

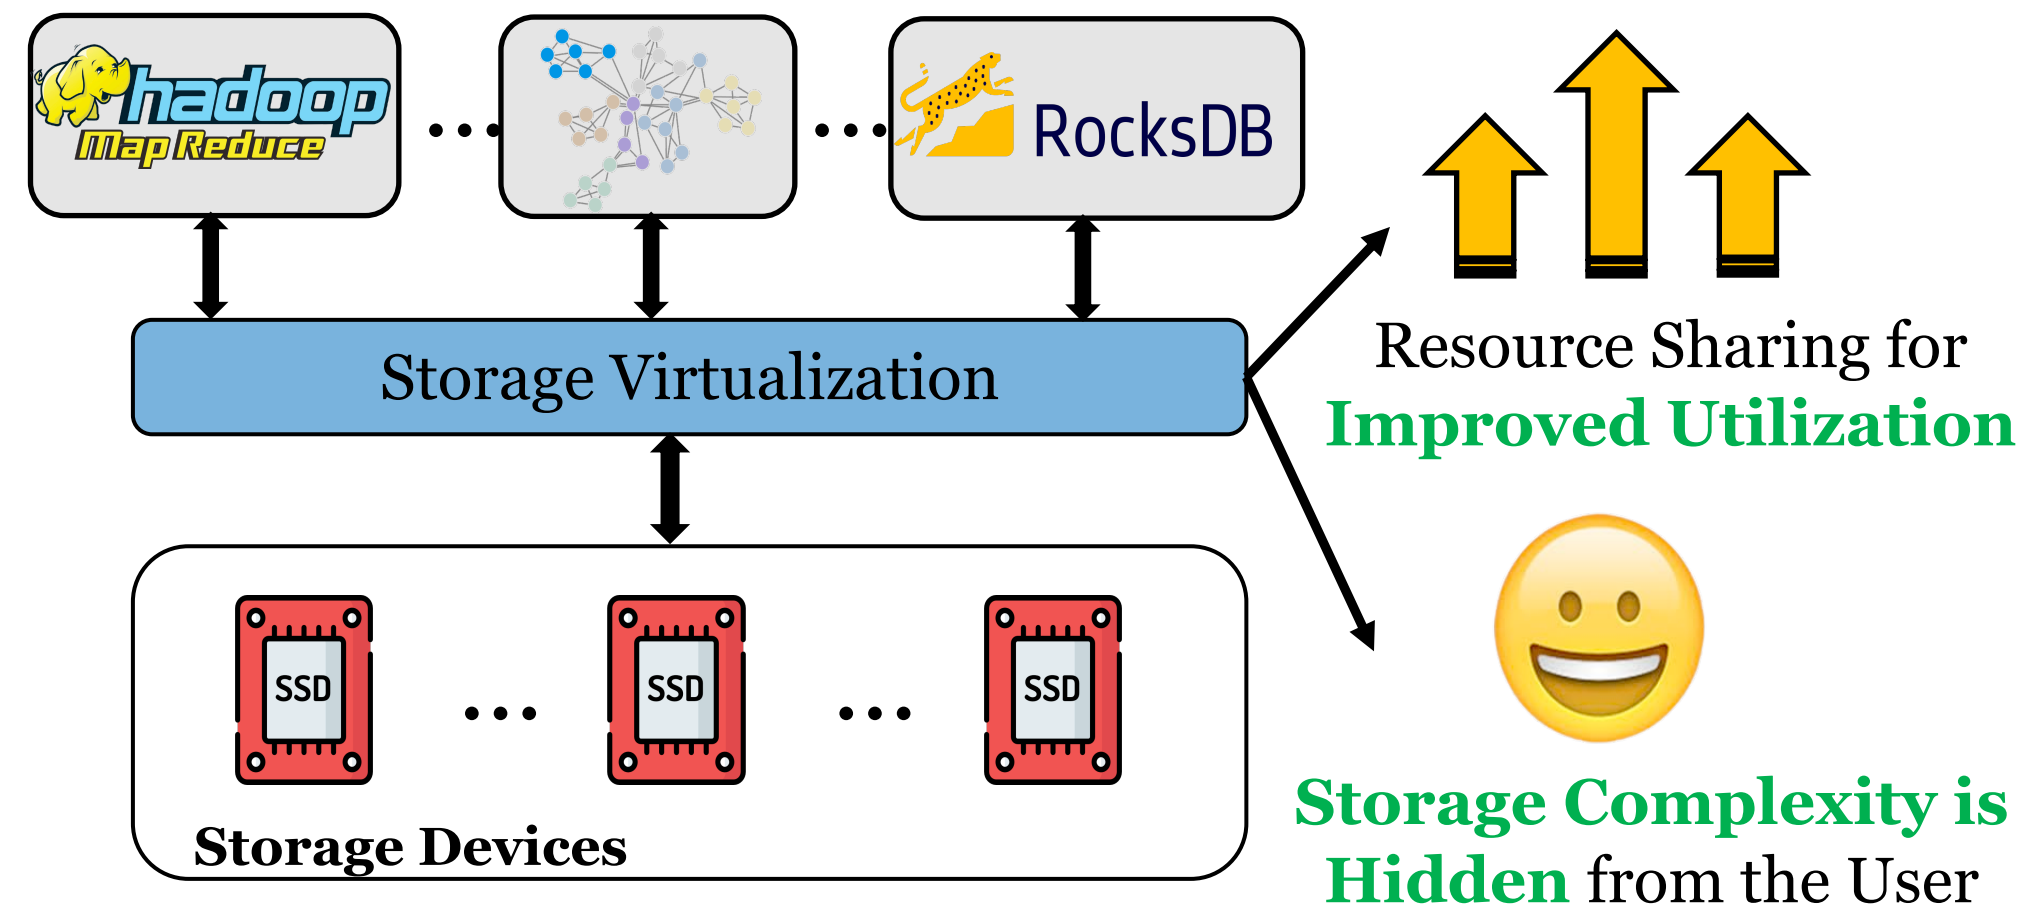
BlockFlex: Enabling Storage Harvesting with Software-Defined Flash in Modern Cloud Platforms

