you serve hundreds of thousands or even millions of queries per second, NSD is the world's leading name server. This makes it ideal for Top Level Domain implementations, DNS root servers and anyone in need of a fast and optimised authoritative name server. Currently, three DNS root servers and many top-level domain registries use NSD as part of their server implementation. NSD also strives to be a reference implementation for emerging standards of the IETF.

OpenDNSSEC

After ending support for version 1.4 of OpenDNSSEC in 2019, this year we saw continued upgrades and deployments of version 2.1 by large DNS operators – including TLD operators – who depend on a fully managed DNSSEC signing solution.

In addition to the existing deb package, the community also made an rpm package file available, which makes it a lot easier to install OpenDNSSEC on Red Hat based Linux platforms.

After putting considerable effort into outreach over the last two years, cooperation with the user community has become intensive and is running smoothly. Operators ask us for help with their upgrades and deployments, and provide feedback to further improve OpenDNSSEC. Internetstiftelsen (.se) and SIDN in particular have been enthusiastic users and supporters of OpenDNSSEC.

The 2.1.9 and 2.1.10 minor releases provided further stability improvements (including bug fixes), and in particular more robust PKCS#11 interactions with various HSM setups.

About OpenDNSSEC: OpenDNSSEC is a policy-based zone signer that automates DNSSEC key management and the signing of zones. The main goal of the project is to make the Domain Name System Security Extensions (DNSSEC) easy to deploy, thereby driving the adoption of DNSSEC and enhancing Internet security.

SoftHSM

The SoftHSM project, to which NLnet Labs contributed in the past, was incorporated as a project under the Commons Conservancy in 2019. The long-term goal of this step was to keep the project sustainable and allow new partners to make significant contributions. The last release of the software, however, was version 2.6.1, published in 2020. We will keep maintaining the project, but we have not developed any new activities related to SoftHSM.

About SoftHSM: SoftHSM was developed to provide a software-based solution for people who wish to use OpenDNSSEC but are not willing or able to invest in a new cryptographic hardware device. It provides a software implementation of a generic HSM with a PKCS#11 interface. SoftHSM has been designed to connect directly to OpenDNSSEC, but thanks to its standard PKCS#11 interface can also be used by other cryptographic products.

DNS(SEC) Libraries

LDNS

In 2021, LDNS saw two minor releases (versions 1.8.0 and 1.8.1), which incorporated a long list of bug fixes, some of which were contributed by the user community. But we also took the opportunity to implement new functionality: support for ZONEMD was added to ldns-signzone and ldns-verify-zone. We also made a draft implementation for the SVCB and HTTPS record types (we explained the mechanisms and value of ZONEMD, SVCB and HTTPS in the Unbound section above).

We will continue to maintain LDNS, with no plans for major changes in the near future.

About LDNS: LDNS is a C library to simplify DNS programming. It supports all low-level DNS and DNSSEC operations. It also defines a higher-level API, which allows a programmer to quickly create or sign packets, for example. Developers can use LDNS to easily create RFC-compliant software and build proof-of-concepts for various Internet Drafts.

We do not strive for LDNS to be a comprehensive library that supports every (emerging) standard. The software includes a DNS lookup utility named drill (an alternative implementation to BIND's dig). As drill has nothing in common with either NSD or BIND, it allows for debugging and testing using an independent code base.

getdns and Stubby

In 2021, we released version 1.7.0 of the getdns library and version 0.4.0 of the Stubby resolver. These packages included several bug fixes and some improvements to the user experience.

About getdns: getdns is a modern asynchronous DNS API and library. It implements DNS entry points from an interface design developed and vetted by application developers, which was consolidated in an API specification. This implementation is developed and maintained in collaboration between NLnet Labs, Sinodun and No Mountain Software. Although the code is written in the C programming language, bindings for several other languages are available. The software is published under the New BSD License.

About Stubby: Stubby is a local DNS Privacy stub resolver. It is built on the getdns library and is available for UNIX-like systems as well as Windows (the latter as a binary). Stubby uses DNS-over-TLS (DoT) to encrypt DNS traffic sent from a client machine (typically a desktop or laptop) to a DNS Privacy recursive resolver service, thereby improving end-user privacy.

Net::DNS(::SEC)

2021 saw a series of minor releases of Net::DNS(::SEC) (versions 1.30-1.33). The updates provide several improvements in functionality, plus a few bug fixes in the code and documentation.

About Net::DNS(::SEC): NLnet Labs is a long time contributor to and maintainer of Net::DNS(::SEC), a DNS library written in the Perl scripting language. It consists of the Net::DNS resolver and the Net::DNS::SEC addon. The latter adds DNSSEC support to Net::DNS. Net::DNS::SEC must be downloaded as a separate package from CPAN, because the two components may have mutually incompatible dependencies.

