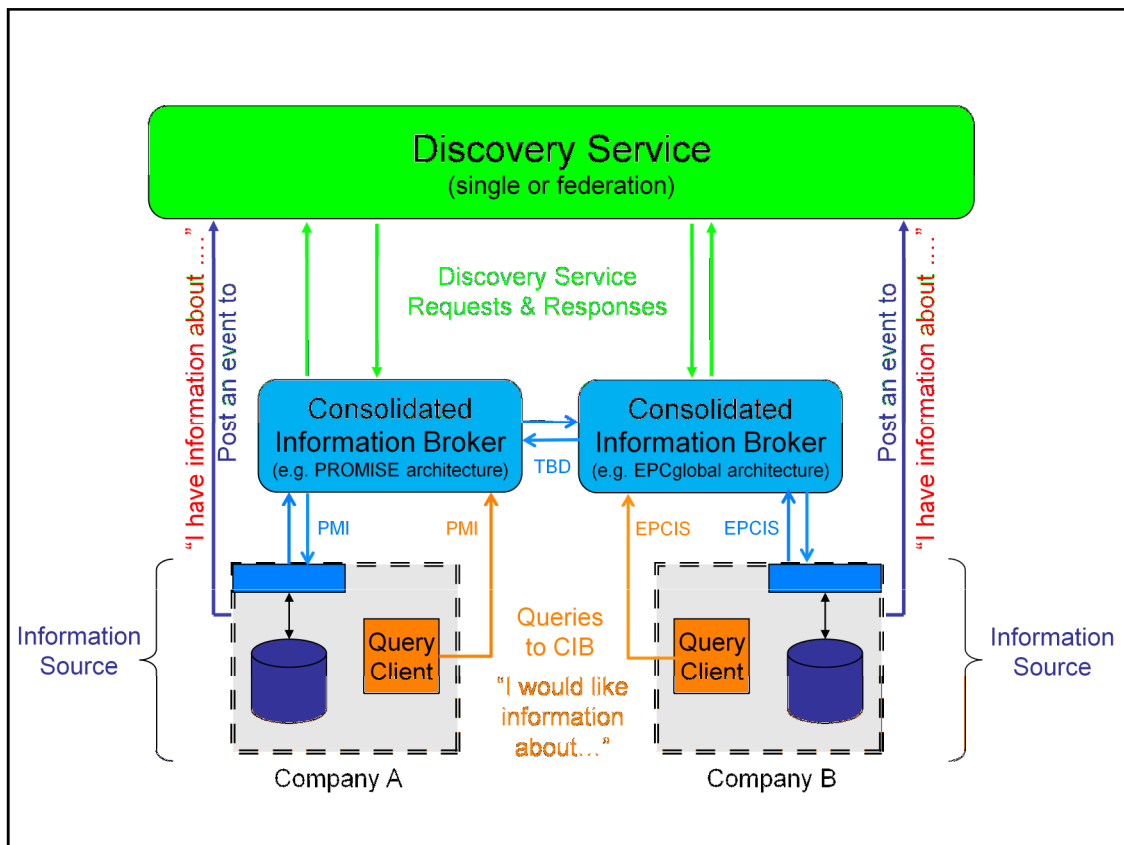


Let me first introduce the main components in this diagram:

1. Company A and Company B are shown as "Information Sources" (we can change any of the terms as necessary) but they are also "Information Consumers", i.e. they can request information as well as supply information. Companies A and/or B could be manufacturers, logistics companies, retailers, service providers, maintainers, dismantlers, recyclers, government agencies; in fact they could be any kind of organization that has a legitimate right to supply and request information, subject of course to access rights and authentication.
2. The two cyan boxes I have called tentatively "Consolidated Information Brokers", or one might call them "information services" (but let's find something appropriate). This kind of service might be provided in-house by large enterprises, but in the case of SMEs, it might be outsourced to a service provider. I see this kind of service as a way to insulate the provider and requester of information from the complexities of dealing potentially with multiple information sources and multiple discovery service providers. I have suggested that a CIB might follow an architecture like EPCglobal or PROMISE, but it might support both, and of course there might be other architectures to consider. The diagram suggests an interconnection between two CIBs that support different architectures, the intention being that Company A and Company B may choose one CIB with which to interact (using either EPCIS or PMI (PROMISE Messaging Interface)), and if some of the data might exist in one or other format, then the CIB might handle the transformation, making it a simpler job for Company A or B, who then needs only to deal with one information request/response interface.



3. The "Discovery Service", single or federated, is an organization that will eventually use the results from ESDS to supply "information about information", though probably not strictly metadata in this case.

For me, the most important intention at the moment is to convey the RELATIONSHIPS between these components rather than convey any more precise connection (interfaces etc.).

I see an advantage to create a separation between the supply and usage of "discovery services" and the "information services" (the CIBs) that use the "discovery services" in order to consolidate information requests on behalf of "users" such as Company A and B.

I see a simpler way to manage trust and authentication if the Discovery Service is able to identify and control its authorised users. The DS should be able to accept publishing events ONLY from recognized publishers with valid certificates. It should be impossible for an unauthorized publisher to introduce counterfeit "referral" data.

Furthermore, if only CIBs are allowed to use the DS to "discover" information on behalf of its users (Companies A & B), then it too can reject requests from users without appropriate credentials, it can limit their access to only data for which they have certificated access rights, and it should prevent "crawlers" from misusing tracking information, since only legitimate and rigorously certificated CIBs should have "unrestricted" use of the DS.

As shown in my tentative diagram, the "Information Sources" (Companies A and B) publish the existence of data ("I have information about..."). According to the eventual functionality of ESDS, this could be "info about unique ID", it could be more granular and list specific data items about unique ID, or it could be information about a series.

As requesters of data, Companies A & B are only permitted to request data via a CIB. Their requests can be managed by the CIB according to access rights certificates that they own, and the CIB can "conceal" all other data to which the requester is not entitled, but more importantly it conceals ALL "discovery" information from the information requester.

That's enough high-level proposition for now... I look forward to all comments and reactions...