# The history and future of the DNS

Anne-Marie Eklund-Löwinder and Ulrich Wisser, Internetstiftelsen

# The history and future of the DNS

## DNS: the early days

More than 30 years ago, more precisely in 1983, two engineers at the University of Southern California, Jon Postel and Paul Mockapetris, created a key component that has become a cornerstone and vital part of the internet infrastructure of today; the Domain Name System or DNS.

These engineers ran the first successful test of a system that made it possible for computers to find each other online and send information between each other without having to manually search for the address of each individual machine.

After these early attempts to make it easier to reach hosts across the internet, a collection of engineers led by Paul Mockapetris got together and created a description of the Domain Name System. This work took place within the Internet Engineering Task Force (IETF) and the RFC series (Request for comments). The first-generation DNS was accepted as an IETF Internet Standard and described in two RFCs (Request for comments): RFC 1034 and RFC 1035.

More specifically, the term Domain Name System refers to two different things. Firstly, it refers to the protocol used to convert human-readable labels (such as computer hostnames) into numeric IP addresses. Computers on the internet locate each other using numbers, not letters. Secondly, it refers to the activity to build a global service using that said protocol to enable communication via the internet.

The two documents mentioned above mark the beginning of the DNS definition. They describe a fully-functional protocol and include some early data types to manage. Internet mail (SMTP) was defined about the same time, and there were serious attempts to get e-mail to make good use of the DNS.

Today, users surf the internet, visiting websites to get information, to send electronic mail, to do their online banking, et cetera by simply clicking on a link or entering an address into the intended field, which may for instance look like https://www.internetstiftelsen.se or info@internetstiftelsen.se. Invisibly, and most of the time even without the user's knowledge, each activity starts a time-critical and sensitive process before the actual resource can be accessed, whether a web server or an e-mail server or something else.

In order to do that, the user's computer must be able to find the unique address of the recipient's server. The fact that each computer needs a unique address does not mean that it always has the same unique address.

In this way, internet addresses may be considered as temporary and not something that one should refer to in the communication between computers, or between computers and a user. The use of dynamic IP addresses, DHCP, greatly contributes to this. Domain names are more constant than IP addresses. Every time a user enters a domain name in their e-mail client or web browser, a process that translates the user-friendly domain name to the computer-friendly IP address starts in order to locate the resource at the other end. To find a particular server with a particular service, you use domain names, and the DNS helps to tell which IP address that server or service has at that moment.

So, in the beginning, we used domain names to name servers. Eventually, the domain names came into use also for addressing electronic mail. Today, we use domain names to identify all kinds of services - something that may not have been the original intention.

The DNS is in itself a distributed database of information that devices use to look up domain names from IP addresses and vice versa. The information that constitutes the internet's DNS is provided through a network of thousands of name servers, each responsible for referring users to the internet in the right direction so that they can reach what they want.

Domain names are hierarchically organized and distributed according to a strict global hierarchy with a tree structure. Each node in the tree can have zero or more sub nodes.

## How does it actually work?

Normally, each computer or local network connected to the internet gets assistance from a name server (DNS server) to which other nearby computers can turn to ask their questions. The name server provides the answers, either from their own database or, if the requested information is not there, it retrieves the information on the internet by being referred across different name servers until it reaches the answer and can return it to the requester. In these cases, the local name servers cache and save the answer for a while, in case a computer asks the same question again in the near future. This avoids unnecessary traffic on the internet.

The domain name system does not do searches, just lookups. One must therefore have a well-defined, unique search key to get an answer. All web address entries are stored in the DNS. When a user enters a web address into their browser, parts of that address serve as a lookup key in the DNS, and thus one can get the web server's IP address in response. It is also possible to store other types of information in the DNS, like the IP address for the mail server for a particular domain.

The process begins with a component called "resolver", which forms part of the user's application. When someone enters something that looks like a domain name in for example the browser, a query is sent to the local name server in the user's own system. Sometimes the resolver can respond directly, but sometimes it must refer to another name server. The following conversation illustrates how it works.