Domain Crate

In 2021 we released version 0.6 of Domain Crate. After completing a major overhaul of the structure and refactoring the entire library last year (in version 0.5), the changes in this release are more limited. We moved/merged some functionality from one crate into another, and fixed some bugs. We also implemented some new functionality: support for Extended DNS Errors (as described in the NSD section above), and traits for generic type conversions of octets between octets sequences.

About Domain Crate: Domain Crate is a DNS library written in the Rust programing language. It contains an ever-growing set of building blocks for including DNS functionality in applications. These blocks currently include the basic data structures and functionality for creating and parsing DNS data and messages, support for signing and verifying messages using the TSIG mechanism, experimental support for reading data from DNS master files (also known as zone files), experimental and as yet incomplete support for DNSSEC signing and validation, and a simple Tokio-based stub resolver.

Routing Software

In 2021, the development of our routing software portfolio continued mature and evolve. The year was especially significant because it marked the first time our portfolio went beyond projects aimed at security through Resource Public Key Infrastructure (RPKI); from here on, it also introduced more generic BGP tooling. Most notably, the routecore project offers BGP routing building blocks, while the rotonda project features a modular, analytical BGP Engine.

In the area of RPKI, our Relying Party software Routinator 3000 and Certificate Authority software Krill each had 16 releases and the RTRTR data proxy continued to evolve. This year also saw the introduction of JDR, a service to explore, inspect and troubleshoot RPKI.

Project funding

One of the reasons we could develop our RPKI toolset with full force is because several organisations in the industry decided to support us, either financially or with infrastructure. The National Internet Registry of Brazil, NIC.br, pledged to support the development of Krill and Routinator for two years, enabling us to dedicate full-time staff towork on the toolset.

APNIC (the Asia Pacific Network Information Centre), the regional Internet address registry (RIR) for the Asia-Pacific region, also supported the continued development of our RPKI toolset, funding the development of Hardware Security Module (HSM) support for Krill.

Additional income came from several organisations, including Internet Service Providers, Internet Exchanges, Tier 1 Carriers and cloud providers purchasing support services.

Furthermore, DigitalOcean, Fastly and Amazon Web Services provided us with their services free of charge so we could set up an automated test platform for the software, host analysis tools, and make our production platform as resilient as possible.

Routinator

Routinator is Relying Party software, also known as an RPKI Validator. Operators can use it to download and validate the global RPKI dataset and feed the result into their routers, or use it elsewhere in the BGP decision making process.

Routinator connects to the Trust Anchors of the five Regional Internet Registries (RIRs) (APNIC, AFRINIC, ARIN, LACNIC and RIPE NCC), downloads all of the cryptographic material in their repositories and validates the signatures. It can feed the validated information to hardware routers supporting Route Origin Validation, such as Juniper, Cisco and Nokia, as well as serving software solutions like BIRD and OpenBGPD. Alternatively, Routinator can output the validated data in a number of useful formats, such as CSV, JSON and RPSL.

In 2021, the first major release was 0.9.0. This keeps the last valid data from a publication point and falls back to using that if an update to the publication point does not have a valid manifest or the data does not match the manifest. This data is stored in a key-value database rather than

directly in the file system. In addition, if an RRDP repository is unavailable for a certain time, Routinator will fall back to rsync. The time since the last successful update before this fallback happens is randomly chosen for each repository between the refresh time and a user-configurable upper limit.

Also in this release, the size of downloaded RPKI objects is limited to 20 Mbytes by default; this applies to both RRDP and rsync. Routinator includes additional TALs for various commonly used RPKI testbeds. Prometheus metrics and JSON status have been greatly extended, with more detailed counters for individual valid and invalid object types. They are also now available on a per-repository basis in addition to the existing per-TAL basis. Prometheus metrics and JSON status can now optionally include per-client RTR metrics, though this is disabled by default to avoid accidentally leaking information about the local network topology.

The major rearchitecting of the data storage in version 0.9.0 caused memory usage to grow to an unacceptable level. As a result, in version 0.10.0 data is stored directly in the file system again. This release maintains all the previous improvements to robustness, while its memory usage is lower than ever. This release also includes a metadata object in the json and jsonext output formats specifying the time the dataset was created. The maximum number of over delta steps performed when updating an RRDP repository is limited to 100 by default.

