

ABSTRACT: The Entity Name System is a service which aims at providing globally unique URIs for all kinds of real-world entities such as persons, locations and products, based on descriptions of these entities. Because entity descriptions available to the ENS for deciding on entity identity—Do two entity description refer to the same real-world entity?—are changing over time, the system sometimes has to revise its past decisions: One entity has been given two different URIs or two entities have been attributed the same URI. The question we have to investigate in this context is then: How do we propagate entity decision revisions to the clients, which make use of the URIs provided by the ENS? In this paper we propose a solution which relies on labelling the URIs with additional history information. These labels allows clients to locally detect deprecated URIs they are using and also merge IDs referring to the same real-world entity without needing to consult the ENS. Making update requests to the ENS for the detected URIs only considerably reduces the number of update requests, at the cost of a decrease in uniqueness quality. We investigate how much the number of update requests decreases using URI history labeling, as well as how this impacts the uniqueness of the URIs on the client. For the experiments we use both artificially generated entity revision histories as well as a real case study based on the revision history of the Dutch and Simple English Wikipedia.