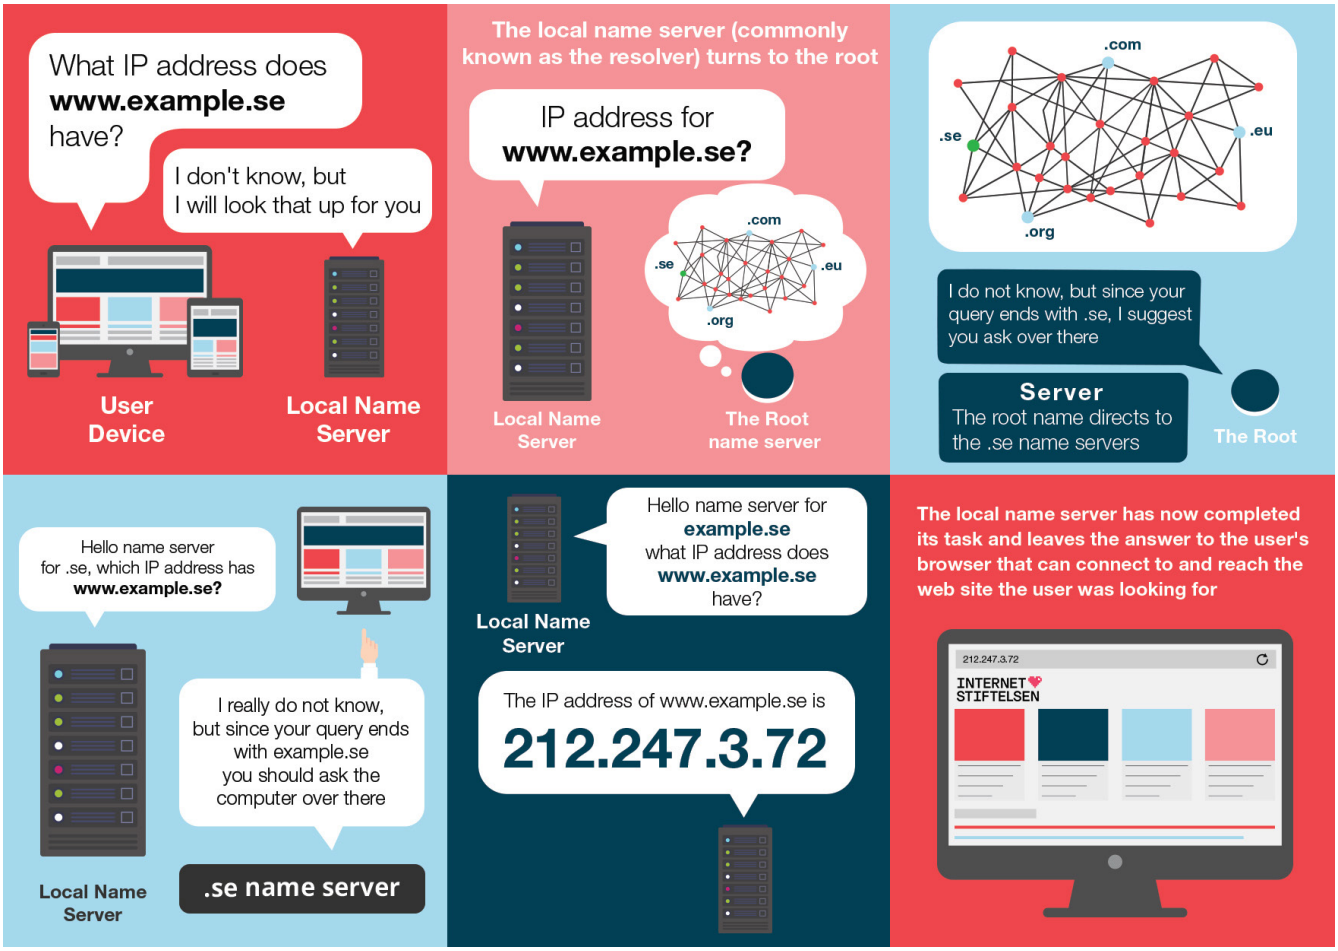

## Step by step improvements of DNS

For the early implementations of the DNS, the best way to provide for continuity was to have multiple servers answering multiple queries. One server - called a master – controlled a number of slave servers. Each of the slave servers got instructions to make contact with the master periodically to check if the data had changed.

