NLNETLABS

ANNUAL REPORT 2022

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 16 people, consisting almost exclusively of developers and researchers, with minimal overhead. We attract talented people who want to make a difference to 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 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 2022 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. Alongside bug and security fixes, stability improvements and smaller features, we have been working to implement PROXYv2, Extended DNS Errors (EDE), and an ACL per interface in Unbound. Support for DoQ is planned for next year.

NSD now supports IXFR-out and zone verification, and the software can run on small hardware devices such as the Raspberry Pi. Next year we will implement the eXpress Data Path (XDP) framework.

OpenDNSSEC received only minor updates this year, e.g. improved support for RFC 9276. For next year, our plans include implementing CDS/CDNSKEY, and migrating the current (external) OpenDNSSEC website onto our own site and documentation system. Support for ZONEMD and high-availability (fail-over) setups will also be released next year.

The most notable new feature of Domain Crate is support for SVCB and HTTPS record types. Ideas and concepts from the Connect by Name project, e.g. a stub resolver written in Rust, will also be developed further in the Domain Crate project.

Our routing security software – Routinator, Krill, RTRTR and Rotonda – continued to mature and evolve. Routinator got support for TLS to the RTR and HTTPS servers, and for BGPsec router keys. One of the new features planned for next year is to support TCP-AO for RTR connections.

Krill has improved support for migrations of pre-0.9 installations, and now comes with the new CLI tool krillup, which allows you to prepare and test the upgrade in advance. It also features BGPSec Router Certificate Signing, and PKCS#11 support for Hardware Security Modules (HSMs) for key operations. We are currently working on the renewing and extending the UI.

RTRTR now comes with a SLURM unit, which can be used to specify local policies and exceptions to the global RPKI data (per RFC 8416), and with an RTR-TLS unit to send RTR data over TLS connections.

Rotonda has seen steady progress this year: The rotonda-store component, the in-memory database for prefixes, has been tested in a large-scale pre-production environment. Routecore, the Rust library with parsers for BMP and BGP packets, has seen additions for parsing IPv4 and IPv6 unicast, and is evolving into a generic library that can take a wide variety of inputs. Work will continue on adding more BGP sub-protocols, such as IPv4/IPv6 multicast. We also started work on a domain-specific language called Roto for filtering, querying and configuring BGP applications. The product launch of Rotonda is planned for February next year.

DNS(SEC) Software Projects

Unbound

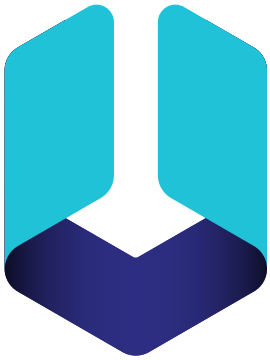

In 2022 we published three significant releases of Unbound (1.15, 1.16 and 1.17) and several maintenance updates (including bug and security fixes). New features include support for PROXYv2, Extended DNS Errors, and an ACL per interface. We also started the implementation of DNS-over-QUIC (DOQ), which is planned for release next year.

We have started using GitHub milestones to manage new releases of the Unbound software. This will provide insight into planned new features, their status and in what release they will become available to users. Version 1.17.1, which will be published early next year, will be the first release under this new regime.

PROXYv2 allows propagation of the client's original IP address and other connection information through HAProxy's TCP proxy. This is typically used with Unbound instances running behind a DNS load balancer. The implementation of PROXYv2 was sponsored by SUNET and Apple (the latter of whom agreed to upgrade its SLA agreement from Silver to Gold). Although the implementations of PROXYv2 and an ACL per interface were done at the request of industry partners, both are features that can be used more generally in DNS services.

Extended DNS Errors (EDE) 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. Extended error reporting is especially useful when validating DNSSEC; several large users of Unbound have already indicated that they will be using this new feature.

QUIC, defined in RFC 9000 and adjacent RFCs, is a new protocol alongside TCP and UDP. It brings together the efficiency of UDP, the reliability of TCP and the security of TLS in a single internet transport protocol. In a similar way to HTTPS/3, DNS-over-QUIC (DoQ, defined in RFC 9250) provides a native mapping of DNS transport on QUIC.

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 are enough to set up a resolver for your machine or network.

NSD

In a series of releases (versions 4.4, 4.5 and 4.6), we added support for IXFR-out and CreDNS. Version 4.4.0 added configuration options to minimise the memory usage of NSD, allowing the software also to run on small hardware devices such as the Raspberry Pi.