Benjamin Reidys* Jinghan Sun* Anirudh Badam† Shadi Noghabi† Jian Huang

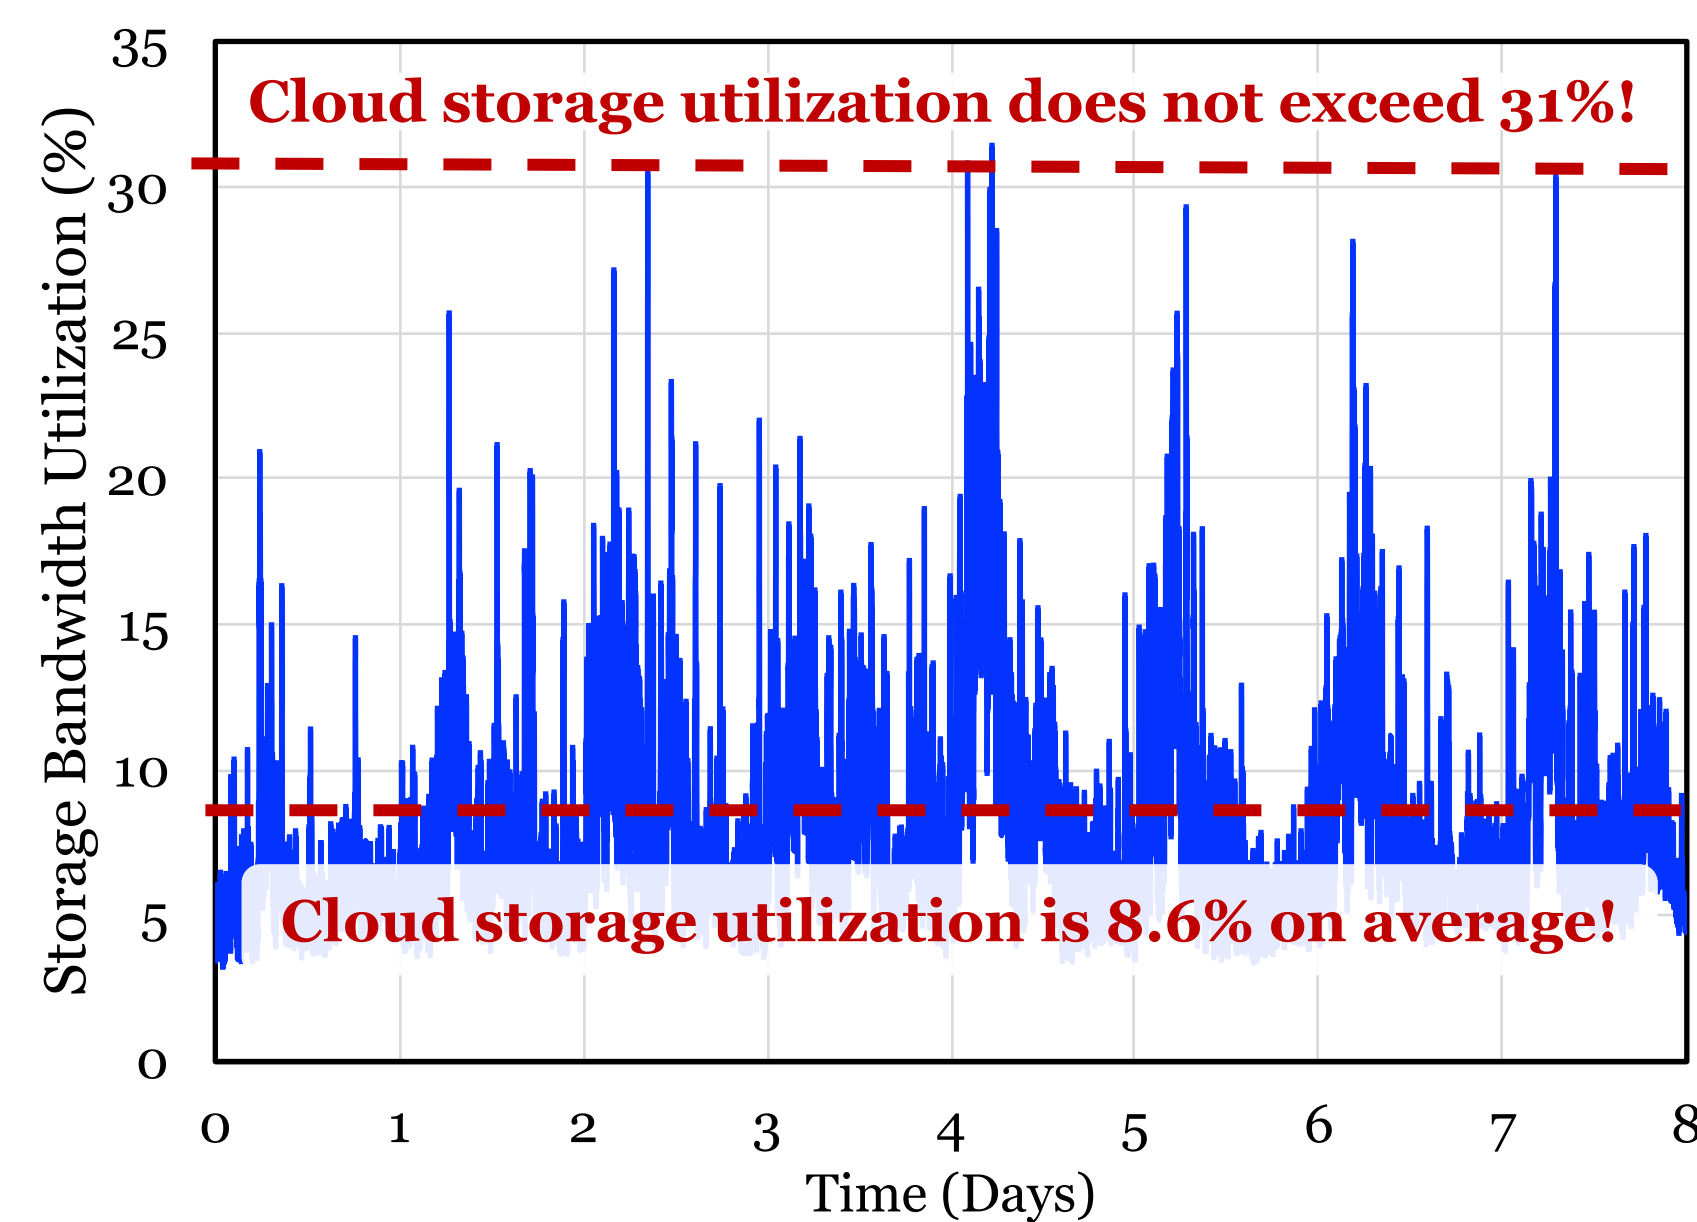
*Co-primary authors University of Illinois at Urbana-Champaign †Microsoft




