



UCL Department of Computer Science  
CS M038/GZ06: Mobile and Cloud Computing  
2010–2011, Term 2  
Kyle Jamieson and Brad Karp

One-pager: Xen (Barham et al., 2010)

Due: Start of lecture, 11th March 2011

*Instructions: in your own words, answer the following question as succinctly as possible (in 200–500 words, but shorter answers within this range are encouraged). Quoting figures or text from the assigned reading or from any other source is specifically prohibited.*

In Section 4.5, the authors argue that Xen is very memory-efficient when running many domains (guest OSes) concurrently, and that one can easily run 100 guest OSes on a single server without running out of memory.

Suppose Xen were being used to host 100 virtual machines, each of which runs XenLinux and the OSDB-IR database benchmark described in Section 4.1. Would this configuration likely scale in memory to 100 virtual machines when running atop the Xen design described in the paper? Why or why not? How about if the 100 virtual machines each ran XenLinux and the SPEC INT2000 CPU benchmark described in Section 4.1? Why or why not? Be specific in your answer: if there is some type of data that will consume a lot of memory, state what this data is and why it consumes a lot of data. Justify your answer in light of the characteristics of the benchmarks.