

Dynamic Block-level Cache Management for Cloud Computing Systems

Dulcardo Arteaga, Douglas Otstott, Ming Zhao

School of Computing and Information Sciences, Florida International University

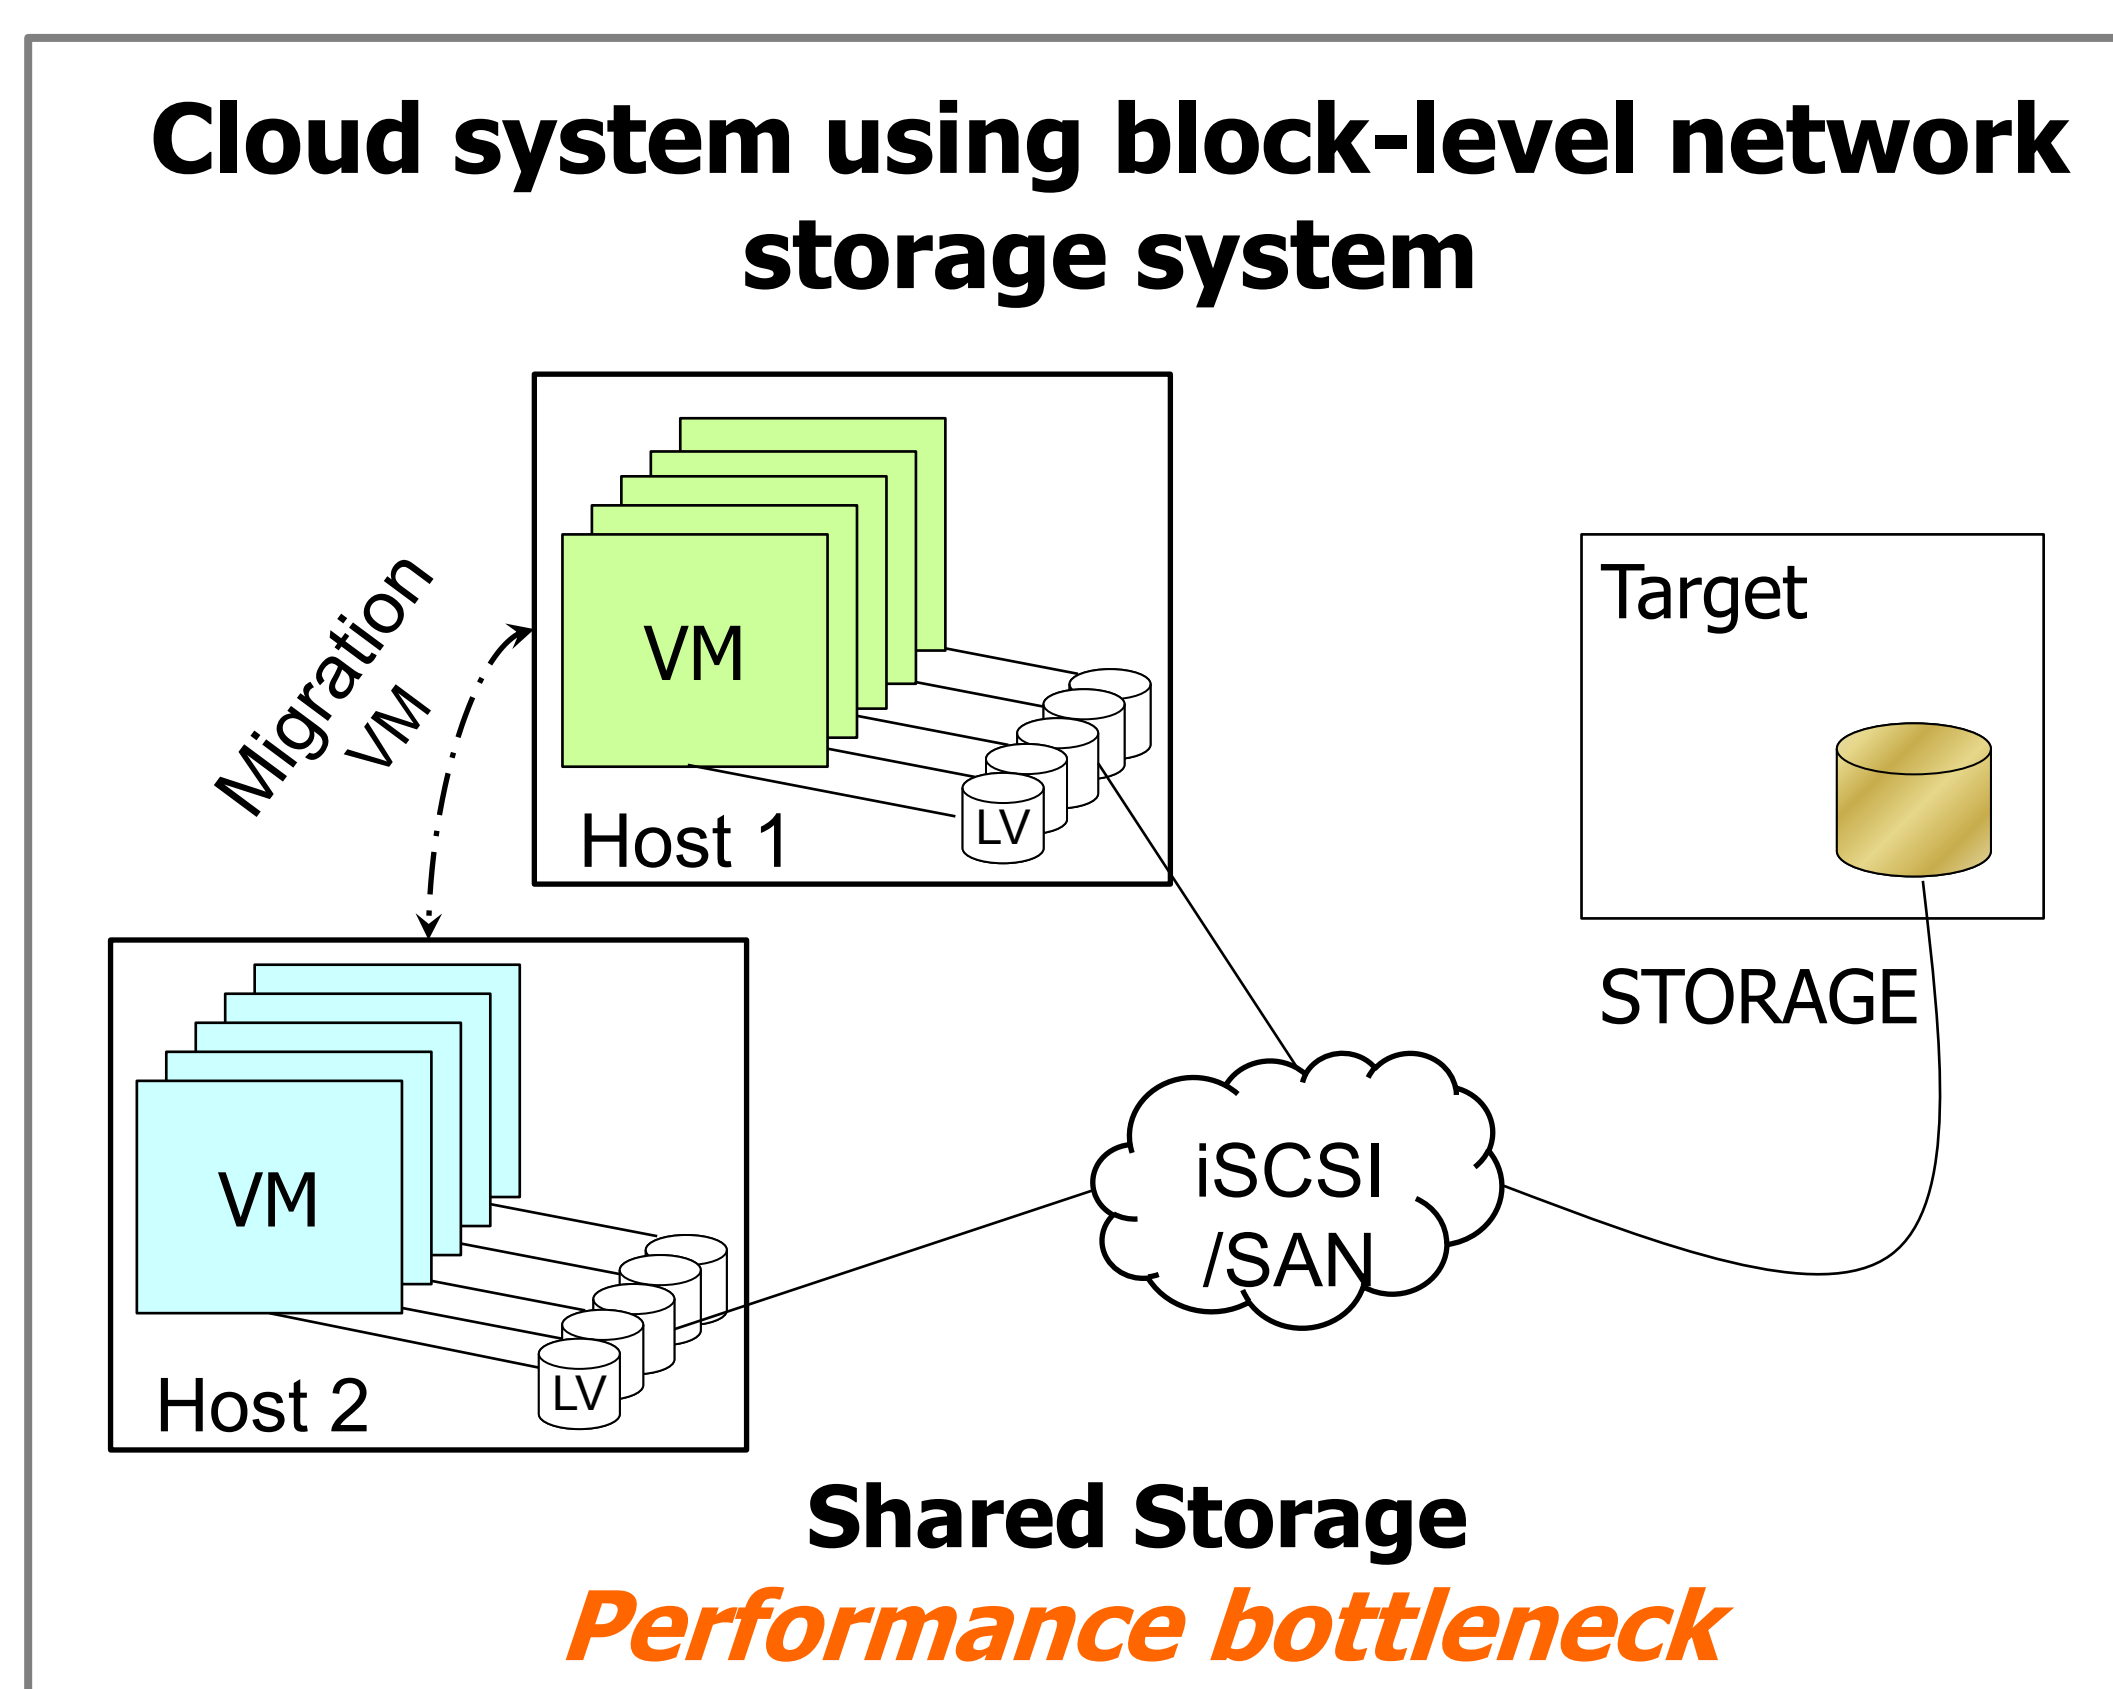
Background

Goal

- Improve I/O performance of virtual machines (VMs) in cloud computing systems using caching

Background

- Block-level network storage (iSCSI, NBD, SAN) commonly used in cloud systems
 - Fast VM migrations
 - Improved data availability
- Scalability becomes serious issue as the size of cloud systems continue to increase
 - Bottleneck in shared network storage
 - Performance interference across VMs



Proposed Solution

Dynamic block-level client-side caching for cloud computing systems

- Unified cache shared by co-hosted VMs to achieve full resource utilization
- Ability to differentiate and isolate block I/Os from different VMs
- Cache policies for:
 - cache replacement
 - write-through vs. write-back
- Maintain data consistency by periodically flushing write-back-cached data