BACKGROUND & MOTIVATION



Low Storage Utilization in the Cloud




Underutilized Allocated Storage



20% of VMs do not use any allocated storage

50% of VMs use only 26% of storage

Underutilized Unallocated Storage

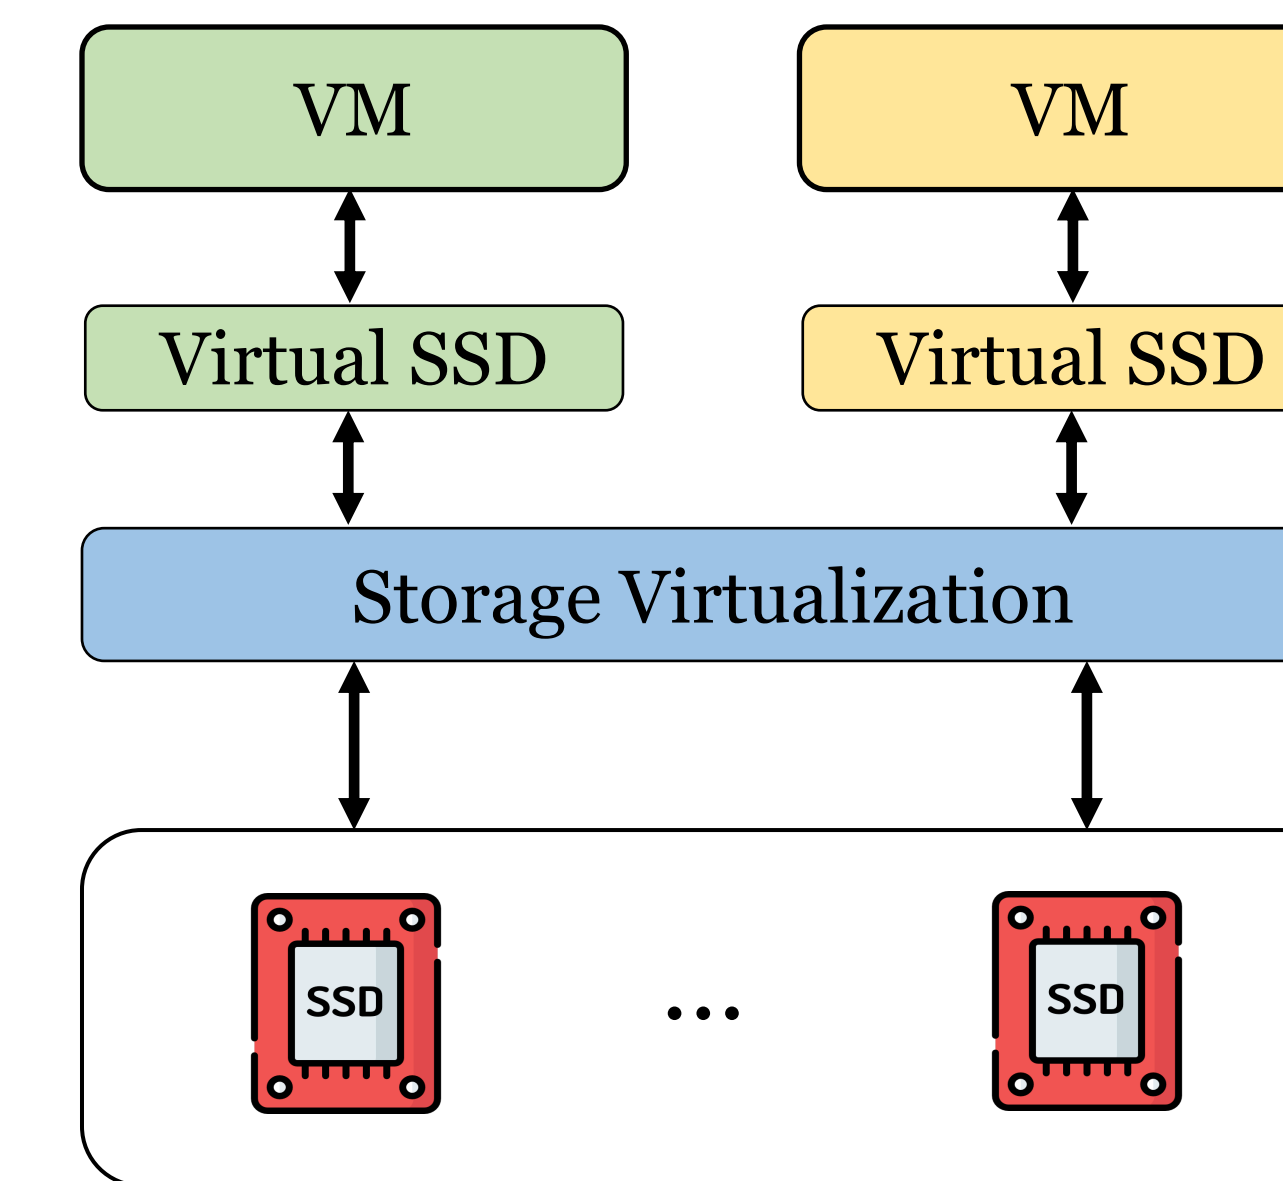
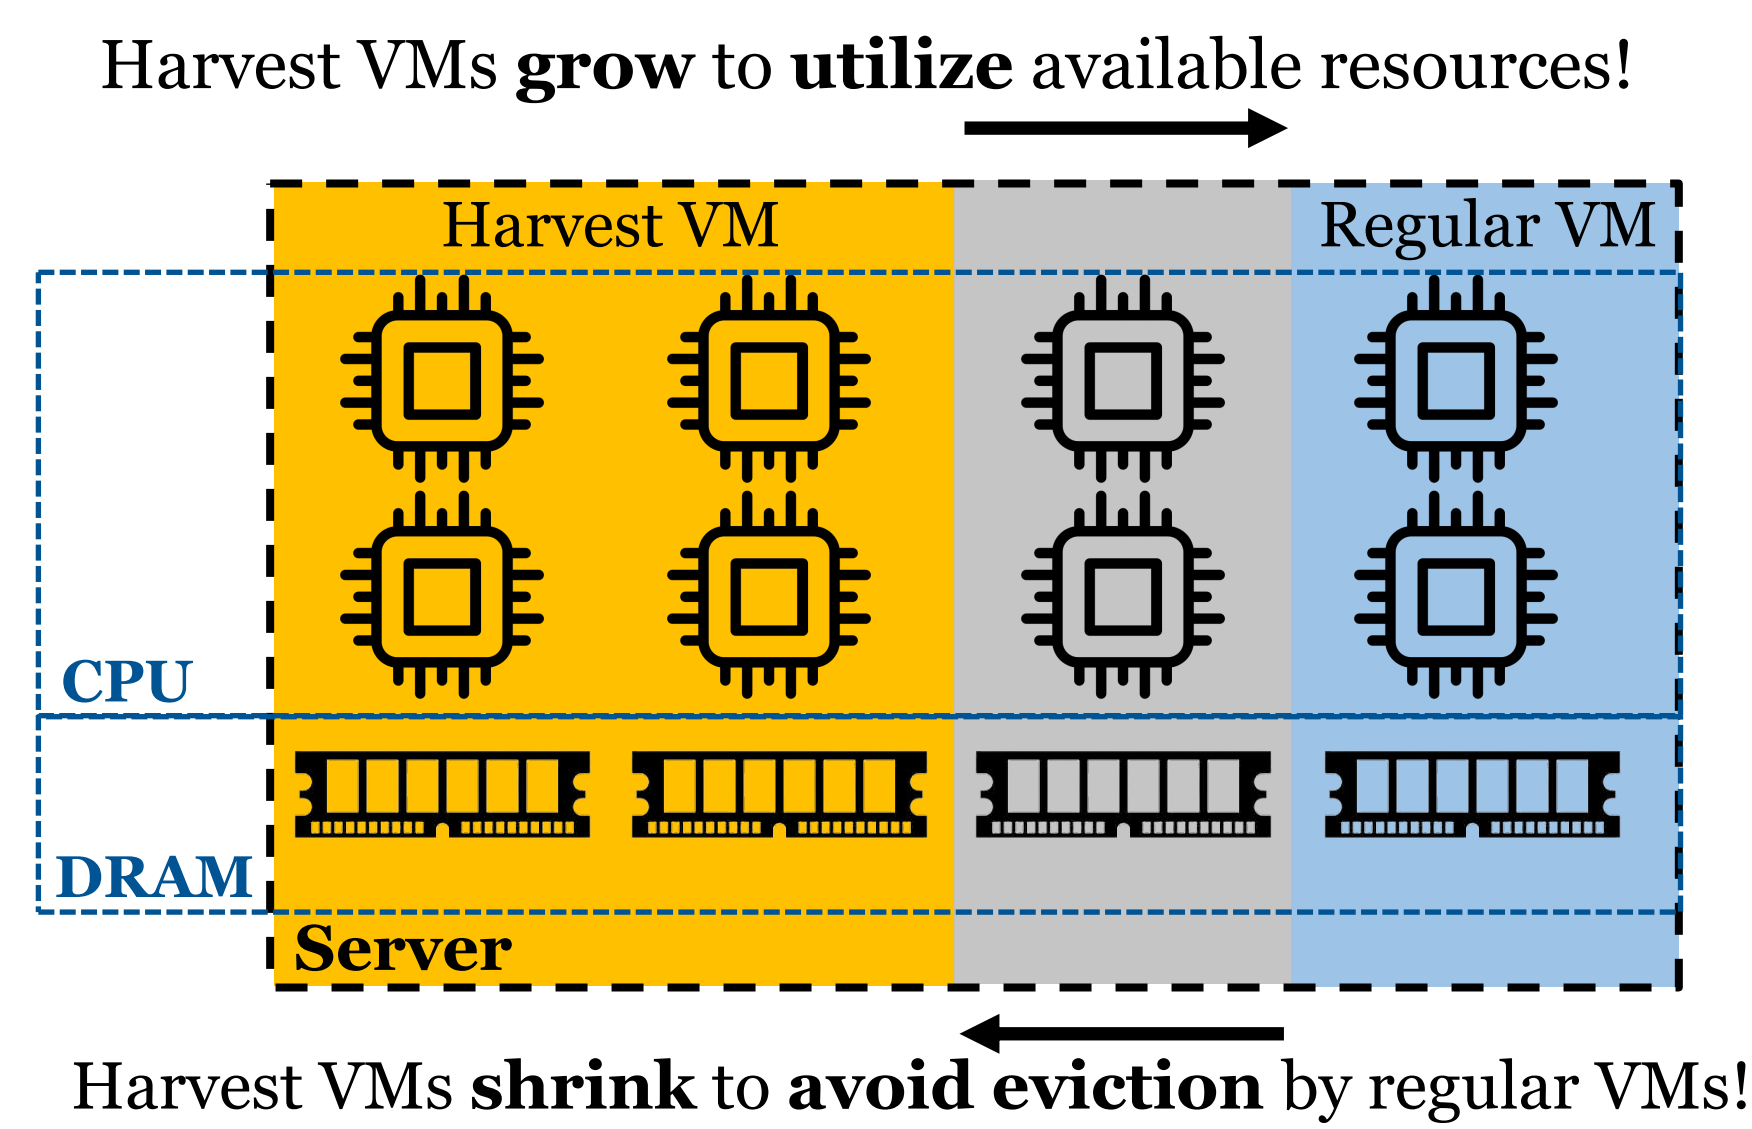


