

Distributed Systems

Remzi Arpaci-Dusseau (Wisconsin)
Wyatt Lloyd (USC)
and The Entire Group

Requirements (1)

Workloads

- Benchmarks
- Traces

Reproducibility

- Mechanisms to enable
- How to quantify level of reproducibility?

Requirements (2)

Platform diversity

- Hardware: CPU, network, storage
- Software: OpenStack, Eucalyptus, etc.
- Local and wide area resources
- Need for some (limited) uniformity across data centers?

Resource controls

- Host: Rack-level, datacenter-level
- Network/storage resources

Requirements (3)

Fault injection

- Ability to mimic realistic failure scenarios (machines, disks, networks)

Isolation levels

- Bare metal, intermediate forms?
- Make contention information available?

Requirements (4)

Running at scale

- Interactively (to develop)
- Batch (for final #s)

Instrumentation and monitoring

- Different granularities
- Ways to integrate different data sources
- Tracing and archival
- Use the infrastructure as something to study

Requirements (5)

Ease of development

- Easy to make simple experiment
- Easy to scale experiment

Incentives and expectations

- What can I expect to get, when?
- Why should I play nice?

Requirements (6)

Mechanisms for sharing (or, Community building)

- Course materials
- Experiments
- Case studies
- Applications
- Everything