The last release of 2021 fixes several vulnerabilities found by Koen van Hove, a researcher at the University of Twente. Specifically, Routinator now limits the maximum length an RRDP request can take; this prevents a possible issue where an RRDP repository maliciously or erroneously delays a request and subsequently a validation run. A new configuration setting limits the length of a chain of CAs from a trust anchor, fixing a possible vulnerability where a CA creates an infinite chain of CAs. Lastly, support for the gzip transfer encoding for RRDP has been removed because gzip in combination with XML provides multiple ways to delay validation.

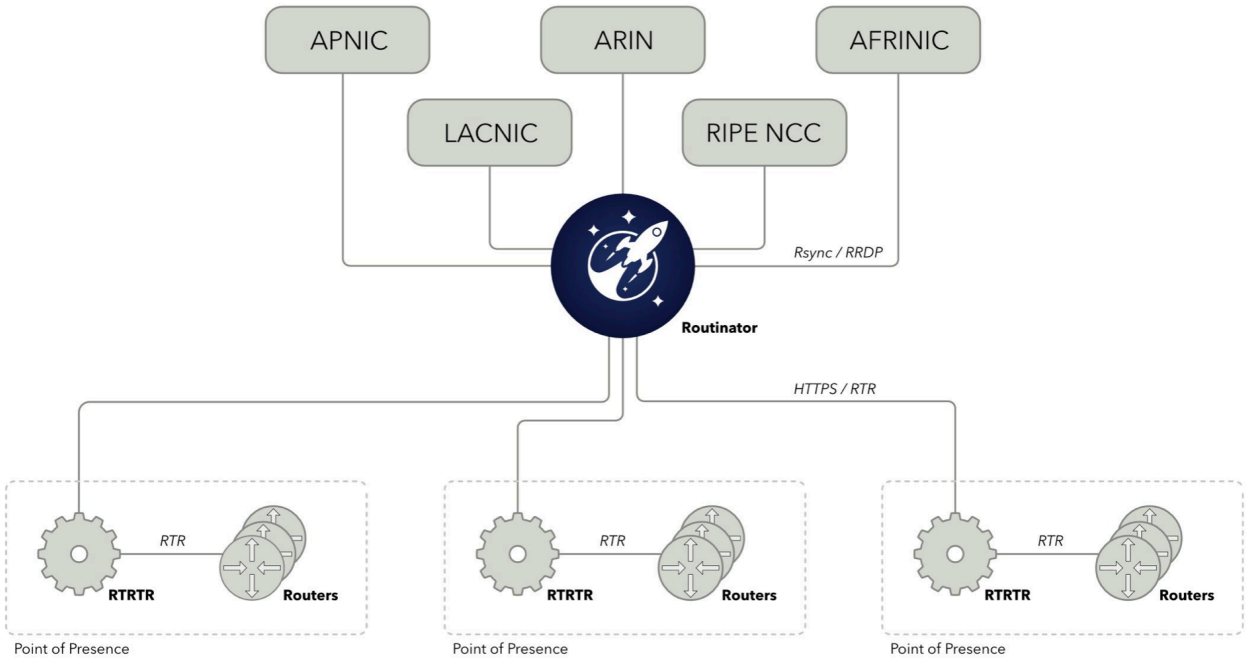

RTRTR

RTRTR is an RPKI data proxy designed to collect Validated ROA Payloads from one or more sources in multiple formats and dispatch it onwards. It provides the means to implement multiple distribution architectures for RPKI, such as centralised RPKI validators that dispatch data to local caching RTR servers.

RTRTR can read RPKI data from multiple RPKI Relying Party packages via RTR and JSON, and in turn provide an RTR service for routers to connect to. The HTTP server provides the validated data set in JSON format, as well as a monitoring endpoint in plain text and Prometheus format.

In 2021, the "json" unit gained support for the modified JSON format used by newer versions of rpki-client. We also added a "slurm" unit that can be used to manipulate payload sets based on local exception files defined in RFC 8416. Lastly, a new "rtr-tls" unit and target can send RTR data over TLS connections.

Krill

With Krill, operators can generate and publish RPKI cryptographic material to authorise their BGP announcements. Until recently, operators were largely dependent on the hosted RPKI systems that each of the five Regional Internet Registries (RIRs) provide. Krill lets organisations run RPKI on their own systems as a child of one or more RIRs. It can also run under a different parent, such as a National Internet Registry (NIR), and in turn act as a parent for other organisations.

The implementation will support running the CA both upwards and downwards. Upwards means that operators can have multiple parents, such as ARIN, RIPE NCC, etc., simultaneously and transparently. Downwards means that the CA can delegate to child organisations or customers who in turn run their own CA. This makes Krill ideal for National Internet Registries and Enterprises.