70% of servers have unsold space

50% of servers have 20% of space unsold

STORAGE HARVESTING IS CHALLENGING

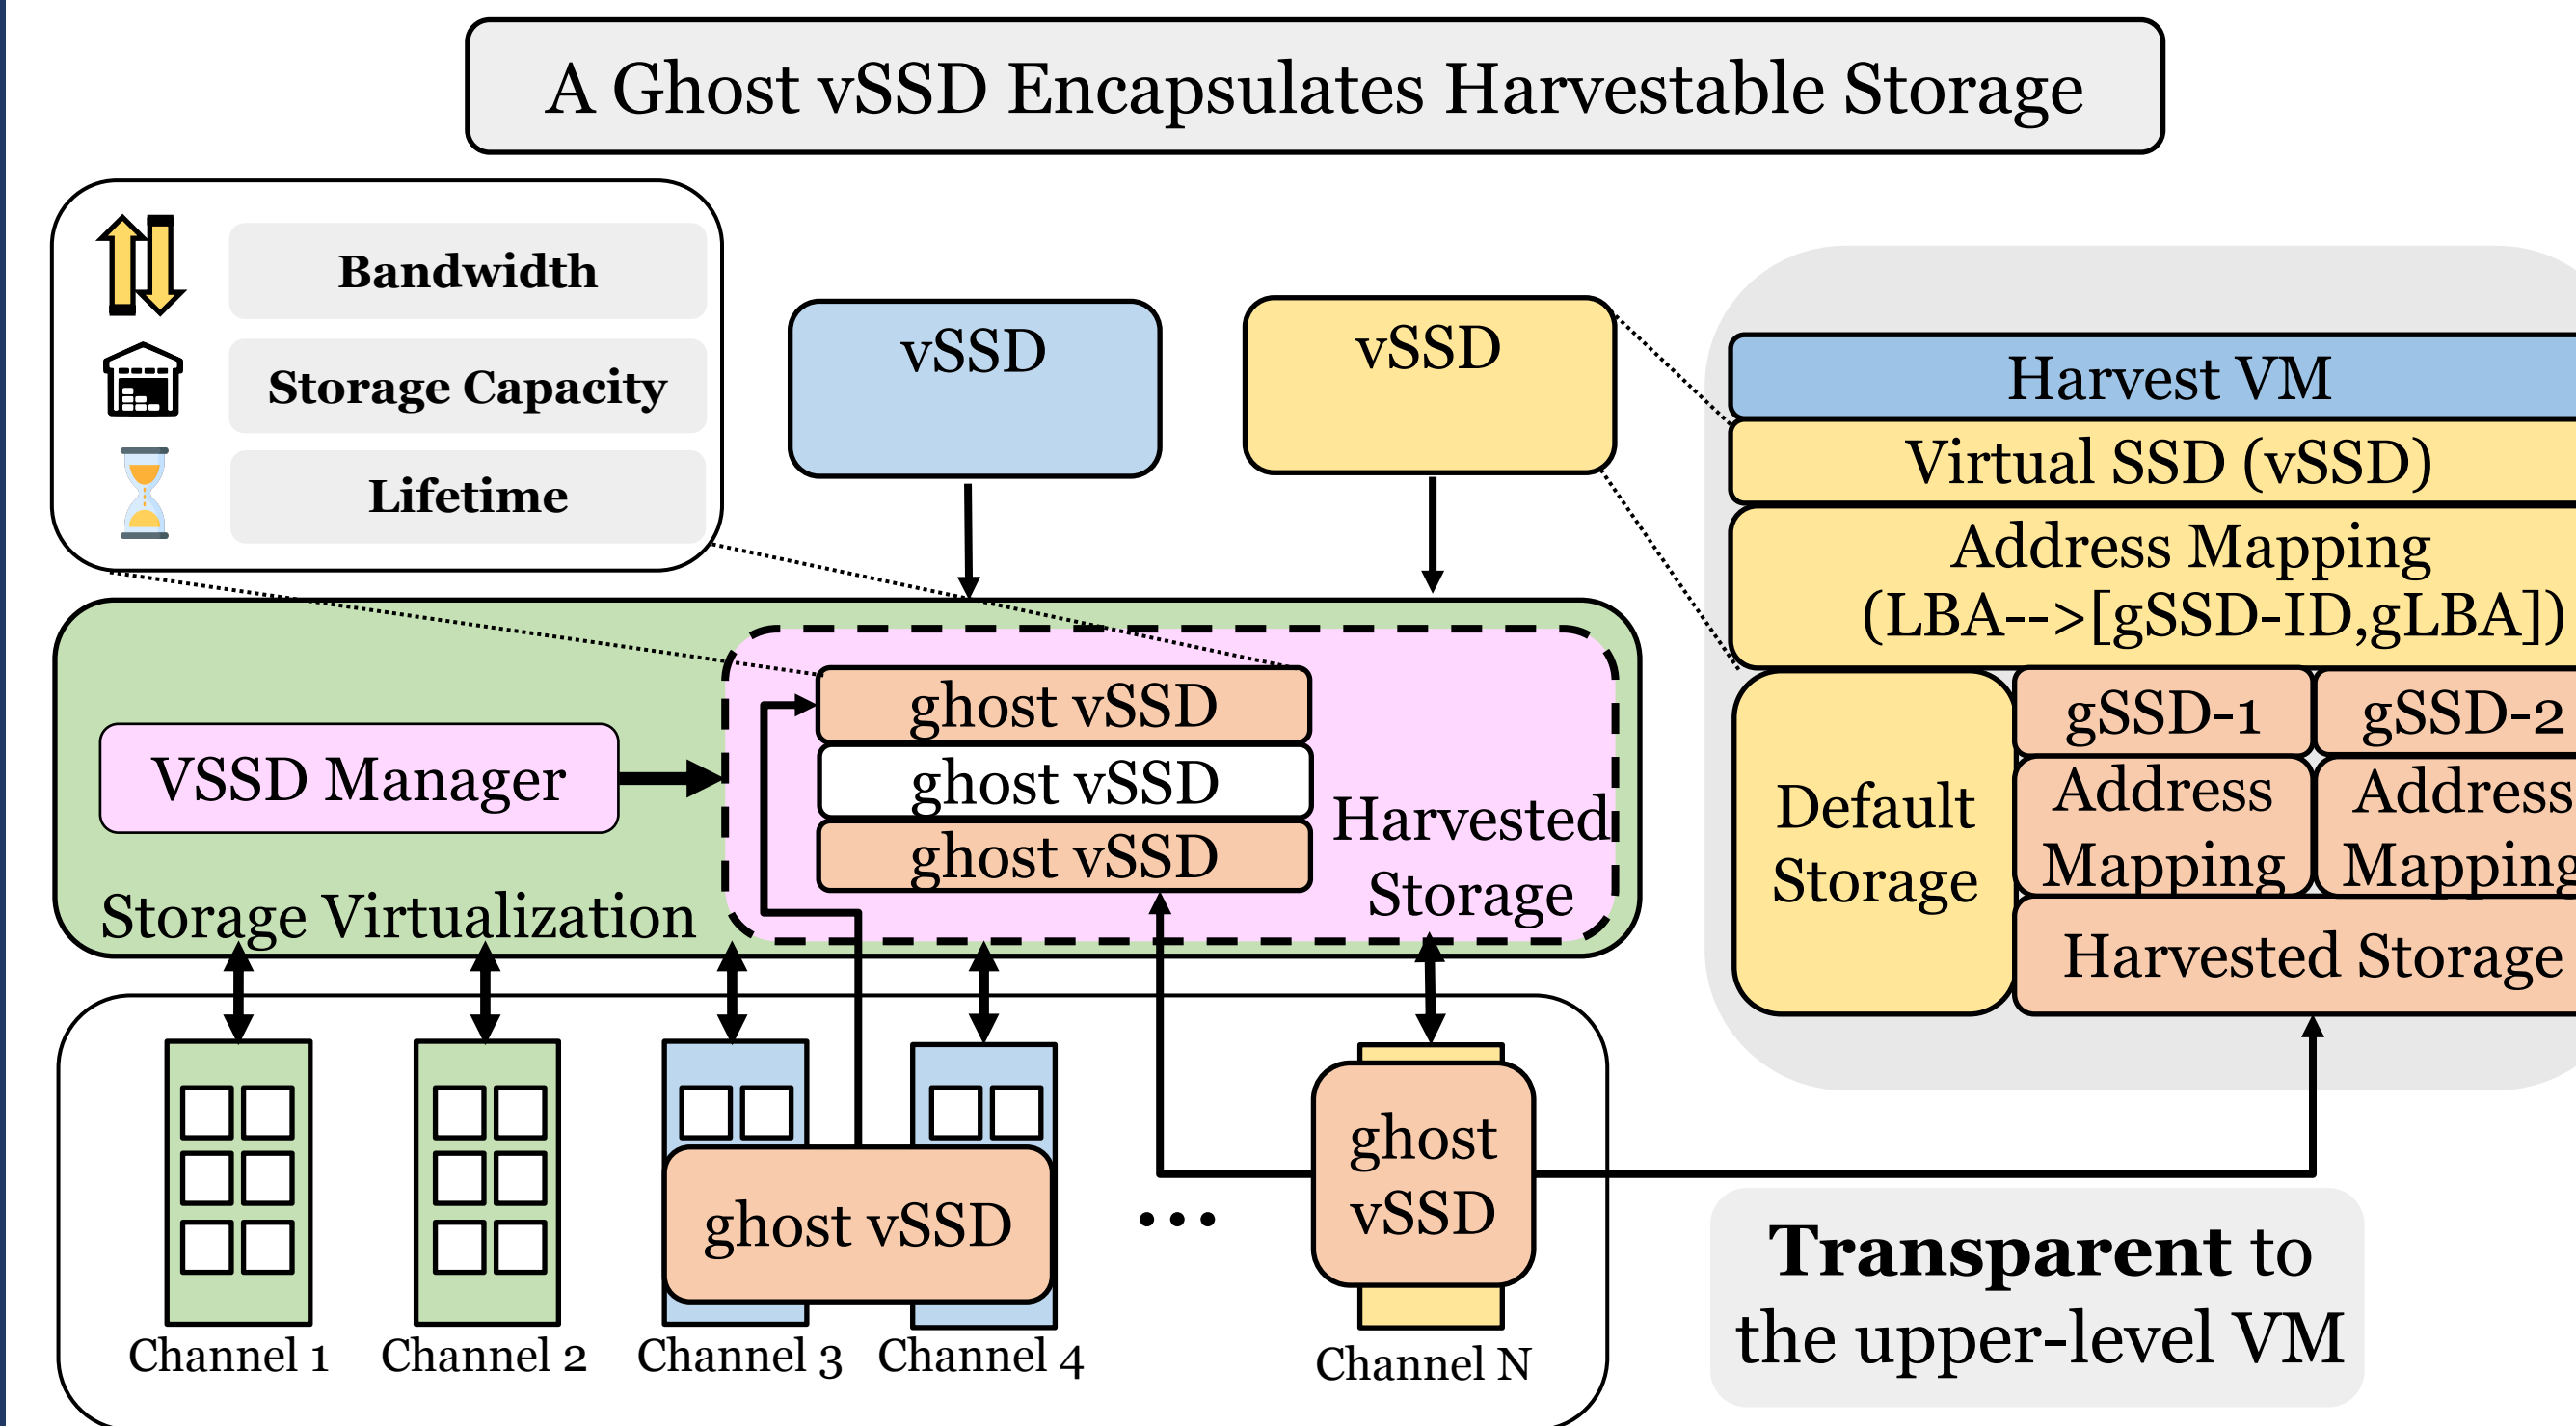
Harvest VM: Promising to Improve Resource Utilization



