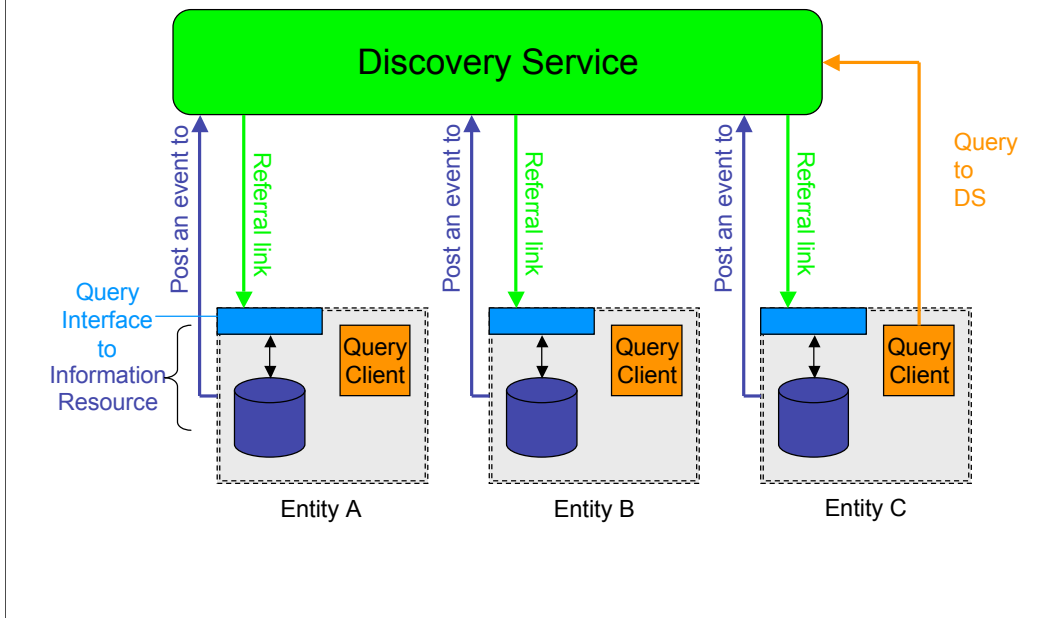


Simple scenario: multiple organizations use the same DS



- Discovery Services provide a lightweight referral service to enable the gathering of information that may be held by multiple entities. Examples of entities may include companies within a supply chain or across the lifecycle of a product but may also include other kinds or organizations and even individuals who wish to announce that they hold information (including photos, videos, text, commentary, reviews, ratings) on a particular topic.
- The information is considered to be fragmented because parts of it are held by multiple entities rather than being centralized. However, there is a logical connection between these fragments, in the sense that they are all related to a common topic. The common topic could be a particular keyword or discussion topic or it could be the individual identity of a physical object or information object.


We can envisage at least two main modes of interaction between an entity and a Discovery Service:

- 1) publishing: an entity (acting as a 'publisher') announces (posts) that it holds some information about a particular topic or ID
- 2) querying: an entity (acting as a 'client') makes a request to a Discovery Service to seek the addresses of other entities that hold information for a particular topic or ID.

It is important to note that the following are assumed to be outside of the scope of a Discovery Service:

- a) Definition of the data model and query interface that is accessed via a URL referral provided by a Discovery Service
- b) Performing of proxy requests in order to gather information from other entities

## Complex scenario: federation of multiple Discovery Services



- For reasons of scalability and performance, it may not be practical to rely on a single Discovery Service for all kinds of topics. There may be appropriate ways to partition the lookup/referral problem into a number of Discovery Services that cater for specific clusters of topics, families of identifiers, different industry sectors or different geographic regions.
- However, if we assume such a partitioning into multiple complementary Discovery Services, there may still be a need for co-operation among multiple Discovery Services in a federation, specifically to assist a client entity that initially only knows of one Discovery Service - but unfortunately not the relevant one that can provide a referral.
- To support such co-operation, we may also need to consider a 'bootstrap protocol' that allows Discovery Services within a federation to 1) be aware of each other's existence, URL address and topical scope; 2) help a client locate a relevant Discovery Service. It may be possible to achieve this via appropriate indexing and lookup of the scope of candidate Discovery Services, based on the semantic classification of the topic or the structure of the identifier. Failing that, multi-cast broadcasts could be considered among a peer-peer network of Discovery Services to seek a relevant DS.
- The client entity may be able to contact the relevant DS directly - or we may need to consider a federated trust model in which their well-known/regular DS provider (e.g. DS 2 for entity X) provides a security token that vouches for them when they interact with other Discovery Services within the federation.