About one decade passed before the publishing of the first major update to the DNS protocol. That was an addition of a more dynamic way to keep the DNS data up to date with the use of two new mechanisms; NOTIFY and Incremental Zone Transfer (IXFR).

NOTIFY was a real game changer. Rather than having the master wait until a slave came to check for new data, the master could send a NOTIFY message to the slaves, to get them to acquire the new data.

On top of that, IXFR made a marked change to the way data was distributed. With AXFR, which was the first generation DNS way to do it, the entire zone with all the data records travelled from the master to the slaves with the changes included. IXFR changed that model by enabling only the changes to be sent.

The next important improvement in the evolution of the DNS was dynamic updates defined in RFC 2136. In early DNS, to change even just one single record, the administrator would have to go to the master server, edit the file, and then get the master to reload the file (before waiting for slaves to update).

Dynamic updates allowed an administrator to edit the live zone, even remotely. Administrators did not need to log into the master. With that change followed a greater insight. Dynamic updates reused the original message format for another purpose. This led to subsequent efforts to update the DNS, not being afraid to redefine fields in the protocol, such as the Extension Mechanisms for DDS (EDNS) in RFC 2671, which defined extensions that added further modernization to DNS.

After the addition of NOTIFY, IXFR, and dynamic updates, the evolution of the DNS protocol began to unravel. More code was added here and there, but no one properly reviewed the protocol to check for structural integrity.

This period came to be documented in RFC 2181 and RFC 2308. RFC 2181 was simply titled “Clarifications to the DNS Specification” and dealt with some data issues that were considered to be overlooked. RFC 2308 covered answers that said “no” and helped document terminology which is still used today.

After the finalisation of the “reforms” documented in RFC 2308 and RFC 2181, the next top focus of DNS modifications was Secure DNS or DNSSEC and would remain so for many years to come.

## Potential weaknesses in the DNS

Although attacks against the DNS are not as common as, for example, virus attacks, they do occur, and they are becoming more common by the day. However, how often they occur is difficult to say.

Because the DNS is a distributed database, each domain holder or name service provider (on behalf of the domain holder) manages its own part of the database. The local administration makes it easier to keep the database up to date. However, there is nothing to prove who provides what information, which means that it is quite possible to forge information and put it in the DNS database and thus fool users. It is therefore not possible to know for sure whether the information you receive in response to the database is reliable or not.

In 1990, the security researcher Steven Bellovin described cache poisoning for the first time, but the report was held back until 1995.

False DNS information opened up opportunities to steal information from others or to interfere with various kinds of transactions, for example intercepting e-mail or redirecting internet shoppers.

## Domain Name System Security Extension

In the original implementation of the DNS there was no way to verify that the information retrieved from the DNS was genuine and undamaged. The problems and the need for a security supplement to the DNS service were known for a long time. Work on developing such a security supplement has been going on for a number of years and is known today as DNSSEC. The name Secure DNS also comes up. The security supplement is based on the use of cryptographic techniques for electronic signatures.

DNSSEC (Domain Name System Security Extension) is a more secure way of doing lookups of internet addresses, for example web and e-mail. In contrast to the usual domain name system (DNS), lookups with DNSSEC are cryptographically signed, which makes it possible to ensure that they come from the right source and that the content has not been tampered with during transmission.

As the first Top Level Domain (TLD) in the world to adopt and offer DNSSEC, Internetstiftelsen (the Swedish TLD registry) signed its zone file in 2005. Starting with early adopters, a proof of concept was performed during 2006, and in February 2007 .se offered DNSSEC as an additional service to its registrants (domain name customers). The aim was that .se's DNS service should not only be highly robust and available but also trustworthy.