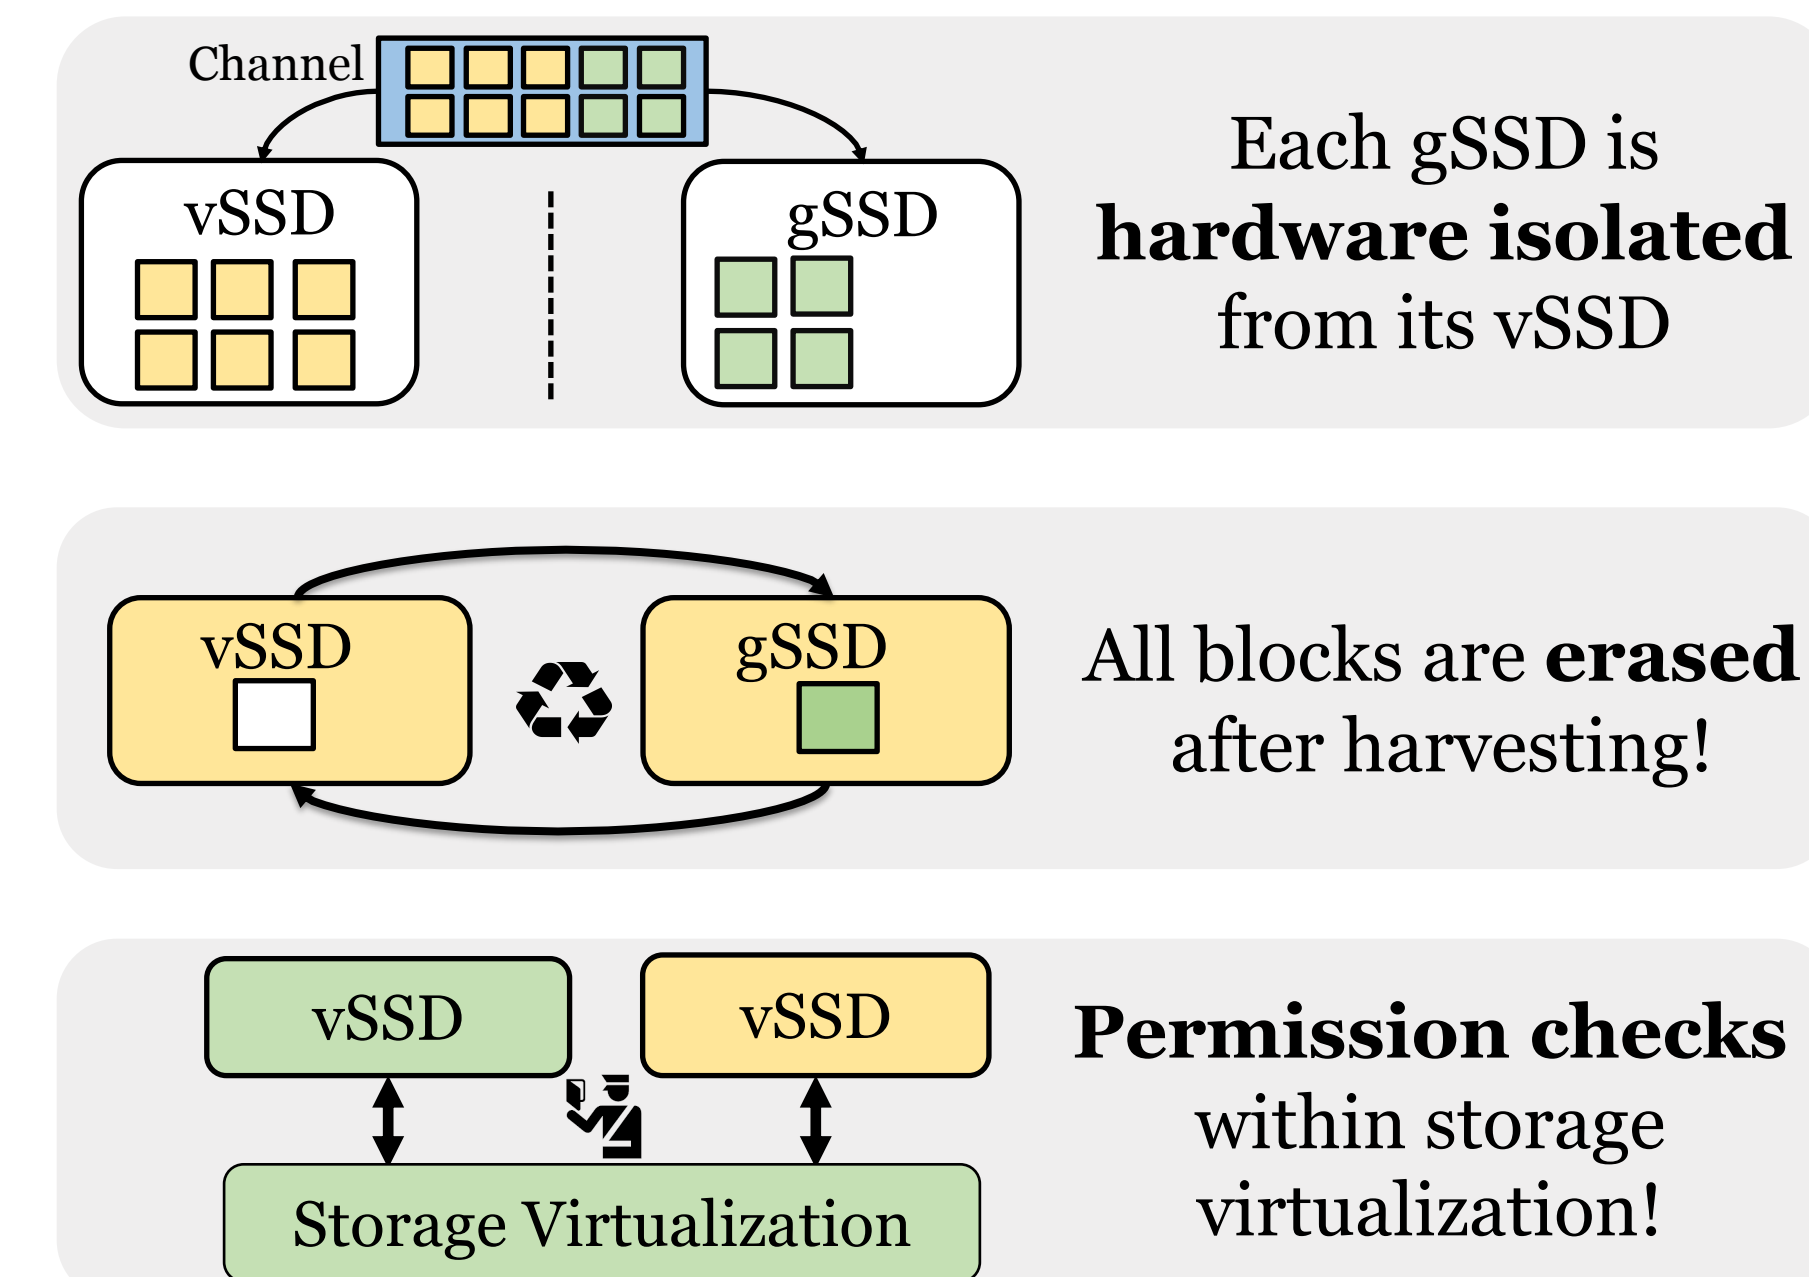
- 1 **Insufficient Virtualization Support**
- 2 **Data Security Concern**
- 3 **Storage Performance Interference**

