Target practice #6.
Ontologically neutral HTTP Semantics - noodling - JAR 2009-06-13

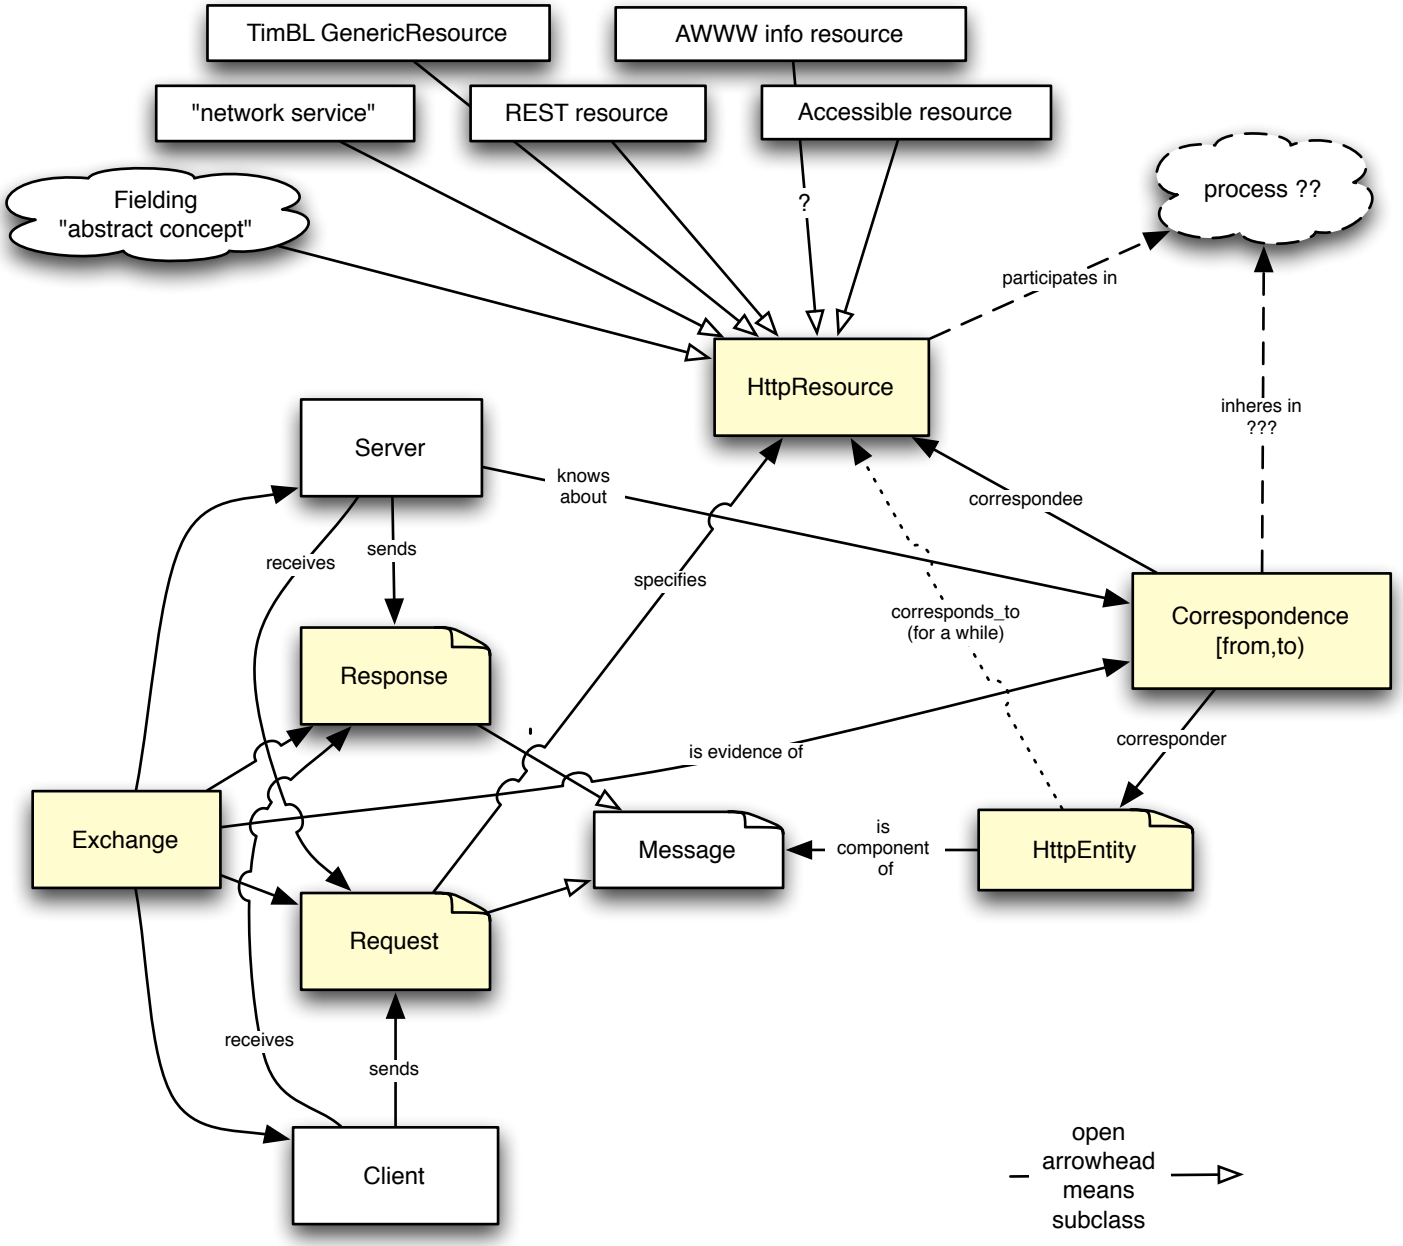

The server knows somehow about a current correspondence between the "requested resource" and some entity, so it reports on this correspondence by placing that entity (including perhaps last-modified and expires times) in the response it sends back to the client. (If it knows several correspondences it chooses one somehow - CN, User-agent, cookies, whim, etc.) Other servers may know different correspondences for the same resource.

The correspondence reflects (is a quality of?) of some situation (process?) involving the resource and possibly other factors (we can't rule those out). Those factors might include the server, as when the resource "resides on" the server, or on some independent resource involved in encoding and/or decoding correspondences.

The shaded boxes are the more important ones.