A publication server is included in Krill, but can also be run as an independent component. This means organisations can host published certificates and ROAs themselves, or let a third party such as a Content Delivery Network do it on their behalf.

Krill is intended for:

- organisations who do not want to rely on the web interface of the hosted systems that the RIRs offer, but require RPKI management that is integrated with their own systems;
- organisations who need to be able to delegate RPKI to their customers or different business units, so that that they can run their own CA and manage ROAs themselves;
- organisations who manage address space from multiple RIRs; using Krill, they can manage all ROAs for all resources seamlessly within one system;
- organisations who want to be operationally independent from their parent RIR, such as an NIR or an enterprise.

In 2021, the first major release (0.9) had many changes under the hood compared to the previous version. The updates included multi-user support in the user interface for both local users and OpenID Connect, reduced disk space usage and growth over time, consistency in API and naming, and improvements to the Publication Server.

With Krill's multi-user support it becomes possible to give people in your organisation individual access rights to your CA, without needing to share a password. If you have an OpenID Connect provider you can integrate Krill with it.

In subsequent smaller releases, Krill gained the ability to warn about ROA configurations for resources no longer held, improved reporting of I/O errors, and improved status reporting and monitoring. Lastly, experimental support for Autonomous System Provider Authorization (ASPA) objects was included. ASPA lets operators verify the AS_PATH attribute of routes advertised in the Border Gateway Protocol.

JDR: Explore, Inspect and Troubleshoot RPKI

Working with RPKI can be quite complex. Implementing Relying Party (RP) or Certificate Authority (CA) software requires knowledge and understanding of a significant number of RFCs. The end-user deploying and running such software, is normally spared this deep-dive into the land of standards. That is, as long as everything works as expected.

Once things do not work as expected, finding the cause can be challenging, as there are so many (moving) parts involved. The RPKI is a distributed repository with possible delegations, containing objects created with different pieces of software, transported via one of several ways, to be interpreted by yet again a plethora of libraries and software. And while most software will try to offer concise logging to the user in case of any unexpected situation or error, the focus of these packages is often not the troubleshooting part.

This is where JDR comes in. Just like RP software, JDR interprets certificates and signed objects in the RPKI, but instead of producing a set of Verified ROA Payloads (VRPs) to be fed to a router, it annotates everything that could somehow cause trouble. It will go out of its way to try to decode and parse objects: even if a file is clearly violating the standards and should be rejected by RP software, JDR will try to process it and give the end-user as much troubleshooting information as possible.

In 2021 we began ongoing work to implement time-related validity checks ("Is this object already valid, or perhaps past its due date?") and revocation checks ("Is this certificate listed on a valid CRL?"), as well as support for the RPKI Repository Delta Protocol (RRDP). The development of JDR was partly funded by the RIPE NCC Community Projects Fund.