IXFR-out allows incoming changes to zone files (IXFR-in, from upstream) to be propagated in a similar way to other name servers (downstream). Previously NSD could only do this using full AXFR-out transfers, which is not optimal for specific setups in the DNS infrastructure.

IXFR-out was also the prelude to the implementation of zone verification, which allows zones to be checked for DNSSEC validity before publication. The development of this feature was partly sponsored by SIDN. Several parties, including Netnod and GoDaddy, are already using the new zone verification feature in their production systems, and others plan to do so.

We also started work that will increase the performance and scalability of NSD. Next year we will implement the eXpress Data Path (XDP) framework, a new network technology able to efficiently handle data over 10+ Gbps networks in applications. We already have a proof of concept for a new zone parser that is more efficient than the current implementation for loading zone files, and this will become available next year. We also have a proof of concept for an adaptive radix tree as the new core datastructure of NSD, which will be more efficient in time and memory consumption.

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 IETF standards.

OpenDNSSEC

After ending support for version 1.4 of OpenDNSSEC in 2019, over recent years 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.

Last year, in addition to the existing deb package, the community also made an rpm package file available, which makes it much easier to install OpenDNSSEC on Red Hat based Linux platforms.

After we have put considerable effort into outreach over the last three 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 2.1.11 and 2.1.12 minor releases improve support for RFC 9276 (with new recommendations for the salt and iteration values of NSEC3), and fix several bugs.

For 2023, our plans include implementing CDS/CDNSKEY and migrating the current (external) OpenDNSSEC website onto our own site and documentation system (last year we migrated the documentation of Unbound and NSD to the Read the Docs platform).

ZONEMD has already been implemented, and we are currently rewriting parts of the database interaction code to facilitate high-availability (fail-over) setups. Both features will be released in 2023.

CDS and CDNSKEY are two relatively new record types (defined in RFC 8087) that allow operators of signed zones to have the DNSSEC parameters in the parent zone updated (based on the existing DNSSEC chain of trust). This allows for automated key rollovers and other updates without the need to upload DNSSEC data through EPP or a web interface.

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

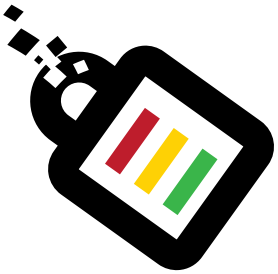

The SoftHSM project, to which NLnet Labs contributed in the past, was incorporated as a project under the Commons Conservancy in 2019. The long-term goal of this step was to keep the project sustainable and allow new partners to make significant contributions. The last release of the software, however, was version 2.6.1, published in 2020. We will keep managing 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 it can also be used by other cryptographic products.

DNS(SEC) Libraries

LDNS

In 2022, ldns saw two minor releases (versions 1.8.2 and 1.8.3) incorporating a 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 the SVCB and HTTPS record types is now compiled in by default, and EDNS0 Option handling and conversion into presentation format is now supported.

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 proofs of concept 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 2022, we released several minor versions of the getdns library and the Stubby resolver. Stubby 0.4.2 includes several bug fixes and some improvements to the user experience. Getdns 1.7.3 and Stubby 0.4.3 fix a problem in the systemd service configuration file.

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 through a collaboration between NLnet Labs, Sinodun and No Mountain Software. Although the code is written in C, bindings for several other programming 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)

2022 saw a series of minor releases of Net::DNS (version 1.20) and Net::DNS::SEC (versions 1.34-1.36). 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 longtime 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 2022 we released version 0.7 of Domain Crate, in which we refined existing functionality, implemented new features, and fixed some bugs. the most notable new feature is support for SVCB and HTTPS record types (which was added to Unbound, NSD and ldns last year).

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.

About Domain Crate: Domain Crate is a DNS library written in the Rust programming 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.

Connect by Name

Connect by Name is an NGI Zero project aiming to implement a library that allows a software developer to set up internet connections from an application in the most private and secure manner using well established and open standards (e.g. DNSSEC validation, IPv6 (Happy Eyeballs), TLS, and TLS+DANE).

The project was completed in October, delivering a proof-of-concept library that allowed us to gain experience with the technologies involved. In addition an IETF Internet Draft was written to standardise the configuration of local DNS proxy instances exchanging traffic between applications and DNS servers in the most secure way. The Draft was presented at the IETF 115 DNSOP WG meeting.

As part of this project, DNS over HTTPS (DoH) was implemented in the getdns library.

The ideas and concepts from this project will be developed further in the Domain Crate project.