DESIGN AND IMPLEMENTATION

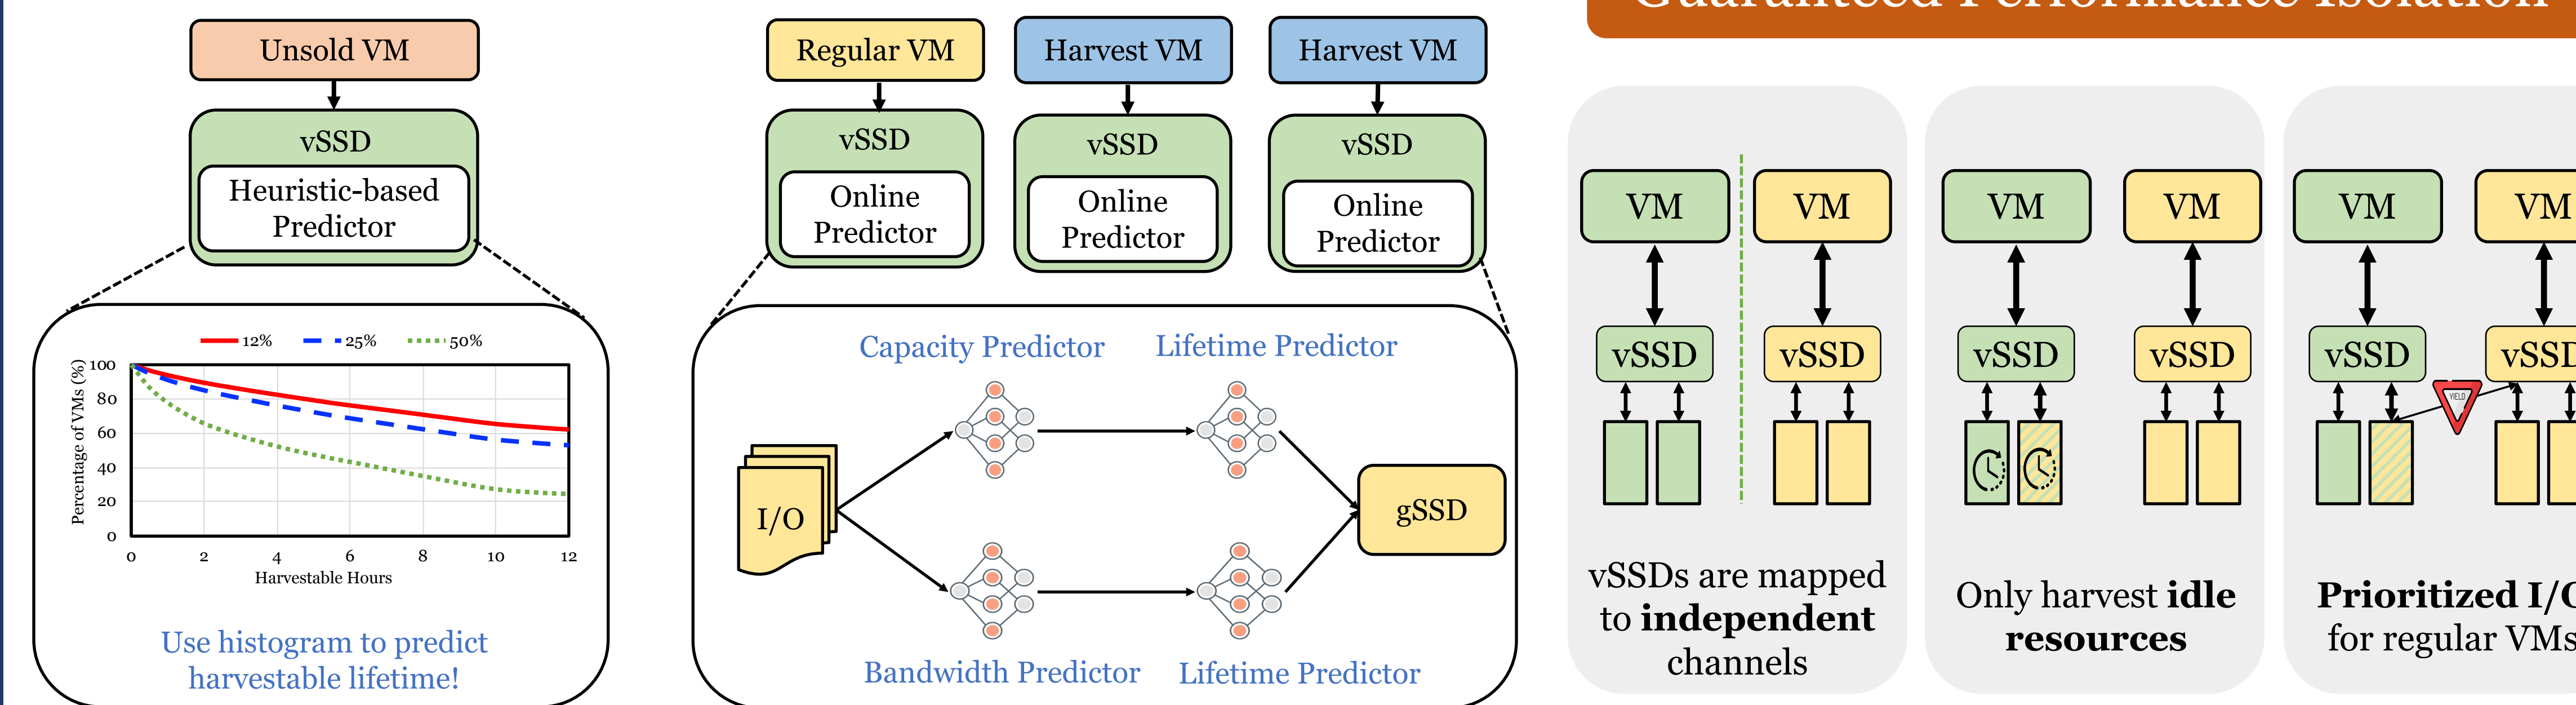
1 Rethink Storage Virtualization with Ghost vSSD



