DNS Service Discovery

Proposal for BDI

Topsector Logistiek

Colophon

DNS Service Discovery
Proposal for BDI

Author
Remco van 't Veer
remco@jomco.nl

December 2022

Index

1 Basic Data Infrastructure

The Basic Data Infrastructure (BDI) framework is a specific implementation of federated data sharing principles focused on coordination challenges in sectors like logistics, construction, engineering & contracting.

The BDI is a specific implementation of the generic approach of 'federative data sharing' in so called data spaces. Federative data sharing is widely considered to be an attractive option to address the challenges for fully exploiting the business potential for the emerging data economy.

The FEDeRATED framework and the BDI which is derived from the FEDeRATED concepts are focused on challenges in the physical economy where coordination between many actors and proof of compliance in Business-to-Business (B2B) and Business-to-Government (B2G) interactions are dominant drivers. For example in global supply chains with import-export controls, in construction, or in engineering & contracting.

One of the challenges in sectors that are diverse and international is how to organize machine-to-machine discovery of parties that adhere to the framework, of the availability (if any) of IT-based services as defined in the framework, and the location of resources.

This study proposes a mechanism based upon the existing Domain Name Service (DNS) of the Internet.

2 Service Discovery

When parties want to share data, it is crucial that they can easily find and access each other's data. This requires a mutual agreement on what to share, how to share it, and the location of the data.

This study proposes Service Discovery - a mechanism allowing systems to easily find each other - using DNS[1].

[1] *Domain Name System, the distributed naming system for computers*

3 DNS

DNS is one of the oldest internet standards still in use. Organisations use it to make a part of their infrastructure publicly accessible. Think of email and websites. DNS is also the backbone of many federated services including VoIP[2], LDAP[3] and XMPP[4] (email is a federated service as well).

A DNS zone[5] may be administered by the organisation itself, or the administration may be outsourced. It gives them free reign in choosing which 'names' constitute 'their part of the internet'. Each zone can be subdivided into 'subzones', administered by their own departments, thereby creating a hierarchical structure.

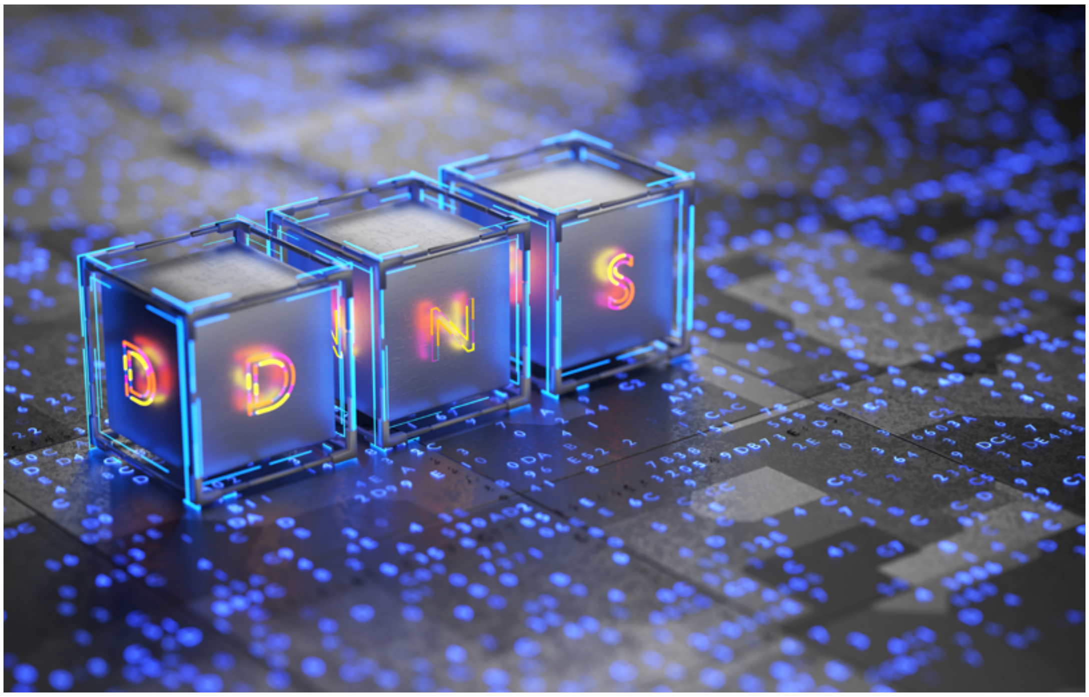

[2] *Voice over IP, internet telephony*
[3] *Lightweight Directory Access Protocol, for finding user information*
[4] *Extensible Messaging and Presence Protocol, often used for chat and IoT applications*
[5] *A specific portion of the DNS namespace*

4 How?

Email works (basically) as follows. For each domain with associated email addresses, the organisation will configure MX-records[6] pointing to servers that can receive email. When an email is sent, a DNS server is queried as to which *MX*-records exist in the receiving domain, after which the server is contacted and the email delivered. It is also possible for *TXT*-records[7] to have been created, containing cryptographic keys used in authenticating the email's origin.

DNS record types are standardised. *MX*-records are only used for email and nothing else. Fortunately, *TXT*-records allow much freedom and multiple records can be registered with the same name, allowing lists of information to be queried. For email, this is what happens when the first server to which an *MX*-record points is unavailable, after which the next server on the list is tried.

Below, we detail the various components of service discovery. Then, in section 4.5, we provide a concrete example of an organisation hosting several services.

The ingredients required for service discovery are as follows:

4.1 Well-known subdomain

By introducing a 'well-known subdomain', we can add a point with a predictable name at which the available services can be queried: Service Discovery.
For example: `_bdi.acme-corp.com.` This can be the root level of a DNS zone and assigned to the appropriate administrators.
Please note that this is a *domain name* rather than a *host name*. A host name, as used for websites and email, may not contain the underscore symbol ('_'). Using the underscore indicates this to be a 'special' record.