Routing Software

Routinator

In 2022 we published two new releases of Routinator and several maintenance updates (including bug and security fixes). Version 0.11 implemented TLS to the RTR and HTTPS servers, and BGPsec router keys. Version 0.12 comes with a restructured TAL configuration, which will use the bundled RIR TALs directly unless told otherwise. These release also incorporate many more smaller changes and additions.

At the end of 2022, over 2200 organisations were using Routinator, bringing its market share to more than 75 percent (and still growing fast). Independent statistics on the RPKI Validator market can be found at: https://dataplane.org/rpki.html.

About Routinator: Routinator is Relying Party software (written in the Rust programming language), also known as an RPKI Validator. The full-featured application is designed to be secure and highly portable. It is a lightweight implementation that can run effortlessly on almost any operating system using minimalist hardware.

The Routinator system periodically downloads and validates the global RPKI dataset. Routers can connect to Routinator to fetch validated data via the RPKI-to-Router (RTR) protocol. The built-in HTTP server offers a user interface and API endpoints for various file formats (e.g. CSV, JSON and RPSL), as well as logging, status and Prometheus metrics.

The user interface allows users to validate prefixes against Autonomous System Numbers (ASNs) found in BGP announcements. Next to that it allows users to lookup related prefixes for the prefix they are searching for. These related prefixes can be more or less specific prefixes, prefixes routed in BGP, or prefixes that are allocated by one of the five Regional Internet Registries.

For larger networks we have developed RTRTR (discussed below) as a companion to Routinator. This makes it possible to centralise validation performed by Routinator and have RTRTR running in various locations around the world to which routers can connect.

RTRTR

RTRTR saw three new releases in 2022 (0.2.0, 0.2.1 and 0.2.2). Version 0.2.0 introduces a SLURM unit, which can be used to specify local policies and exceptions to the global RPKI data (per RFC 8416); and an RTR-TLS unit, which sends RTR data over TLS connections. The two subsequential minor versions contain some bug fixes and smaller changes.

About 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.

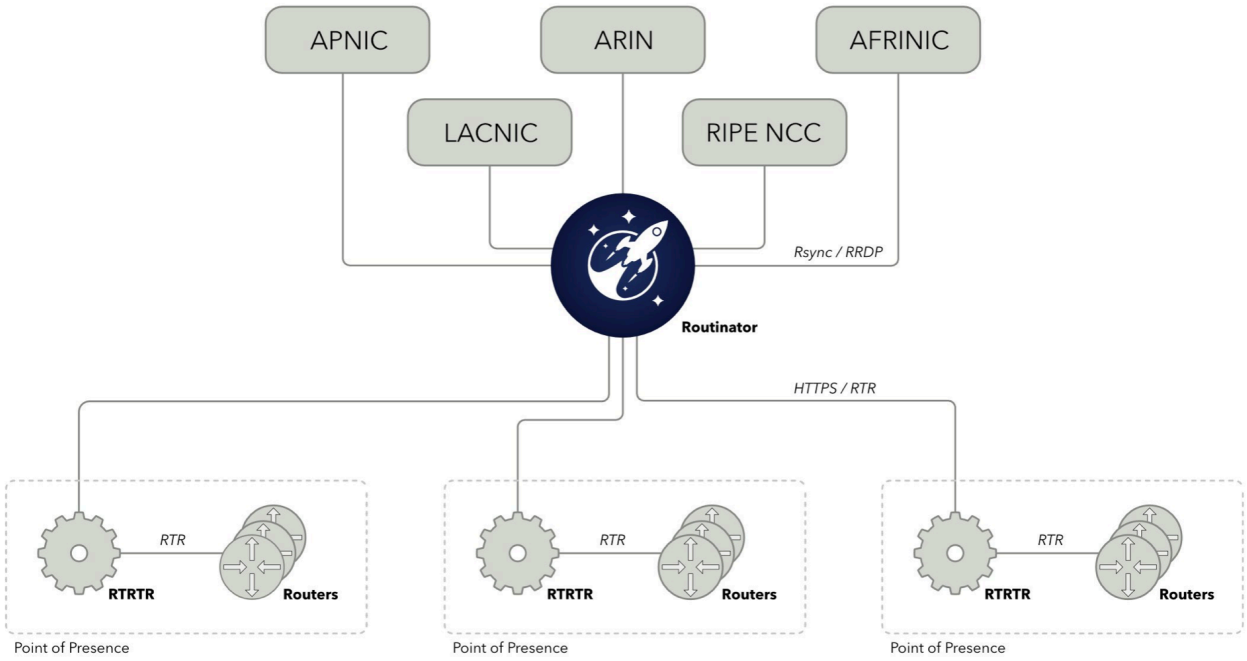

For larger networks, RTRTR is an ideal companion to Routinator. For example, it is possible to centralise validation performed by Routinator and have RTRTR running in various locations around the world to which routers can connect.

