

## Questions for Geiger

### General

- What is the semantic gap?
- Which types of OS services should instead migrate into a VMM?

### Background

- Why/when is estimating working-set size useful?
- How does Geiger go beyond what ESX server does?
- Why is implementing a second-level cache challenging?
- What kinds of policies lead to poor outcomes?

### Mem Techniques

- Geiger monitors cache promotions and evictions; what are they?
- What are some of the hard cases for a VMM to monitor?
- What is stored in the ADL (associated disk location) structure?  
Why is this important?
- Why is page reuse important for Geiger?
- How does Geiger detect reuse?
- Why are COW pages important? How does Geiger detect COW activity?
- Overall, what types of VM events need to be visible to Geiger for it to function?

### FS Techniques

- What is the aliasing problem? How does it arise for a journaling file system?
- How does Geiger handle this problem?
- Are there other types of file systems that would be more difficult to handle?
- What is the block liveness problem?
- Why is it hard to figure out block liveness by simply monitoring disk traffic?
- How does Geiger gain more precise information about block liveness?
- How general is this technique?

### Evaluation

- What causes lag in Geiger's page eviction detection?
- Is the journal alias detector needed?
- What technique was important for real applications?
- Which overheads of Geiger are the most conceding?

### Case Studies (and you thought the evaluation was the end!)

- How does MemRx figure out a thrashing VM's working set size?
- What assumptions are made about OS replacement policy?
- Is MemRx effective?
- Is Eviction-based cache replacement effective?