

DISTRIBUTED

TRACING -2

NOTE ABOUT THIS WEEK'S TASK


- NOT ABOUT TRACING: INSTEAD FOCUS ON SETTING UP INFRASTRUCTURE
 - ↳ USEFUL FOR THE FUTURE: WE WILL BUILD ON THIS/USE IT
- RISK ° MIGHT FIND IT DOESN'T RUN ON MACHINES YOU CAN ACCESS

THAT IS FINE ° POST A QUESTION ON CAMPUSWIRE

↳ I WILL POST A LINK TO JOIN CLOUDLAB, AND INSTRUCTIONS.

APOLOGY FOR TODAY

DISTRIBUTED TRACING FROM THE TOP



DECISIONS

- SAMPLE AT TRACE GRANULARITY

- SAMPLING PROCEDURES

$\frac{1}{1000}$

$\frac{1}{10000}$

- ~~Need~~


Uniform

1 in 10000

- Reservoir -

PROBLEMS / CONCERNS WITH SAMPLING

o TAIL BEHAVIOR



- GARBAGE COLLECTION

- SSD COMPACTION

- UPDATES & PERIODIC TASKS

- ...


o DESIRABLE TRACING BEHAVIOR?


Trace the tail.

o OTHER EDGE CASES

o DAPPER ASSUMPTION FOR TAKE BEHAVIOR

ALTERNATES (HINDSIGHT)





DAPPER VS HINDSIGHT OVERHEADS.

ERGONOMICS

↳ MINIMIZE EFFORT NEEDED TO ADD TRACING.

WHY?
○

How?

ALTERNATE APPROACHES

- ~~FAT~~ / DTRACE / wPROBES

Philosophical
PRO
More versatile(?)
CON
More work(?)
← → Less ver.

- Challenges to using these approaches