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 dataset in JSON format, as well as a monitoring endpoint in plain text and Prometheus format.

Krill

In 2022 we published no less than eight releases of Krill, including three feature releases (0.10.0, 0.11.0 and 0.12.0).

Version 0.9.5 improved support for migrations of pre-0.9 installations. Upgrades are now more secure, and downgrades can be performed automatically in case of problems. With 0.9.5 came the new CLI tool krillup, which allows you to prepare and test the upgrade in advance. We also adjusted the CA parent refresh logic; decreasing the checking frequency from 10 minutes to 24–36 hours decreases the load on CA parents with many children.

Version 0.10.0 introduced BGPSec Router Certificate Signing, and PKCS#11 support for Hardware Security Modules (HSMs) for key operations.

Version 0.11.0 allowed an optional comment for each ROA configuration, and the ability to show the ROA objects for each ROA configuration. Both features are not yet available in the user interface (UI), but will be added in the upcoming make-over.

For the renewal and extension of the UI (which is almost ready for launch now) we have contracted the company Tweede Golf, based in Nijmegen. We aim to turn this relationship into a

long-term arrangement, providing us with a flexible pool of development capacity for the core of our Rust projects.

Version 0.12.0 vastly reduced the CPU usage of Publication Servers for big RPKI repositories.

Other minor improvements, changes, and bug and security fixes were incorporated along the way.

The implementation of the PKCS#11 support for HSMs was sponsored by NIC.br and APNIC. At the request of NIC.br, support for the Key Management Interoperability Protocol (KMIP) was added as well (they also purchased a support contract).

These are important features for Enterprise users and RIRs, making Krill a serious option for these types of organisation. As a result, ARIN, APNIC and RIPE NCC have all introduced Hybrid RPKI services based on Krill's Publication Server. Large parties including Charter, AT&T, Microsoft, Google and Verizon intend to use these services (Charter has already published its first RPKI objects and asked for specific components), boding well for the sale of Krill support contracts.

The other two RIRs, AFRINIC and LACNIC, plan to migrate their RPKI (CA) infrastructure to Krill next year. Deployment at LACNIC required the implementation of Trust Anchor support, which is almost complete. LACNIC sponsored the implementation of this functionality and purchased a support contract for Krill. In the meantime, the software company Tweede Golf has been commissioned by Open Netlabs to develop High Availability (HA) functionality (also sponsored by LACNIC) to allow for redundant deployments. AFRINIC, which has smaller financial resources, hopes to share the benefits of this new functionality for its deployment of Krill.

Krill is now actively being used by more than 1700 organisations in Brazil and over 1200 organisations in Indonesia.

About Krill: Krill is an RPKI Certificate Authority (CA) that lets you run Delegated RPKI under one or multiple Regional Internet Registries (RIRs). Through its built-in publication server, Krill can publish Route Origin Authorisations (ROAs) on your own servers or with a third party.

The software supports running the CA both upwards and downwards. Upwards means that an instance can have multiple parents, such as ARIN and RIPE NCC, 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 (NIRs) and Enterprises.

A publication server is included in Krill, but can also be run as an independent component. This allows organisations to host published certificates and ROAs themselves, or let a third party such as a Content Delivery Network (CDN) 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; or
- need to be able to delegate RPKI to their customers or business units, so that that they can run their own CA and manage ROAs themselves; or

- manage address space from multiple RIRs; using Krill, they can manage all ROAs for all resources seamlessly within one system; or
- want to be operationally independent from their parent RIR, such as an NIR or an enterprise.