### Planning and Development of DNSSEC

To provide DNSSEC as a service, several issues had to be considered. Many of them were the same regardless of whether the service is provided by a TLD or a small DNS Name Service Provider that just runs DNS for a few domains. Systems, policies and routines for key management and the signing of the DNS data had to be developed. When .se developed its service, the main goal was to keep the high availability of its ordinary DNS services and, at the same time offer a highly secure new DNSSEC service. Since no suitable products were available for key management and zone signing, .se had to develop its own system.

Another challenge for such a pioneer was to encourage the market to want DNSSEC. Back in 2006 .se carried out market research among its registrants and found a very positive attitude towards having DNSSEC. This attitude was confirmed in the on-going contacts and discussions with registrants. Unfortunately, it is not enough simply to have registrants with a desire to get DNSSEC and a TLD to provide it. Each registrant also needed a DNS Name Service Provider. Since the DNS and DNSSEC administration is in a distributed fashion, each registrant also needs a DNS Name Service Provider. The task for a TLD in this context is to provide the addresses to the registrant's DNS Name Servers. It is not the TLD's responsibility to handle the registrant's DNS data (for example the IP address for the registrar's internet resources such as web and e-mail servers).

The DNS Name Service Provider is the party who actually handles the registrant's DNS and DNSSEC data. Today, most registrants do not run their own DNS Name Server. They instead have an external DNS Name Service Provider, which could be a registrar, a web hosting provider or some other outsourcing partner. Not all of them offer DNSSEC today, which is a problem for the wider deployment of DNSSEC.

Because of the complexity of DNSSEC, the DNS Name Service Providers need easy-to-use and reliable administrative tools. For the deployment of DNSSEC, a good supply of commercial and open source tools is crucial. Some are already available, but more scalable, and better tools are still needed.

### Creating user value

DNSSEC is not the solution to any of the top priority security issues on the internet, like malicious code and malware such as trojans and worms, distributed through phishing and spam. Nevertheless it is an interesting new layer of infrastructure. DNSSEC increases the possibility of supporting other defence methods. Like all new infrastructure, the value increases with the number of active users.

The real value of DNSSEC is obtained when an internet user actually validates the answers from the DNS look-ups to ensure that they originate from the right source. This can be achieved in different ways. The validation is made by the users' local DNS resolver. For the ordinary internet user, a resolver is typically provided by the users' internet Service Provider, an ISP. For the Swedish DNSSEC development, it has been really encouraging to note that the major Swedish ISPs turned on DNSSEC validation at a very early stage and are validating

DNSSEC signatures for their customers.

Another conceivable use for DNSSEC is to securely store and distribute other security attributes used by other applications while using DNS as a repository. Currently there are a number of opportunities in this area, see the section “The coming year”.

In 2007 a number of ccTLDs were concerned about the slow progress of DNSSEC deployment efforts globally. They believed that the successful deployment of DNSSEC was crucial for the continued stability and security of the internet. As this was contingent upon a signed DNS root zone, they urged IANA and ICANN to speed up and improve their efforts, and migrate to a signed root zone relatively rapidly.

Fully aware that the discussions relating to the signing of the root have been taking place over the last 3-4 years, they believed that the internet had by then reached a point where the absence of a signed root zone was no longer only “merely unfortunate”. Rather, the absence of a signed root zone contributed directly to the development of inferior alternatives, thereby confusing the community and jeopardising the long-term success of DNSSEC deployment.

While .se has continued its work to make DNSSEC become a natural part of the DNS, used by all important Swedish domains and supported by useful applications, they have also worked to encourage the root zone to be signed as well.