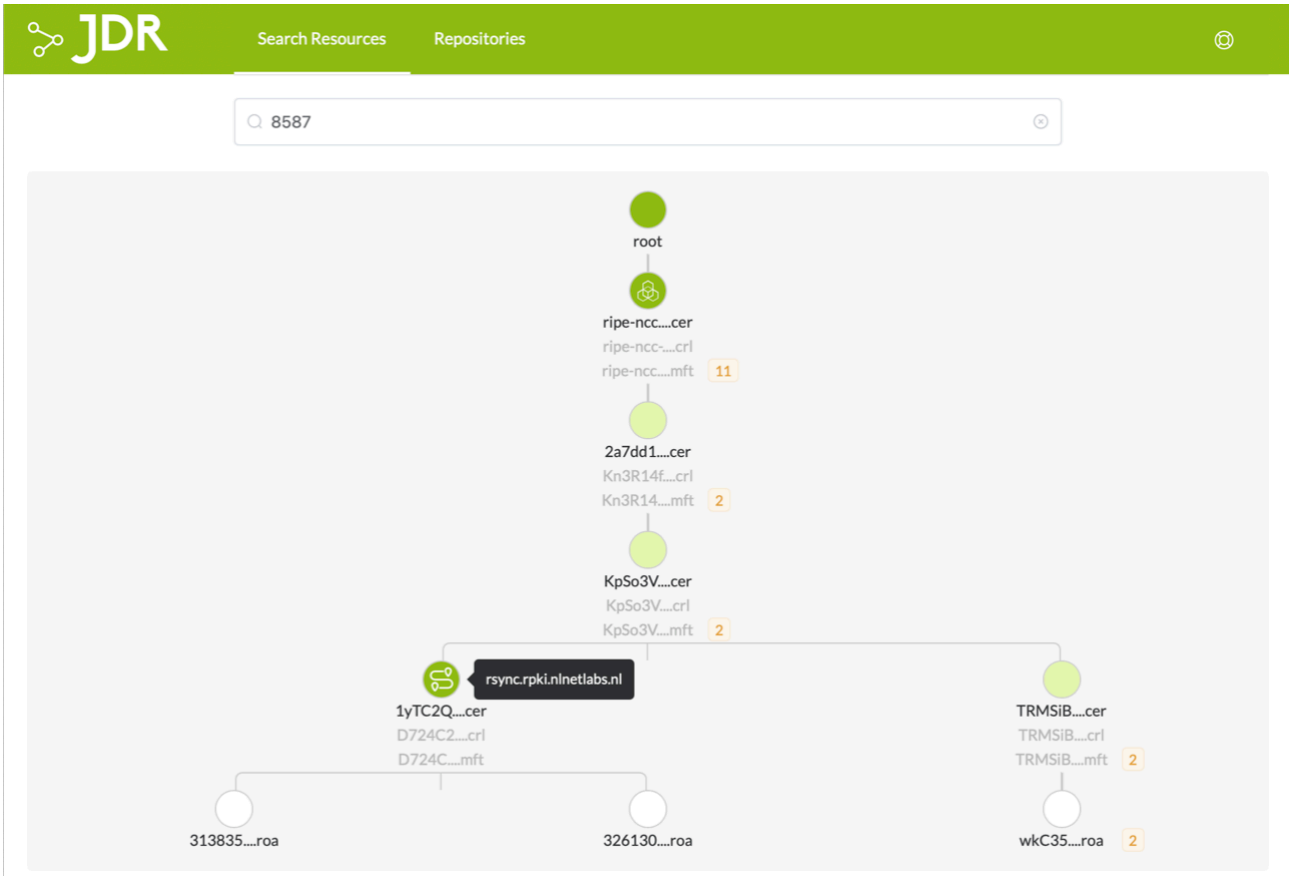

BGP routing building blocks

In 2021 we launched the Rotonda project[1], which aims to create a modular, analytical BGP routing engine. As with all of our other routing software, it is written in Rust, a fast, memory safe programming language. Rotonda will eventually consist of several components: first is the rotonda-store which was released in May. It handles the storage and retrieval of IP prefixes, using a tree bitmap as the data structure to store IPv4 and IPv6 prefixes. Other components we will release in the future will handle the protocols, the runtime and the command-line interface.

Routecore is a Rust library with fundamental building blocks for BGP routing – that is, types and traits for applications that need to deal with data related to BGP and routing. Routecore saw its first release at the end of 2021 and will grow more complete over time.

[1] https://blog.nlnetlabs.nl/donkeys-mules-horses/

Research

Introduction

Research is an essential part of NLnet Labs' mission (read our research vision here). As in previous years, we continued our research efforts in collaboration with both the academic community and industry. In this section we discuss our key research highlights of 2021.

Route Origin Validation of DNS resolvers

The Border Gateway Protocol (BGP) is responsible for routing on the Internet. BGP lacks built-in trust and security measures, making it vulnerable to IP prefix hijacking and route leaks. To defend against these threats, the Resource Public Key Infrastructure (RPKI) standard has been developed in the IETF. RPKI secures the Internet's routing infrastructure by signing and validating prefix origin data.

In the RPKI system, Route Origin Authorizations (ROAs) provide attestable statements specifying which prefix is authorised to be originated from which Autonomous System Number (ASN) in BGP. Route Origin Validation (ROV) is the process of using the data from the ROAs in RPKI to determine whether a route announced in the BGP is valid, invalid, or unknown.

There are, however, still situations where an organisation may indirectly fall victim to prefix hijacks, even if its own AS is RPKI-protected. A good example of this is the Amazon Route 53 BGP exploit, in which the prefixes of Amazon's authoritative DNS servers were hijacked. In this case, any AS with a DNS resolver not protected by RPKI would receive a valid but malicious response from the hijacked authoritative DNS server, even if the AS from which the query originated was RPKI-protected. So, for end-users to be fully secure, in addition to the network in which they reside, their DNS resolvers must also be based in RPKI-protected networks.

In this research project, we will:

- Measure the uptake of Route Origin Validation of DNS resolvers. We will do that by scheduling long-running measurements targeting authoritative name servers hosted on an RPKI beacon.
- Measure the uptake of Route Origin Validation of authoritative name servers. This will be accomplished by sending queries to the authoritative name server operators (drawing up an inventory from OpenINTEL data; see below for more information on this project) originating from an RPKI beacon.