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

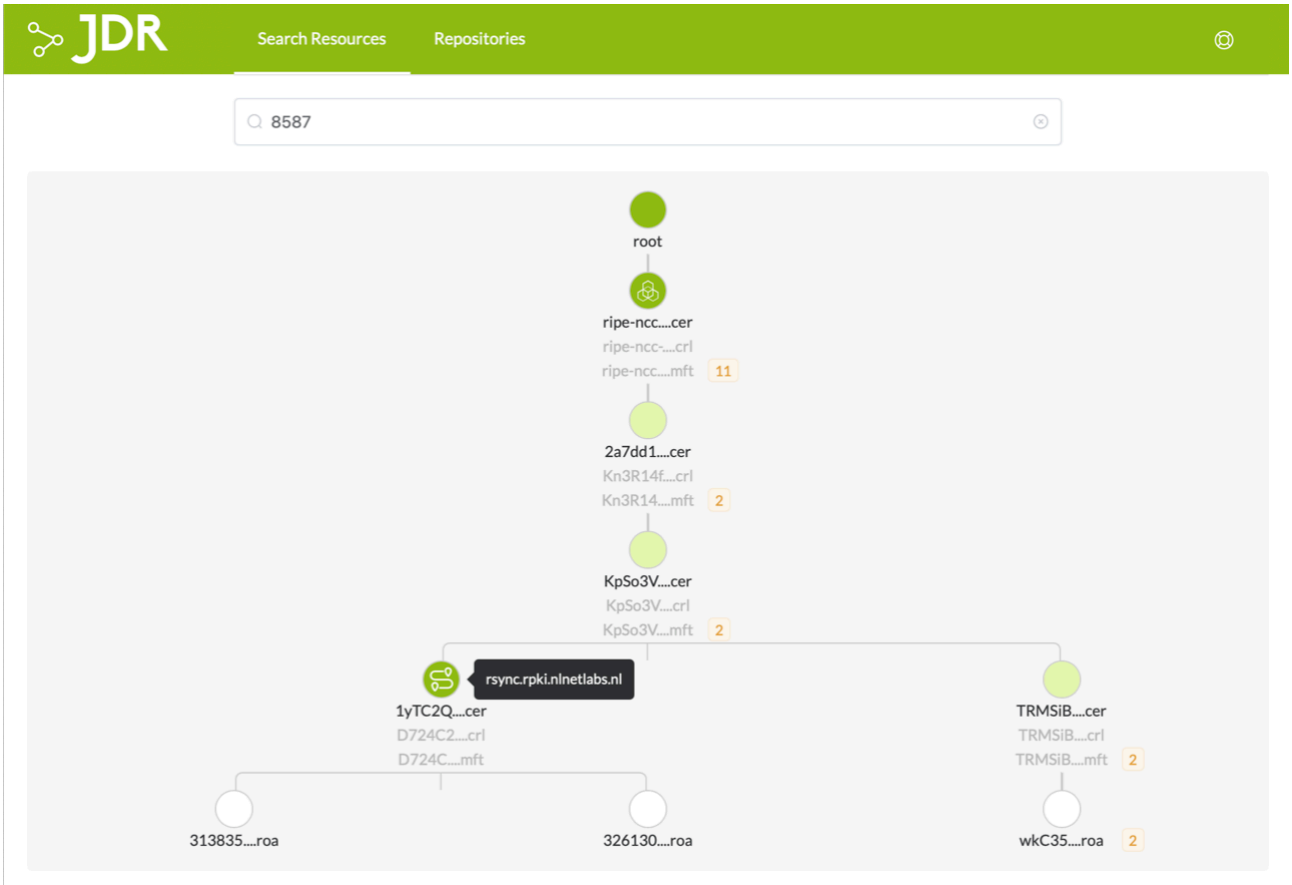

by RP software, JDR will try to process it and give the end-user as much troubleshooting information as possible.

In 2022, JDR has not seen major changes in terms of functionality. With the increasing size and distribution of the RPKI repositories, the deployment and parts of the codebase have been adapted to handle this increased load, and operate in a more resilient way.

Rotonda

The Rotonda project that was launched in 2021, has seen steady progress in 2022. The rotonda-store component, the in-memory database for prefixes, has been tested in a large-scale pre-production environment. Routecore, the Rust library with parsers for BMP and BGP packets, has seen additions for parsing IPv4 and IPv6 unicast, and is evolving into a generic library that can take a wide variety of inputs. Work will continue on adding more BGP sub-protocols, such as IPv4/IPv6 multicast, L2VPN, etc. New releases for both rotonda-store and routecore have been issued in 2022.

We also started work on a domain-specific language called Roto for filtering, querying and configuring BGP applications. Roto is a statically typed, compiled language, that is not Turing complete. Roto draws inspiration from router configuration languages from different vendors, both open source and proprietary, and from RFC standards, such as YANG based configuration.

The work we did on Roto this year includes research, design, and developing a prototype of the compiler and one of the virtual machine. Both the compiler and the virtual machine will be included in Rotonda.

We are now working towards the product launch of Rotonda, expected in February next year at the NANOG conference in Atlanta, together with our development partner Arelion, one of the largest Tier-1 providers in the world. Our regular design bureau can be delivering a logo and corporate identity at any time.