4.2 Discovery

Now that we have a starting point, we can configure it for discovery. We use *TXT*-records as these give us complete freedom in adding additional 'special' records to our 'well-known sub domain'. We adhere to these formats:

- `_bdi.<subdomain>` *TXT*-records
  Provides a list of all services for the subdomain, using the format below

- `_<service>._<proto>._bdi.<subdomain>` *TXT*-records
  With *proto* the internet protocol used by the service (tcp or udp), and *service* the type of service (ldap, mqtt, etc.)
  It returns a list of services of that type (instances), using the format below

- `<instance>._<service>._<proto>._bdi.<subdomain>` *TXT*-records
  With *instance* a name that can be freely chosen

[6] *Mail exchanger records*
[7] *TXT-records contain 'free' text notes*

4.3 SRV-records

We now have knowledge of the services available, but we don't yet know where to find them. After all, underscores are not allowed in host names, and it is inadvisable to create direct *A*-records[8] or *CNAME* -records[9] for this server. We use *SRV*-records[10] instead.
These contain:

- *target*
  The hostname or IP address of the service

- *port*
  The port number of the service

- *priority* and *weight*
  Values used to select a record when multiple endpoints are available for the service

Should multiple *SRV*-records exist for a service, then each of the endpoints should point to the same data. The *priority* and *weight* values are used to select the best candidate from a list (see RFC 2782 for more information).

We now know where to find the information, the protocol to use, the server and the port number.

4.4 TXT-records

For some services, the information in the *SRV*-record may not suffice. The solution is to create additional *TXT*-records that have the same name as the *SRV*-record. The content of these *TXT*-records depends on the specific service and protocol - for example, the necessary rights for using the service.

Information in these records consists of attributes that have a name and a value, separated by an '=' symbol, with attributes separated by a ';' symbol. Although all (special) characters are allowed in *TXT*-records, DNS service providers often limit their use. Unfortunately, there is no consistency between providers. Most of them will, however, support all normal letters, digits, spaces, and the '+ / = ;' symbols[11].
An example:

```
quality=high; resolution=seconds
```

[8] *Address record, containing the IP address of a hostname*
[9] *Common Name Record, the alias of a hostname pointing to another hostname.*
[10] *Service records specifies the location of a server*
[11] *Letters and digits means standard ASCII letters and digits.*

4.5 An example

We'll use the ACME Corporation as an example.

They have the 'well-known subdomain':

- `_bdi.acme-corp.com`

The associated TXT-records let us know that they have SPARQL and MQTT endpoints

- `_bdi.acme-corp.com. 3600 IN TXT`[12]
  `'_sparql._tcp._bdi.acme-corp.com'`
- `_bdi.acme-corp.com. 3600 IN TXT`
  `'_mqtt._tcp._bdi.acme-corp.com'`

We zoom in on MQTT to find:

- `_mqtt._tcp._bdi.acme-corp.com. 3600 IN TXT`
  `'warehouse-status-events._mqtt._tcp._bdi.acme-corp.com'`
- `_mqtt._tcp._bdi.acme-corp.com. 3600 IN TXT`
  `'logistic-events._mqtt._tcp._bdi.acme-corp.com'`

We find the server with the SRV-record:

- `logistic-events._mqtt._tcp._bdi.acme-corp.com. 3600 IN SRV`
  `'1 1 443 mqqt.logistics.acme-corp.com'`[13]
- `logistic-events._mqtt._tcp._bdi.acme-corp.com. 3600 IN SRV`
  `'2 1 443 mqqt-backup.logistics.acme-corp.com'`

And the additional information:

- `logistic-events._mqtt._tcp._bdi.acme-corp.com. 3600 IN TXT`
  `'queue=main; auth=ishare'`

We now know that, for *MQTT* messages, we must contact *mqqt.logistics.acme-corp.com* (and when not available, to contact *mqqt-backup.logistics.acme- corp.com*) at port number 433, and that everything of interest can be found at queue main, and that authorisation is handled through *iSHARE*.

4.6 Origin

The above is loosely based on DNS-SD[14]. This standard can, however, not serve our goal as it uses *PTR*-records that are very impractical in use. Only few DNS service providers support *PTR*-records as they are mainly used in dealing with unwanted emails. It also presumes more control over the records than allowed by providers: such as the characters used in names and values.

[12] The number 3600 indicates that the record can be stored for one hour (3600 seconds)

[13] The numbers 1, 1 and 433 denote the priority, weight, and port of the service, respectively

[14] See also RFC 6763

5 Security

The DNS platform is secured by means of DNSSEC[15]. Use of DNSSEC is a prerequisite for the safe implementation of service discovery as described above. DNSSEC is a widely recognised standard supported by all DNS service providers.

[15] *Domain Name System Security Extensions - it prevents other entities from impersonating the compromised organisation, thereby providing false information or stealing login credentials.*

6 In conclusion

The mechanism as described is an initial effort towards a service discovery system, meaning that it needs further development. Especially section 4.4, describing the use of *TXT*-records, can be characterised in more detail for each *service type*. And we probably need to come to agreement regarding *instance* names in section 4.2.

As DNS is important for the proper functioning of the internet, it has been tried and tested time and again. For several decades, various federated services have successfully used DNS as a register and for implementing some type of service discovery. Using zones, it is easy to distribute the administrative effort hierarchically. These features make DNS excellently suited for service discovery.

Topsector Logistiek
Ezelsveldlaan 59
2611 RV Delft
+31 15 251 65 65
www.topsectorlogistiek.nl

Topsector
Logistiek