We have been measuring the uptake of ROV protection of DNS resolvers since January 2020 (see this thesis report). The latest results of the measurements of ROV protection of DNS resolvers can also be found on the DNSThought website (here for IPv4 and here for IPv6).

Our intention is to perform ongoing measurements to monitor the state of RPKI protection of DNS resources over the long term. To that end, we have set up an RPKI beacon under our own control. Running our own beacon allows us to carry out these measurements for a longer period of time. In 2021, we continued our work on building and extending this infrastructure. Our measuring infrastructure and beacon is also being used by OARC's CheckMyDNS test platform.

Experimenting with DNS and XDP

In recent years, programmable network devices have received much attention from both academia and industry, and affordable hardware is becoming increasingly available. We think that network-programming technologies such as eBPF and P4 can also be used to improve the performance of DNS resolvers and name servers.

In 2020, we started a SURF-sponsored Research on Networks (RoN) project to assess eBPF's capabilities to improve the performance and stability of DNS resolvers and name servers. In that first phase we looked into the capabilities of the new technologies eBPF and eXpress Data Path (XDP). Using a proof-of-concept implementation, we wanted to find out how we can leverage the power of eBPF/XDP to improve resolver performance, increase name server versatility, and perform low-level measurements on high-speed connections. [blog]

This year, in the second phase of this project, we augmented existing DNS services with XDP programs, for the latter to do the heavy lifting in the kernel. We wanted XDP to quickly return the easy answers early on, while using the more sophisticated functionality of existing name servers and resolvers in user space for more complex tasks.

For this part of the project we selected Response Rate Limiting (RRL) as the use case for XDP. RRL protects against amplified denial-of-service attacks, in which the attacker sends requests to a DNS server, thereby soliciting large responses, while spoofing the source IP address of the request as the victim's, so that the victim is flooded with large answers. We extended RRL using XDP to exclude resolvers with which the name server operators have an established long-term relationship. This resulted in a list of networks that are always exempt from any RRL measures, so service to these networks is guaranteed, while other networks have RRL applied to them. [blog, MSc thesis] The outcomes of this project have been presented at three conferences: UKNOF 47, NANOG81 (video), and OARC34 (video). We also organised a XoT hackathon as part of the DNSOP track at IETF110 (video).

Other Research Highlights

OpenINTEL

In 2018, NLnet Labs joined the OpenINTEL project, whose goal is to serve as the "long-term memory" of the DNS. To this end, it performs daily measurements of over 60% of the global DNS name space. OpenINTEL is built on core NLnet Labs products (LDNS and Unbound). Other project partners are the University of Twente, SURFnet and SIDN.

OpenINTEL was extended in 2020 with a daily measurement of the IPv4 reverse DNS namespace. What distinguishes this initiative from other projects mapping reverse DNS is the fact that this measurement has a high frequency (daily) and also contains specific code to map the delegation hierarchy of the reverse name space, including RFC 2317 delegations of blocks smaller than /24 using CNAMEs.

In 2021, we continued our work on the OpenINTEL project.

Further Reading

You can read more about all the research projects NLnet Labs participates in on our website.

Community Outreach

Standardisation

NLnet Labs actively participates in the internet standardisation efforts of the IETF. In 2021, we contributed to several Internet Drafts in the DNS-related working groups and in the SIDROPS working group. For example, to improve privacy we contributed to DNS query name minimisation, and for security and resilience we contributed to the so-called DNS server cookies to mitigate DDoS and spoofing attacks. In SIDROPS, we contributed to improve operational aspects and provided operational recommendations for delivering resilient RPKI services. Next to contributing to Drafts, NLnet Labs is also an enthusiastic participant in IETF hackathons where the goal is to achieve the second half of the IETF's motto of "rough consensus and running code".

Our long-term commitment to open internet standardisation is also reflected in Benno Overeinder's appointment as one of the co-chairs of the IETF DNS Operations Working Group.

Internet.nl

NLnet Labs is a member of the Dutch Internet Standards Platform (Platform Internetstandaarden). Through this initiative, various partners from the internet community and the Dutch government collaborate to raise awareness of modern internet standards such as IPv6, DNSSEC, RPKI, TLS, STARTTLS/DANE, SPF, DKIM and DMARC.

The website Internet.nl, launched in 2015, is used to educate and entice users, government organisations and businesses to adopt modern internet standards. NLnet Labs was responsible for the development and maintenance of the Internet.nl portal until June 2021, when we handed the project over to the Internet Standards Platform.