In the meantime, we are extending the partnership with Arelion: In addition to the purchased support contracts for Rotonda and Routinator, Arelion will set up a dedicated 10 Gbps connection at its Point of Presence at NIKHEF (at Amsterdam Science Park). A dark fiber connection to our office will allow us to test our implementation continuously and at large scale through a peering session with hundreds of production routers from almost every manufacturer. A similar arrangement with SURF is currently being discussed.

Internet Initiative Japan (IIJ) and several Internet Exchanges have expressed an interest in this project.

About Rotonda: Rotonda aims to create a modular, analytical BGP routing engine, made out of components that can be combined into a BGP application by the end-user. 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 handles the storage and retrieval of IP prefixes using a tree bitmap as the data structure.

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.

Other components are the Roto filter and querying language, and the rotonda runtime, that will handle the protocols, the runtime and the command-line interface.

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 2022.

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, however, 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/ROV 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 originate 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) 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).

To perform ongoing measurements to monitor the state of RPKI protection of DNS resources over the long term, we have set up an RPKI beacon under our own control, replacing the beacon kindly provided by Job Snijders until September 2021. Running our own beacon allows us to carry out these measurements for a longer period of time. In 2022, 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 to measure the Route Origin Validation status of resolvers.

This year, Willem Toorop made four presentations on the beacon:

- 'Measuring Route Origin Validation' at the IEPG at IETF113 on 20 March 2022 in Vienna;[1]
- 'Tools for Measuring Route Origin Validation' at the CENTR 46th Technical Workshop on 21 March 2022 in Vienna;
- 'Het opzetten van een Resource Public Key Infrastructure (RPKI) baken' at NLUUG vj22 on 20 May 2022 in Utrecht;[2]
- 'Setting up an RPKI beacon' at ColoClue Presents on 26 November 2022 in Weesp.

Koen van Hove did a small research project using our beacon, which is presented in a blog post on our own website[3] as well as on the RIPE Labs website.

We supervised two Security and Network Engineering students, who used our beacon for their research project on measuring Route Origin Validation of authoritative DNS servers. They presented their work on 23 October at the 39th DNS-OARC and 47th CENTR Technical Workshop in Belgrade.

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.

In this project's second phase, which started last year, we create a series of example programs to show that DNS service augmentation agnostic of the DNS software in userspace is possible with XDP. To that purpose we augment 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.

The example program presented last year was rate-limiting. This year, we expanded on the idea by creating a DNS-software-agnostic service augmentation that also impacts outgoing traffic: we extended the rate-limiting example of 2021 with a DNS Cookie implementation. This resulted in a blog post published on our own website as well as on the APNIC website. We also had a

---

[1] http://iepg.org/2022-03-20-ietf113/Measuring%20Route%20Origin%20Validation.pdf

[2] https://nluug.nl/evenementen/nluug/voorjaarsconferentie-2022/talks/willem-toorop-het-opzetten-van-een-resource-public-key-infrastructure-rpki-baken/

[3] https://blog.nlnetlabs.nl/measuring-the-impact-of-rpki-rov/

podcast in APNIC's PING about our XDP work up to that point, which was also published on all the usual podcast channels.

With rate-limiting at the XDP layer, DNS Cookies also have to be implemented at the XDP layer, because a valid DNS Cookie allows rate-limiting to be bypassed. The same holds for collecting telemetry and statistics. Catching all traffic this early in the network stack bypasses all the conventional means of collecting statistics. We explored and provided example programs to collect this information at the XDP layer as well, as described in another blog post published on the APNIC website, on our own website, and at RIPE Labs. We also recorded another podcast with APNIC's PING about this work.

Other Research Highlights

2STiC

2STiC, short for Security, Stability and Transparency in inter-network Communication, is a joint research programme in which ten Dutch internet organisations collaborate. Its goal is to develop and evaluate new or improved mechanisms that increase the security, stability and transparency of internet communications, through both extensions of today's internet and emerging internet architectures. Its long-term objective is to establish a collaborative research centre in the field of trusted and resilient internet infrastructures that will help to put the Dutch and European networking communities in a leading position in the field.

