ABSTRACT: The primary goal of the Semantic Web is to use URIs as a universal space to name anything, expanding from using URIs for webpages to URIs for “real ob jects and imaginary concepts” as put by Berners-Lee. This distinction has often been tied to the distinction between information resources, like webpages and multimedia files, and non-information resources, which are everything from real-people and bound books. Furthermore, the W3C has recommended not to use the same URI for information resources and non-information resources, and several communities like the Linked Data initiative are deploying this principle. Yet, the definition put forward by the W3C, that non-information resources are things whose “essential nature is information” seems confusing at best. For example, would the text of Moby Dick be an information resource? While this problem could safely be ignored up till now as a mere distraction, with both the rise of Linked Data, projects like OKKAM, and recent work on modelling HTTP in RDF for error reports, it appears that this problem must be modelled formally. An ontology called IRW (Identity and Reference on the Web) of various types of resources and their relationships, both for the hypertext Web and the Semantic Web, is presented. It builds upon Information Object Lite (an extension of DOLCE Ultra Lite for describing information objects) and IRE and aligns with the work of the W3C in modelling HTTP in RDF and Berners-Lee’s Tabulator ontology. It can be used as a tool to make the Semantic Web more self-describing, and it also allows inference to be used to test for membership in various classes of resource, and so for validating Linked Data.