RPKI Documentation Project

Started the end of 2018 and continued into 2021, NLnet Labs took the initiative to create a comprehensive documentation project for the RPKI ecosystem. This project brings together in-depth information about how RPKI works as well as documentation for tools from different open source organisations. The project has already received community contributions, for example from the developers of the RTRlib toolchain, as well as documentation for operational guidance.

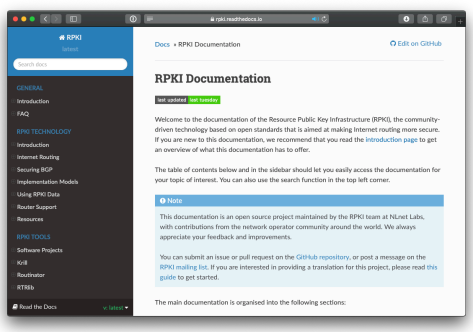

Presentations

NLnet Labs regularly presents at national and international conferences and meetings. In 2021, we mainly attended online meetings and presented at the various IETF, ICANN, RIPE and DNS OARC meetings. A full overview and the slide decks of our presentations can be found on our presentations can be found on our website.

Community Service

We fulfilled the following community positions in 2021:

Organisation	Role	Person
IETF DNSOP Working Group	Co-chair	Benno Overeinder
Forum Standaardisatie	Member	Benno Overeinder
ICANN	RSSAC Caucus member	Benno Overeinder Willem Toorop
ICANN	SSAC member	Jaap Akkerhuis
ICANN	Various advisory roles	Jaap Akkerhuis
ISO	ISO 3166 MA member	Jaap Akkerhuis
Internet Society	Member advisory council	Jaap Akkerhuis
DNS-OARC	Board member	Benno Overeinder
DNS-OARC	PC member	Willem Toorop
Quad9	Board member	Benno Overeinder

Academia

As of March 2021, Ronald van Rijswijk-Deij has been appointed as a Professor of Network and Security in the chair of Design and Analysis of Communication Systems (DACS) at the University of Twente. He will also remain involved with NLnet Labs as a principal scientist, e.g. to supervise a PhD candidate and to steer joint projects. In this capacity he will work with Benno Overeinder in giving direction to NLnet Labs' R&D efforts on both a strategic and a tactical level, collaborate with the people at NLnet Labs, and maintain contact with other parties.

Team

NLnet Labs strives to be a lean organisation, aiming to achieve its goals with minimal management overhead. We value diversity, aiming to employ staff members from a wide range of nationalities, cultures and backgrounds. Our goal is to be as open and inclusive as possible, bound together with our love of open source and open standards (read our Code of Conduct here).

Almost all our staff members are software developers or research engineers. The foundation strives to maintain a compact team, with a healthy mix of experience ranging from junior to senior and people who focus on software development or research. Other responsibilities – management, product development, finance and auditing, staffing and recruiting, sales and marketing – are shared by two people.

Recruiting

Recruiting new staff has become increasingly difficult in recent years. As of March 1, Tom Carpay started as a junior software developer. Finding a senior software developer (also partly ambassador and project manager) for Unbound and other DNS projects took considerable time and effort – the whole year of 2021, in fact. Philip Homburg took up this position in November.

Located at Amsterdam Science Park, NLnet Labs has strong local and international links with academia, research organisations and industry parties. Being part of that ecosystem makes us an appealing employer for developers/researchers with an interest in applied R&D and a love of impactful open-source software.

From the COVID restrictions, we learned that working online from home is feasible if not ideal. We are therefore open to roles that can partly be fulfilled remotely.

Every year, NLnet Labs supervises on average two to four graduating students. We also have room for one or two PhD candidates, though no-one took up the opportunity this year.

Funding

Income From Support and Development

Following the plan of previous years, a key goal for 2021 was to further increase the turnover from support contracts and paid software development. As a non-profit foundation, NLnet Labs is obliged to follow strict tax regulations and is not allowed to offer commercial services. Support and development contracts are therefore offered through Open Netlabs B.V. This company is a wholly owned, taxable subsidiary of the NLnet Labs Foundation. As such, it serves the non-profit public-benefit goals of its parent, and is guided and managed according to the NLnet Labs charter.

Open Netlabs offers support contracts with a service level for our production-grade software packages, such as NSD and Unbound. Customers receive support and early access to security patches, and through their financial contribution also support our mission to provide free and open software for all.

Open Netlabs also provides training and software development in the area of internet security standards, as well as consulting services such as installation and integration support, optimisation and auditing.

In 2021, Open Netlabs generated income from both support contracts and contracted software development. We are thankful that these contributions enable us to build free, open-source software in a sustainable way.