Our participation in the 2STiC consortium focuses on two projects for which we contribute expertise on internet architecture and standards, and review results and papers:

- UPIN (User-driven Path Verification and Control for Inter-domain Networks):
  The goal of UPIN is to develop and evaluate a scalable distributed system that enables users to cryptographically verify and easily control the paths through which their data travels through an inter-domain network like the Internet, in terms of both router-to-router hops and router attributes (e.g. router location, operator, security level, and manufacturer).
- CATRIN (Controllable, Accountable, Transparent: the Responsible Internet):
  The goal of CATRIN is to start up the Responsible Internet, a novel security-by-design concept and extension to the internet infrastructure that enhances the range of actions users have at their disposal to share information securely and confidentially. This will enable higher levels of trust and autonomy for users, organizations, and societies.

ICANN DNSSEC Deployment Metrics

SIDN Labs was commissioned by ICANN to perform a study on DNSSEC Deployment Metrics, to which we contributed about 40 hours.

The project was successfully completed in June, and the final report was published by ICANN in October. The report presents how the DNS community currently measures the implementation and deployment of DNSSEC, and makes recommendations regarding the improvement of measurement activities.

DNSThought

DNSThought is a DNS measurement and data analysis platform that provides longitudinal insight into resolvers and their capabilities. The project leverages the RIPE Atlas data collection measurement infrastructure, giving it a unique perspective from many vantage points on the Internet. Since its inauguration in April 2017, resolver capability measurements have been performed every hour, resulting in a valuable data set for the research and operations community. The continuous measurements allow researchers and operational engineers to study the status of standards and resolver functionality over time.

Last year, the RIPE NCC Community Projects Fund awarded funding of EUR 40,000 for the renewal of the platform. The old website is static and is updated once a day through a script. The new site will be interactive, featuring a datastore and API accessible to researchers and operators.

Willem Toorop has created the backend, and a freelance developer has been working on the frontend. We are currently putting the final touches to the corporate identity and the textual content of the portal. The new site is expected to go live early next year, after which we will present the results at various conferences.

RSSAC028

In 2017, ICANN's Root Server System Advisory Committee (RSSAC) presented various options for renaming the root servers, so that information about them can be secured using DNSSEC, and so that root servers can become independent of the .net top-level domain (TLD) in some cases. The RSSAC put forward five alternative naming schemes.

In partnership with SIDN Labs, we are now conducting a study to establish what implications these five alternative naming schemes would have for the root servers of the Domain Name System (DNS). Adoption of a new naming scheme is under consideration because it could, for example, make the root less dependent on the .net TLD.

We kicked off this project this summer by giving a short introduction and presentation to the Root Server Operators (RSOs). In September we started on the first work package, sending all RSOs a questionnaire.

The project is expected to be completed early next year.

Further Reading

You can read more about all the research projects NLnet Labs participates in on our website at https://nlnetlabs.nl/research/projects/.

Community Outreach

Standardisation

NLnet Labs actively participates in the internet standardisation efforts of the IETF. In 2022, we contributed to several Internet Drafts in the DNS-related working groups and the SIDROPS working group. For example, to improve security and resilience we contributed to the so-called DNS server cookies to mitigate DDoS and spoofing attacks. This year we worked with other open source developers on the catalog zone standard to simplify the configuration and deployment of a large number of domains by a DNS service provider. In SIDROPS, we contributed to improve operational aspects and provided operational recommendations for delivering resilient RPKI services. In GROW we discussed the standards relevant to BGP monitoring and data collection, which we use in the Rotonda project. As wel as 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.

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.

Policy and governance

Advocating for good governance of an open Internet aligns very well with our mission and objectives. We actively contribute to different policy-making bodies, including ICANN and Forum Standaardisatie; and we engage with the EU on regulations such as the forthcoming Cyber Resiliency Act which was presented by the European Commission in fall 2022. In these interactions, we emphasise the importance of embracing open standards and open source. Our advocacy extends beyond the interests of NLnet Labs, representing the broader open standards and open source community to ensure that new regulations uphold the principles we adhere to and do not compromise them.

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.

Presentations

NLnet Labs regularly presents at national and international conferences and meetings. In the first months of 2022 we presented at the last of the online-only meetings that resulted from the COVID restrictions. From March onwards things were starting to go back to normal, and we were physically present at the various IETF, ICANN, RIPE, DNS OARC, CENTR and MORE-IP meetings. A full overview and the slide decks of our presentations can be found on our website at https://nlnetlabs.nl/community/presentations/.