.se, together with a number of Swedish stakeholders, sent a letter to ICANN with a number of strong recommendations. The recommendations were to make the decision to sign the root zone on a firm target date without further delay. They also urged ICANN to urgently publish a road map for reaching that target, that ICANN immediately should enter into necessary negotiations with involved parties, and finally, that ICANN should instruct IANA to take the necessary steps to implement that road map.

Which leads us to the topic of who governs the internet and the role of ICANN and IANA.

## Who governs the internet?

### ICANN

The assumingly most well-known party within internet governance is the Internet Corporation for Assigned Names and Numbers, ICANN in short. ICANN’s role is to oversee the huge and complex interconnected network of unique identifiers that allow computers on the internet to find one another.

In other words, ICANN co-ordinates the unique identifiers of the internet across the world. Without that co-ordination we cannot guarantee that we would have only one global internet. The co-ordination ensures that there are only unique domain names. ICANN co-ordinates how top-level domains can be reached and verifies that domain names are unique to avoid repetition or clashes.

In the same way that you cannot have the same domain name on one TLD (otherwise you never know where you would end up), it is also not possible for two IP addresses to be the same. In the same manner as for the DNS, ICANN co-ordinates how IP addresses are supplied to avoid repetition or clashes. ICANN is also the central repository for IP addresses, from which IP ranges are supplied to regional registries, who in turn distribute them to network providers.

This is commonly termed “universal resolvability” and means that wherever you are on the network – and hence in the world – you receive the same predictable results when you access the network. Without this, you could end up with an internet that worked entirely differently depending on your location in the globe.

ICANN was formed in 1998. It is a not-for-profit partnership of people from all over the world dedicated to keeping the internet secure, stable and interoperable. It promotes competition and develops policy on the internet’s unique identifiers.

## What about root servers?

Root servers are a different case. There is a well-spread myth that there are exactly 13 root servers. Well, that is wrong. More accurately, there are 13 IP addresses on the internet where root servers can be found. The actual servers that have one of the 13 IP addresses can be found in dozens of different physical locations. These servers all store a copy of the same file which acts as the "main index" to the internet's domain name system. It lists pointers to each top-level domain (.com, .se, etc.) where that top-level domain's authoritative name servers can be found.

Root servers are consulted relatively infrequently because once computers on the network know the address of a particular top-level domain they have the option to cache it, checking back only occasionally to make sure the address has not changed. Nonetheless, root servers remain vital for the internet's smooth functioning.

The operators of the root servers remain largely autonomous, but at the same time work with one another and with ICANN to make sure the system stays up-to-date with the internet's advances and changes.

## ICANN structure

ICANN consists of a number of different groups, each of which represent a different interest on the internet and all of which contribute to any final decisions that ICANN makes.

Three supporting organisations represent:

- The organisations that deal with IP addresses, RIRs and LIRs
- The organisations that deal with domain names, gTLDs and nTLDs
- The managers of country code top-level domains (a special exception as explained at the bottom), ccTLDs.

Four advisory committees that provide ICANN with expertise, advice and recommendations. These represent:

- Governments and international treaty organisations
- Root server operators
- Those concerned with the internet's security challenges
- The "at large" community, meaning average internet users.

And finally a Technical Liaison Group, which works with the organisations that devise the basic protocols for internet technologies.

ICANN's final decisions are made by a Board of Directors. ICANN has a President and CEO who is also a Board member and who directs the work of the ICANN staff, who are based around the globe and who help co-ordinate, manage and finally implement all the different discussions and decisions made by the supporting organisations and advisory committees. An ICANN Ombudsman acts as an independent reviewer of the work of the ICANN staff and Board.

## The ICANN decision process

When it comes to making technical changes to the internet infrastructure of domains and IP addresses, a simplified rundown of the process would be as follows.

Any issue of concern or suggested change to the existing network is typically raised within one of the supporting organisations (often following a report by one of the advisory committees) where it is discussed, and a report is produced the published for public review. If the suggested changes impact any other group within ICANN's "ecosystem", that group also reviews the suggested changes and makes its views known. The result is then put out for public review a second time.

