

Hard Drive Failures in a Large Datacenter

by Brendan Sweeney

CSS458, Spring 2013

System

The operating hard drives of a large datacenter. Hard drives are generally grouped in replicated volumes within a SAN so that several failures may occur before any data loss results.



Tools

The system is written and run in Python 2.7.



SimPy 2.3 is used to provide the simulation framework.

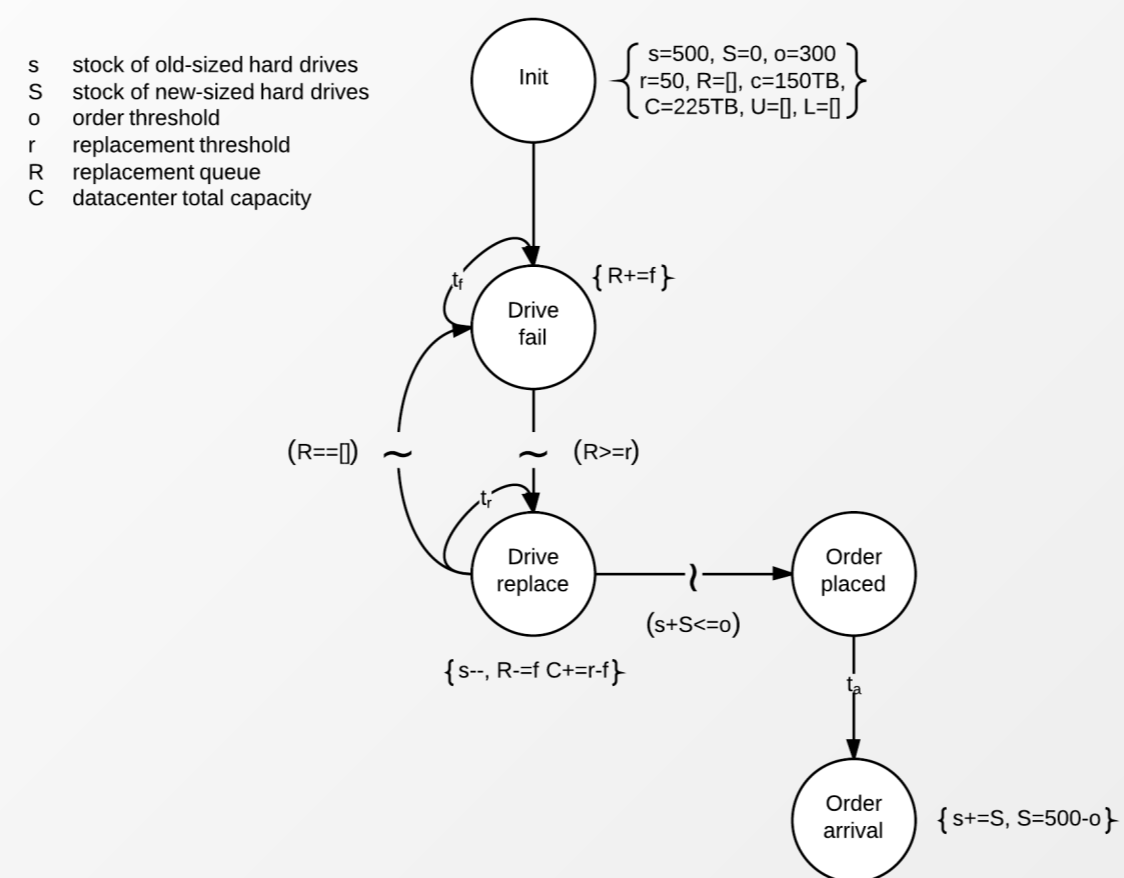
Statistical analysis and presentation are made within LibreOffice 3.6.



Problem

Given a large datacenter with thousands of hard drives in operation, what is the best replacement policy for failed drives that will prevent technicians from spending too much time on the task while avoiding irreparable loss of data?