Community Service

We fulfilled the following community positions in 2022:

Organisation	Role	Person
IETF	DNSOP co-chair	Benno Overeinder
IETF	Hackathon co-chair	Benno Overeinder
Forum Standaardisatie	Member	Benno Overeinder
ICANN	RSSAC Caucus member	Benno Overeinder Jaap Akkerhuis 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
RIPE	DNS WG co-chair	Willem Toorop
RIPE	BCOP TF co-chair	Benno Overeinder
Quad9	Board member	Benno Overeinder
NLUUG	PC member	Willem Toorop

Academia

As of March 2021, Ronald van Rijswijk-Deij has been appointed 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 works with Benno Overeinder in giving direction to NLnet Labs' R&D efforts on both a strategic and a tactical level, collaborates with the people at NLnet Labs, and maintains 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: https://nlnetlabs.nl/conduct/).

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/science. The team now also has two members focused on bridging policy and technology. Other responsibilities – e.g. management, product development, finance and auditing, staffing and recruiting, sales and marketing – are shared by two people.

Recruiting

As of May 1, Maarten Aertsen started as a Senior Internet Technologist (Policy & Governance Expert). He is a member of the R&D team and in that role focuses on policy development adjacent to open standards and open-source software. This year, Maarten facilitated the strategy sessions with the DNS team and worked with Benno Overeinder on the strategic plan for 2023-2025 and the operational plan for 2023. With regard to his external role, Maarten focuses on policy advice regarding open standards and open-source software. In the Netherlands, he will be advising the Dutch government and public sector. In Europe, he will be following the NIS2 implementation and the policy development on a Cyber Resilience Act. Maarten expressed interest to work with IETF and ICANN in the future..

Jeroen Koekkoek, a former colleague, started on March 1 with our subsidiary Open Netlabs, where he works on NSD and other DNS projects.

Recruiting new staff has become increasingly difficult in recent years. With these two new staff members, however, the team is back up to strength and now has sufficient capacity to realise the organisation's plans and ambitions for the near future.

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.

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 2022 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 2022, Open Netlabs generated income from both support contracts and contracted software development. These returns are still growing, and we hope eventually to be able to sustain our development efforts through support contracts, providing continuity to both software branches. 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 2022 for another five years. We are also grateful for the substantial, long-term grants that Infoblox, Verisign and Comcast 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.

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.

LACNIC funded features such as offline Trust Anchor and general improvements for Krill and the use of Routinator for pre-validation before publication.

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.

Financial Results NLnet Labs

Income			
	2021 Actual (k€)	**2022 Actual (k€)**	**2022 Budget (k€)**
SIDN Subsidy	150	150	150
Other donations	313	217	257
Consultancy and other income	138	132	115
Research and projects	165	209	310
Income from Interest	4	2	1
Total	**770**	**710**	**833**

Expenditure			
	2021 Actual (k€)	**2022 Actual (k€)**	**2022 Budget (k€)**
Staff	628	674	680
Housing	62	68	69
Travel	2	31	30
Depreciation	0	3	0
Project Costs	0	1	0
Other Costs	31	41	35
Sub Total	723	818	814
Negative Result Open Netlabs B.V.	-36	-3	0
Project Reservations	83	-105	19
Total	**770**	**710**	**833**

Balance Sheet (k€)			
Assets		**Liabilities**	
Inventory	5	General Reserve	1341
Open Netlabs B.V. Stock and Loans	286	Special Purpose Reserves	35
Receivables	127	Current Liabilities and Accruals	75
Bank and Cash	1033		
Total	**1451**		**1451**

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 2022		
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, 2025
Jochem de Ruig	Treasurer	June 30, 2024

Four board meetings took place in 2022. Benno Overeinder participated in the board meetings in his role as director of NLnet Labs and Open Netlabs BV.

Board members do not receive any compensation for their board work. Expenses may be reimbursed if necessary (EUR 0 in 2022). 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 2022	
Name	**Function(s)**
Cristian Hesselman	- Director of SIDN Labs - Member of ICANN SSAC - Professor University of Twente - Member of the ACCSS Advisory Board - Member of the Supervisory Board of the Enschede Public Library
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 2023

NLnet Labs is currently experiencing a growth phase, marked by the recruitment of several new colleagues to propel us towards our goals and aspirations. This expansion encompasses both NLnet Labs and its wholly-owned subsidiary, Open Netlabs. The strategic hiring of talented individuals is enhancing the strength of our organisation.

There is a positive trend in income growth, particularly evident in the steady increase of recurring income from Open Netlabs for support contracts. Despite using a portion of our reserve to cover operational costs, the ongoing growth positions us favourably for the future.

Our vision for the next five years is to further expand recurring revenue from support contracts to cover all fixed costs (mainly salaries and premises) in both organisations.

.

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