

CSE 551
Design Exercise #2
Robust, Highly Available Web Service

First draft due: noon, Thursday, November 4, 2010

Final draft: 4:30pm, Thursday, November 11, 2010

In class, we have discussed various mechanisms to achieve robustness and high availability, such as transactions, write-ahead logging, RAID, and consensus. This assignment asks you to synthesize those ideas into a coherent system design for a new e-commerce web site, PhDs.com. Targeted at students who want the benefits of increased education without the bother of studying, PhDs.com grants advanced degrees online, after a modest financial contribution to the school. Demand is expected to be brisk! To maximize profit, the web site should be scalable (as demand increases with word of mouth), always available (we don't want to turn anyone away), and it should grant exactly one degree per payment (that is, it should not grant degrees without payment, nor multiple degrees for the same payment). Degrees should be durable, of course – once granted, potential employers should be able to easily validate student achievement. Happy customers make repeat customers, so multiple PhDs per student are possible, but not in the same discipline (how many CSE PhDs does one need, after all).

Please focus on the back end web site(s). You may assume that a separate, highly available system is provided for translating names (e.g., “PhDs.com”) to a set of IP addresses (e.g., one per geographic location). In truth, however, name translation is much the same problem as the one described here, so your answer will do double duty.

In addition, we would like you to describe how you would quantitatively evaluate your design. What principal experiment would you run, and what would you expect it to show?

We are interested in a sketch of your general approach, rather than details -- e.g., try to keep the writeup to 1-2 pages for the first draft, 2-3 pages for the final draft.