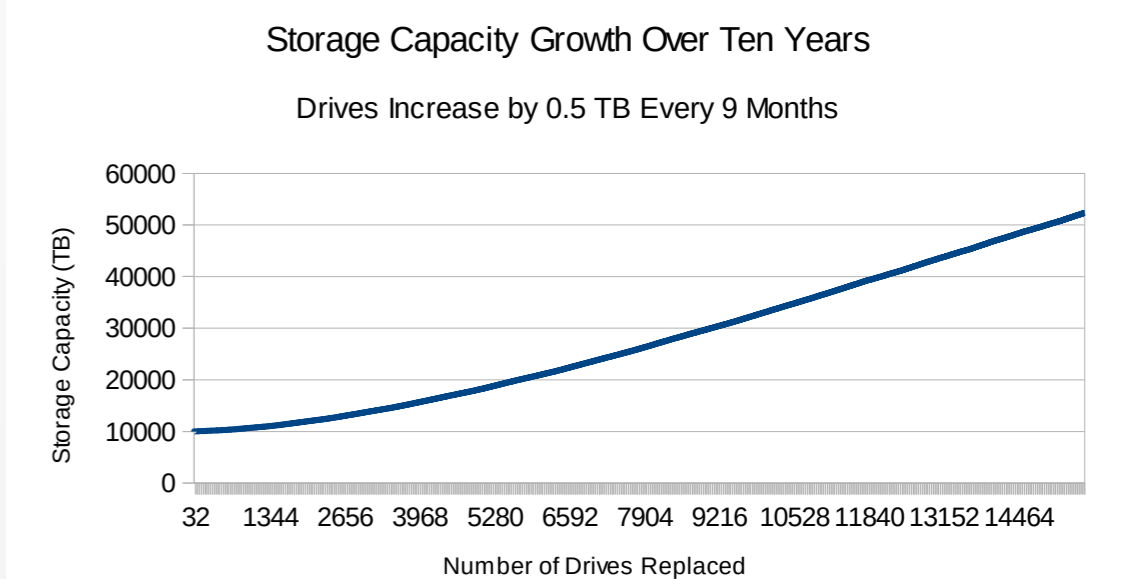
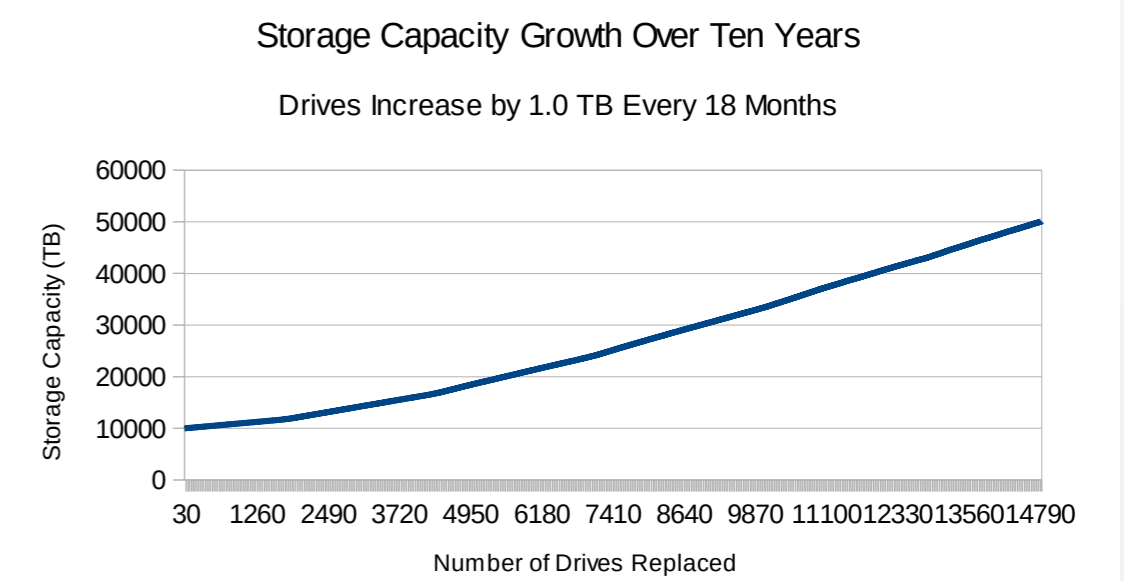
Model



When hard drives fail, they wait for a replacement event to occur, which happens following a certain number of failures. All failed drives are replaced in a single session with the smallest drives available in storage. When drives in storage are replenished, they may have higher capacity than the previous batch. Consequently, the total capacity increases over time as smaller hard drives fail and larger hard drives replace them.

Results

The proportion of capacity increases to frequency determines the rate of total capacity growth.



References

- Pinheiro, Eduardo, Wolf-Dietrich Weber, and Luiz André Barroso. "Failure trends in a large disk drive population." Proceedings of the 5th USENIX conference on File and Storage Technologies. 2007.
- Documentation, Official Python. "Python Tutorial, The." Python Software Foundation (2008).
- Vignaux, Tony, Klaus Muller, and Bob Helmbold. "SimPy Manual." SimPy 2.3.1 Documentation. 2012.