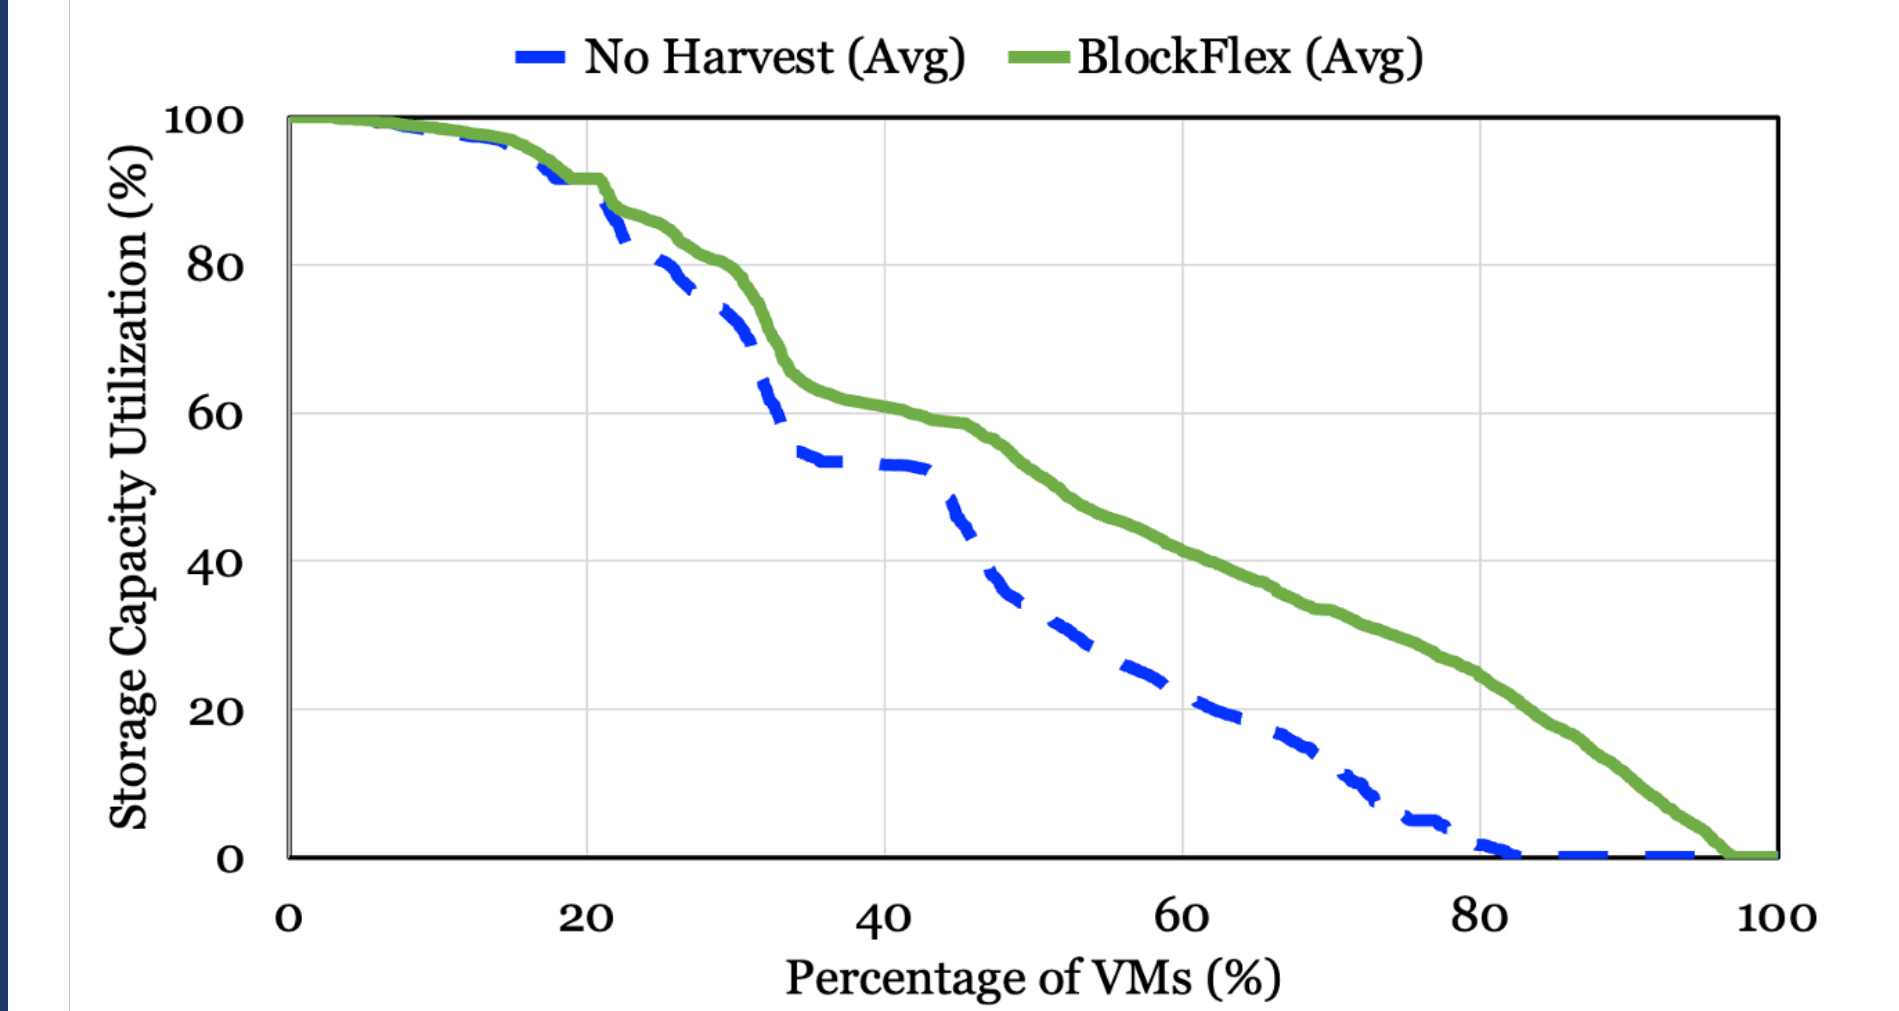
2 Enforcing Data Security



3 Guaranteed Performance Isolation

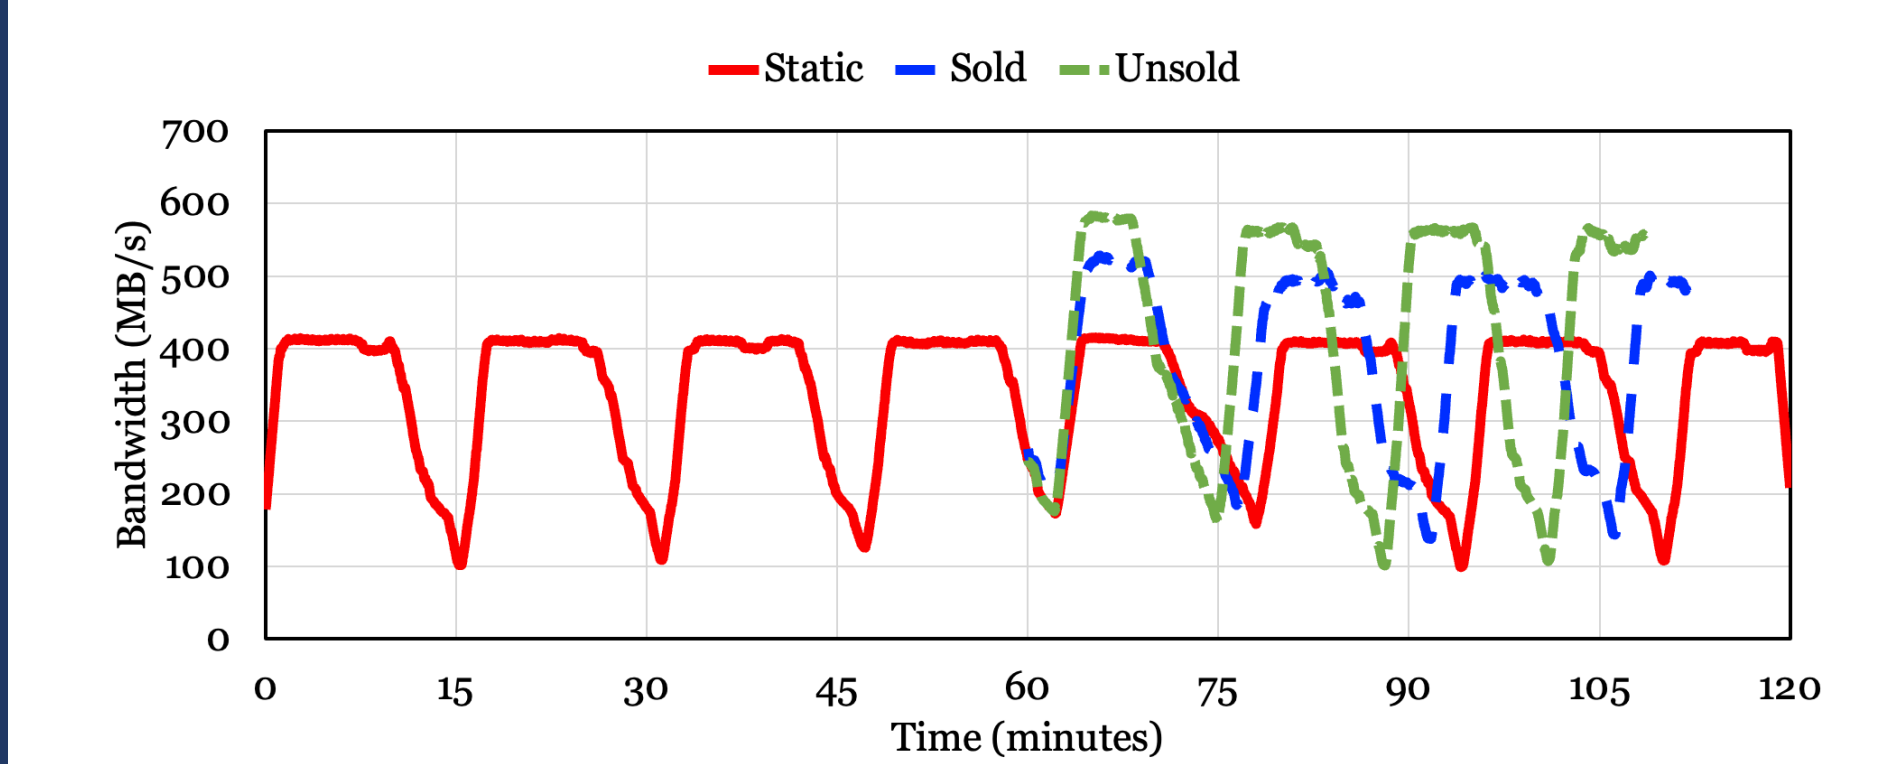


EVALUATION

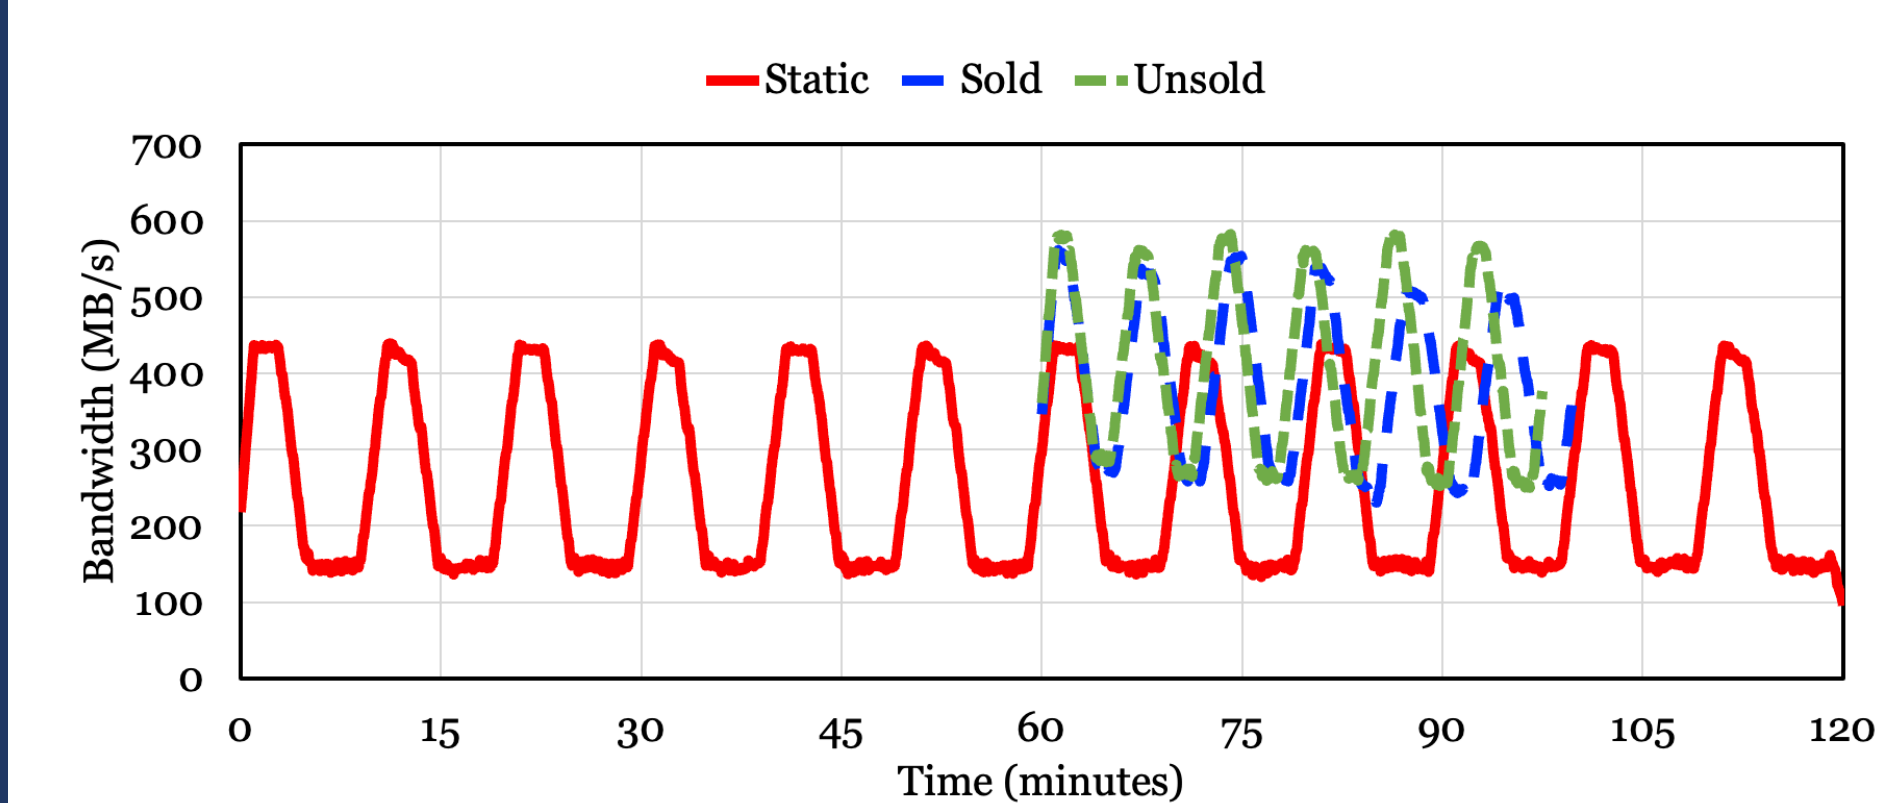


BlockFlex improves the storage utilization by 1.3x on average

ML Data Preprocessing



Graph Computing



BlockFlex improves harvest VM performance by up to 60%



BlockFlex codebase is available at: <https://github.com/platformlab/blockflex>