At the end of that process, the ICANN Board is provided with a report outlining all the previous discussions and a list of recommendations. The Board discusses the matter and either approves the changes, approves some

and rejects others, rejects all of them, or sends the issue back down to one of the supporting organisations to review, often with an explanation as to what the problems are that need to be resolved before it can be approved.

The process is then rerun until all the different parts of ICANN can agree on a compromise, or the Board of Directors makes a decision on a report it is presented with.

### Internet Assigned Numbers Authority (IANA)

IANA is quite likely the oldest internet institution, first documented in 1972, responsible for the global coordination of the internet's unique names and numbers. Since 1998, IANA has been a service provided by ICANN. Previously it was operated in academia under US Government research contracts. Nowadays it run by the PTI department of ICANN. The domain name administration may be a small component, but it is really important and highly visible. Every change in the root zone used to be rubberstamped by the DOC, but a transition took place (30 September 2016) where the US Government decided to hand over full control over the root zone to ICANN.

Essentially IANA is a technical maintenance function. It keeps track of technical delegation details to keep in the DNS root zone, is responsible for the monitoring and coordination of the effective functioning of the DNS, arranges the consultation and research on technical functions and provides neutral services to all TLD managers. This is done independently from their direct involvement in ICANN, and only applies to technical matters; non-technical decisions are not taken by IANA.

## DNS as security enhancer

Despite the fact that DNS puritans were against a broader use of the DNS as a distribution mechanism for other records than the traditional ones, the DNS has increasingly become a repository for security attributes such as certificates, records for secure e-mail, encryption keys and so on. The main reason for that is the possibility to protect the content of the zone file with DNSSEC. Below we give you some examples of how this may be used.

Everybody with an email account has received their fair share of spam. Unfortunately, it is unavoidable. However, over the years the technical community has come up with some clever solutions to at least make it easier to identify spam and make sure that your domain name cannot be abused for unknowingly sending it.

The Sender Policy Framework (SPF) was the first e-mail security protocol to be introduced. It is a simple mechanism to tell the world which name servers are allowed to send mail from your domain. SPF is defined in RFC7208.

Next, DomainKeys Identified Mail (DKIM) introduced signing with cryptographic keys to prove that the domain has been sent through authorized servers either by adding it to the mail header, the body or the entire message. It is defined in RFC 5585, 6376, 5863 and 5617. With DNSSEC, the public keys are protected from man-in-the-middle (MitM) attacks.

And finally, Domain-based Message Authentication, Reporting and Conformance (DMARC) allows you to receive information if someone tries to send messages in your name. DMARC requires both SPF and DKIM to be implemented. All three are configured via records in the DNS with public keys and consequently need protection by DNSSEC to make sure that no-one can do a MitM attack.

The original SMTP standard, RFC821, published in 1982, did send all email messages in clear text over the internet. To protect emails in transit, the STARTTLS standard was introduced. Unfortunately, this standard is not secure itself and some ISPs have used this fact to stop email encryption in their network. The problem is that the sender does not know beforehand if the server supports encryption. And if all that fails the sender will fall back to clear text. DNS-based Authentication of Named Entities (DANE) came to the rescue. DANE records in the DNS, protected by DNSSEC will clearly indicate that a server is able to communicate in an encrypted way. And for good measure it usually even identifies the certificate used.

As you can see, email security has been enhanced with the help of the DNS in combination with DNSSEC in many ways.

## The future of the DNS

Currently the DNS is considered to be one of the basic building blocks of the internet. No DNS service means, for almost all users, no internet. As they say, nothing lasts forever, but how long will the DNS last?

### Near-term changes

Currently the very base of the DNS, the DNS transport, is under reconstruction. If you have followed the discussion you are familiar with DoH (DNS Over HTTPS), DoT (DNS over TLS) and the looming DoQ (DNS Over Quic).

The DNS has from the beginning been a very open protocol. All data is sent in clear text over the internet. That is true for DNSSEC as well, but in this case the data sent is also digitally signed. All data stored in the publicly-accessible DNS should be considered public data. That does not mean that whoever accesses the data should be public too. As such, lots of effort has been put into making the DNS more privacy-friendly.