Grants and Subsidies

Every year since 2012 NLnet Labs has received a generous subsidy from SIDN. This pledge was renewed in 2017 for another five years. We are also grateful for the substantial, long-term grants that Infoblox, Verisign and Internetstiftelsen have donated.

Last but not least, we have also received numerous ad-hoc donations from organisations as well as individuals, for which we are equally grateful.

Financial Results NLnet Labs

Income			
	2020 Actual (k€)	**2021 Actual (k€)**	**2021 Budget (k€)**
SIDN Subsidy	175	150	150
Other donations	314	313	267
Consultancy and other income	153	138	105
Research and projects	178	165	385
Income from Interest	5	4	2
Total	**825**	**770**	**909**

Expenditure			
	2020 Actual (k€)	**2021 Actual (k€)**	**2021 Budget (k€)**
Staff	699	628	725
Housing	60	62	44
Travel	12	2	25
Depreciation	1	0	0
Project Costs	0	0	0
Other Costs	28	31	40
Sub Total	800	723	834
Negative Result Open Netlabs B.V.	-37	-36	0
Project Reservations	62	83	75
Total	**825**	**770**	**909**

Balance Sheet (k€)			
Assets		**Liabilities**	
Inventory	5	General Reserve	1378
Open Netlabs B.V. Stock and Loans	283	Special Purpose Reserves	102
Receivables	253	Current Liabilities and Accruals	98
Bank and Cash	1037		
Total	**1578**		**1578**

Governance

Stichting NLnet Labs was founded on 29 December 1999 by Stichting NLnet. Its board consists of four to seven members with staggered terms. The board's composition and most recent rotation schedule is shown below.

NLnet Labs Board in 2020		
Name	**Role**	**End of Term**
Cristian Hesselman	Chair	June 30, 2024
Marieke Huisman	Secretary	August 30, 2024
Sjoera Nas	Member	September 30, 2023
Andrei Robachevsky	Member	June 30, 2022
Jochem de Ruig	Treasurer	June 30, 2024

Four board meetings took place in 2021. Benno Overeinder participated in the board meetings in his role as director of NLnet Labs and as director of Open Netlabs BV.

Board members do not receive any compensation for their board work. Expenses may be reimbursed if necessary (€0 in 2021). The table below shows the additional functions held by board members and director of Stichting NLnet Labs.

Additional Functions Held By NLnet Lab Board Members and Directors in 2020	
Name	**Function(s)**
Cristian Hesselman	- Head of SIDN Labs - Member ICANN SSAC - Associate Professor University of Twente
Marieke Huisman	- Full Professor University of Twente
Sjoera Nas	- Senior Privacy Advisor, Privacy Company - Advisory Board SIDN Fonds
Benno Overeinder	- See the Community Service section for an overview
Andrei Robachevsky	- Technology Programme Manager Internet Society - Member EU MSP Standardisation
Jochem de Ruig	- Organic wine entrepreneur at Wilde Wijnen - Financial Director, Freedom Internet B.V.

Looking Ahead to 2022

Due to COVID, restrictions on physical gatherings were in place for most of 2021. During the lockdowns, we managed to continue much of our work through conference calls and communication platforms such as Slack and Mattermost. Facilitating social interaction such as informal conversations at the coffee machine was more challenging, but we organised a weekly virtual meeting to maintain social contact. The Discord server we set up for the RPKI community (as an addition to the existing mailing lists) has become a popular platform for more informal exchanges.

Despite the challenging circumstances we were able to keep most current affairs running smoothly. Software projects were started and successfully completed. We interacted with users via e-mail or conference calls. And we participated remotely in online community meetings such as those of IETF, RIPE, OARC and NANOG. Business matters, too, proceeded almost normally.

September was the first time our people attended a meeting again in person (NLNOG Day 2021). The only other in-person meeting this year was the NLUUG autumn conference.

We knew already, of course, that the internet infrastructure plays a fundamental role in contemporary economy and society. The most important lesson the COVID restrictions have taught us is that in a global emergency the current infrastructure is very capable of carrying large parts of these economic and social interactions. It makes us proud to know that the internet infrastructure, and our somewhat hidden contributions to it, held up so well. Nevertheless, we hope and expect that 2022 will be a calmer year in this regard.

Colophon

Editors

NLnet Labs

Design

Richard de Ruijter, Graphic Design & Illustration

Photo Credits

Photo on page 6 by Brett Garwood on Unsplash

Contact

Stichting NLnet Labs
Science Park 400
1098 XH Amsterdam
labs@nlnetlabs.nl
www.nlnetlabs.nl

© NLnet Labs