

Scalability

- Background
 - W3C discovers that HTML DTD, etc. are being accessed up to 500M times/day, often from same client
- Architectural issue
 - There is no normative specification that prohibits or even discourages repeated access to the same, cacheable, Web resource
 - RFC 2616: “*provide an explicit expiration time [...] indicating that a response MAY be used to satisfy subsequent requests*”
- Practical issues:
 - Large organizations can’t find the software that’s doing it, may not be in a position to rewrite/reconfigure code
 - Blocking miscreants: can disable access from entire organizations
 - Tell them to proxy: large corporations report expense of proxying all Web access can be >> savings
 - “Tarpitting”: when repeated access to a given resource is seen from given source, delay before responding – W3C server team reports this is the best compromise so far