DNS over TLS (DoT) and DNS over HTTPS (DoH) are both encrypted and break with many DNS traditions. The DNS has been run mainly on UDP. So much so that over the years many firewalls have decided to block DNS over TCP (which is wrong according to current RFCs). Both DoT and DoH connect via TCP to a resolver, and thereafter a TLS session is started. DoT still transmits good old DNS packets. DoH talks HTTP. Both technologies introduce massive changes to how DNS is done. The most discussed change is the introduction of “trusted resolvers”, a concept that is actually used by DoT and promoted by the browser vendors.

The DNS does not stand still. There has not been one IETF meeting where new proposals for future DNS functionality are not discussed. In fact, DNSOP is one of the IETF’s longest-running working groups and it does not show any signs of slowing down.

### Platform economy

A threat that has been looming for many years now is the obsolescence of domain names through the big platform providers, think Google, Facebook or Amazon. Who needs a domain name if you could be found by Google without one? Many organizations opt not to have their own domain, and instead only have a Facebook account. In fact, a large number of e-merchants do not have their own shop and are Amazon merchants only.

The verdict on this business model is still out. Many merchants will do both, being present on a platform and running their own shop. The platforms are under heavy discussion right now. Unfortunately, there is no hard data on how many domains are not registered because of the presence of the platforms.

### Alternative Name Services

Over the years many attempts have been made with alternative name services. From alternative DNS root services to name services with other technologies, nothing has been able to threaten the dominance of the DNS as we know it. And to be honest, we do not expect anything to come in the nearby future. Nevertheless, we will go through some of the most well-known attempts below.

### OpenNIC

A long running alternative root. Over the years it has had some usage, but the vast majority of users have no access to OpenNIC names. OpenNIC has even included NameCoin and other alternative name spaces in their resolvers.

### NameCoin

The longest running and still active blockchain for DNS names is NameCoin. It builds on Bitcoin technology, so much so that miners can mine Bitcoin and NameCoin at the same time. All registered names go under the TLD .bit, which in fact is not recognized by ICANN or the IETF. This is exactly the biggest problem of NameCoin; the majority of internet users does not have access to the .bit TLD.

### Blockchain Name Services

Blockchains are all the hype currently, and even in this space stable names are needed. Many blockchains have started their own name services which all only work on their respective chain. Often the goal is not an all-purpose name service, but a simple name to blockchain address translation.

The Ethereum Name Service (ENS) builds on DNS specifications of allowed Unicode code points in names but uses its own format to store names on the blockchain. It allows the forward- and backward-resolving of names and addresses.

As you can see, there is a lot of movement in the space of name systems and the coming years will bring big changes for all parties involved.

**Council of European National Top-Level Domain Registries**

CENTR is the association of European country code top-level domain (ccTLD) registries, such as .de for Germany or .si for Slovenia. CENTR currently counts 54 full and 9 associate members – together, they are responsible for over 80% of all registered domain names worldwide. The objectives of CENTR are to promote and participate in the development of high standards and best practices among ccTLD registries. Full membership is open to organisations, corporate bodies or individuals that operate a country code top level domain registry.

This paper is part of a series of articles covering industry research, historical data analysis and the future of technologies such as digital IDs, published over the course of 2019 to mark CENTR's 20th Anniversary. These publications do not necessarily present the views of CENTR or of the CENTR community.

*CENTR wishes to thank and acknowledge the organisations which have so generously contributed to the efforts of its 20th Anniversary:*

**Platinum sponsor**

**Gold sponsor**

**Silver sponsor**

CENTR vzw/asbl
Belliardstraat 20 (6th floor)
1040 Brussels, Belgium
Tel: +32 2 627 5550
Fax: +32 2 627 5559

*To keep up-to-date with CENTR activities and reports, follow us on Twitter, Facebook or LinkedIn*