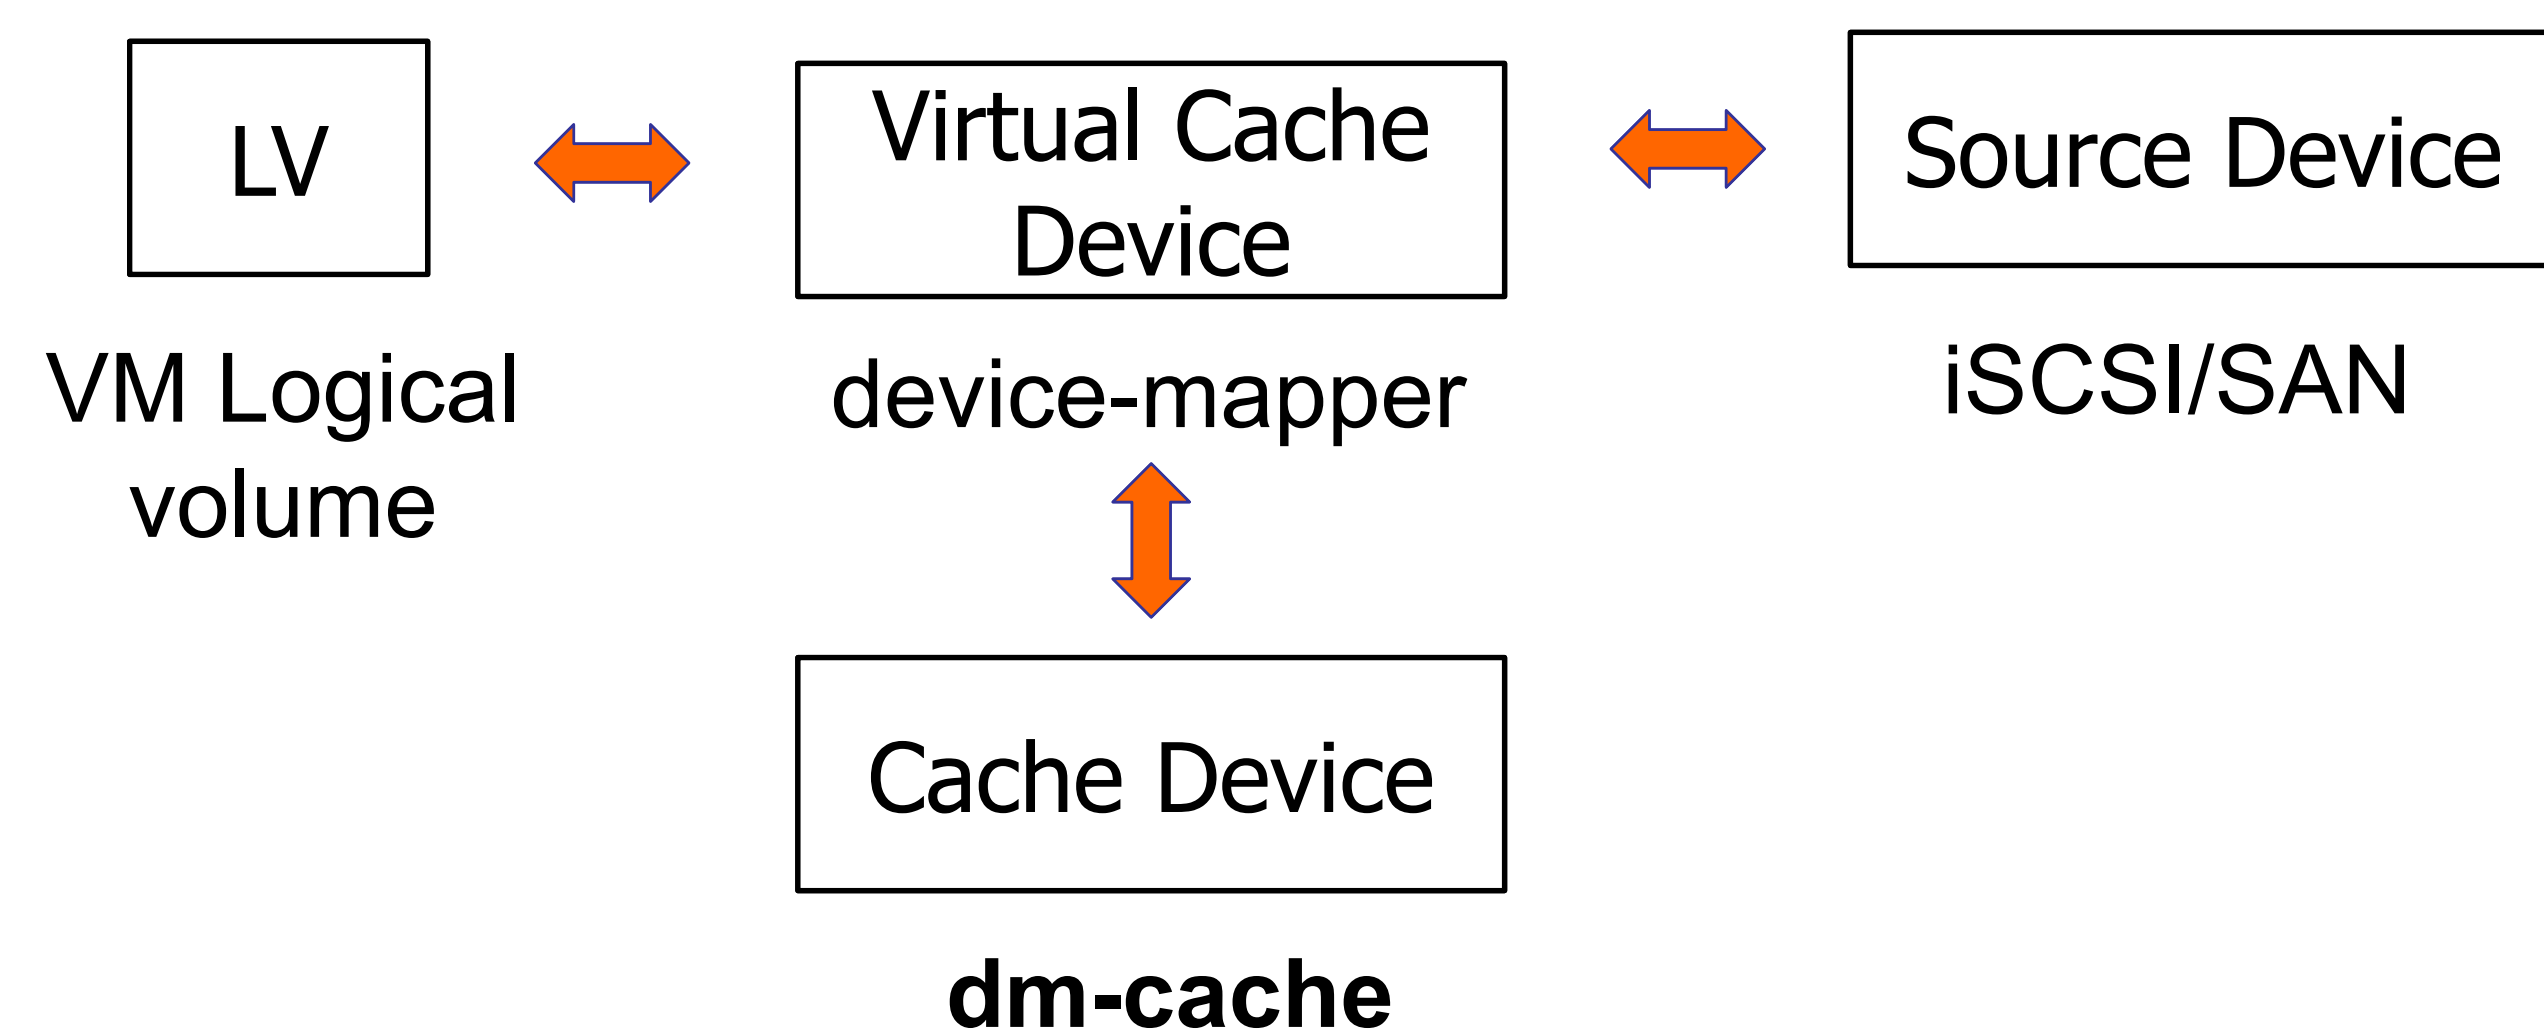
System Design

Block device virtualization based caching

- Device-mapper is a generic framework for creating virtual block devices
- DM-cache** provide a generic block-level disk cache for network storage systems

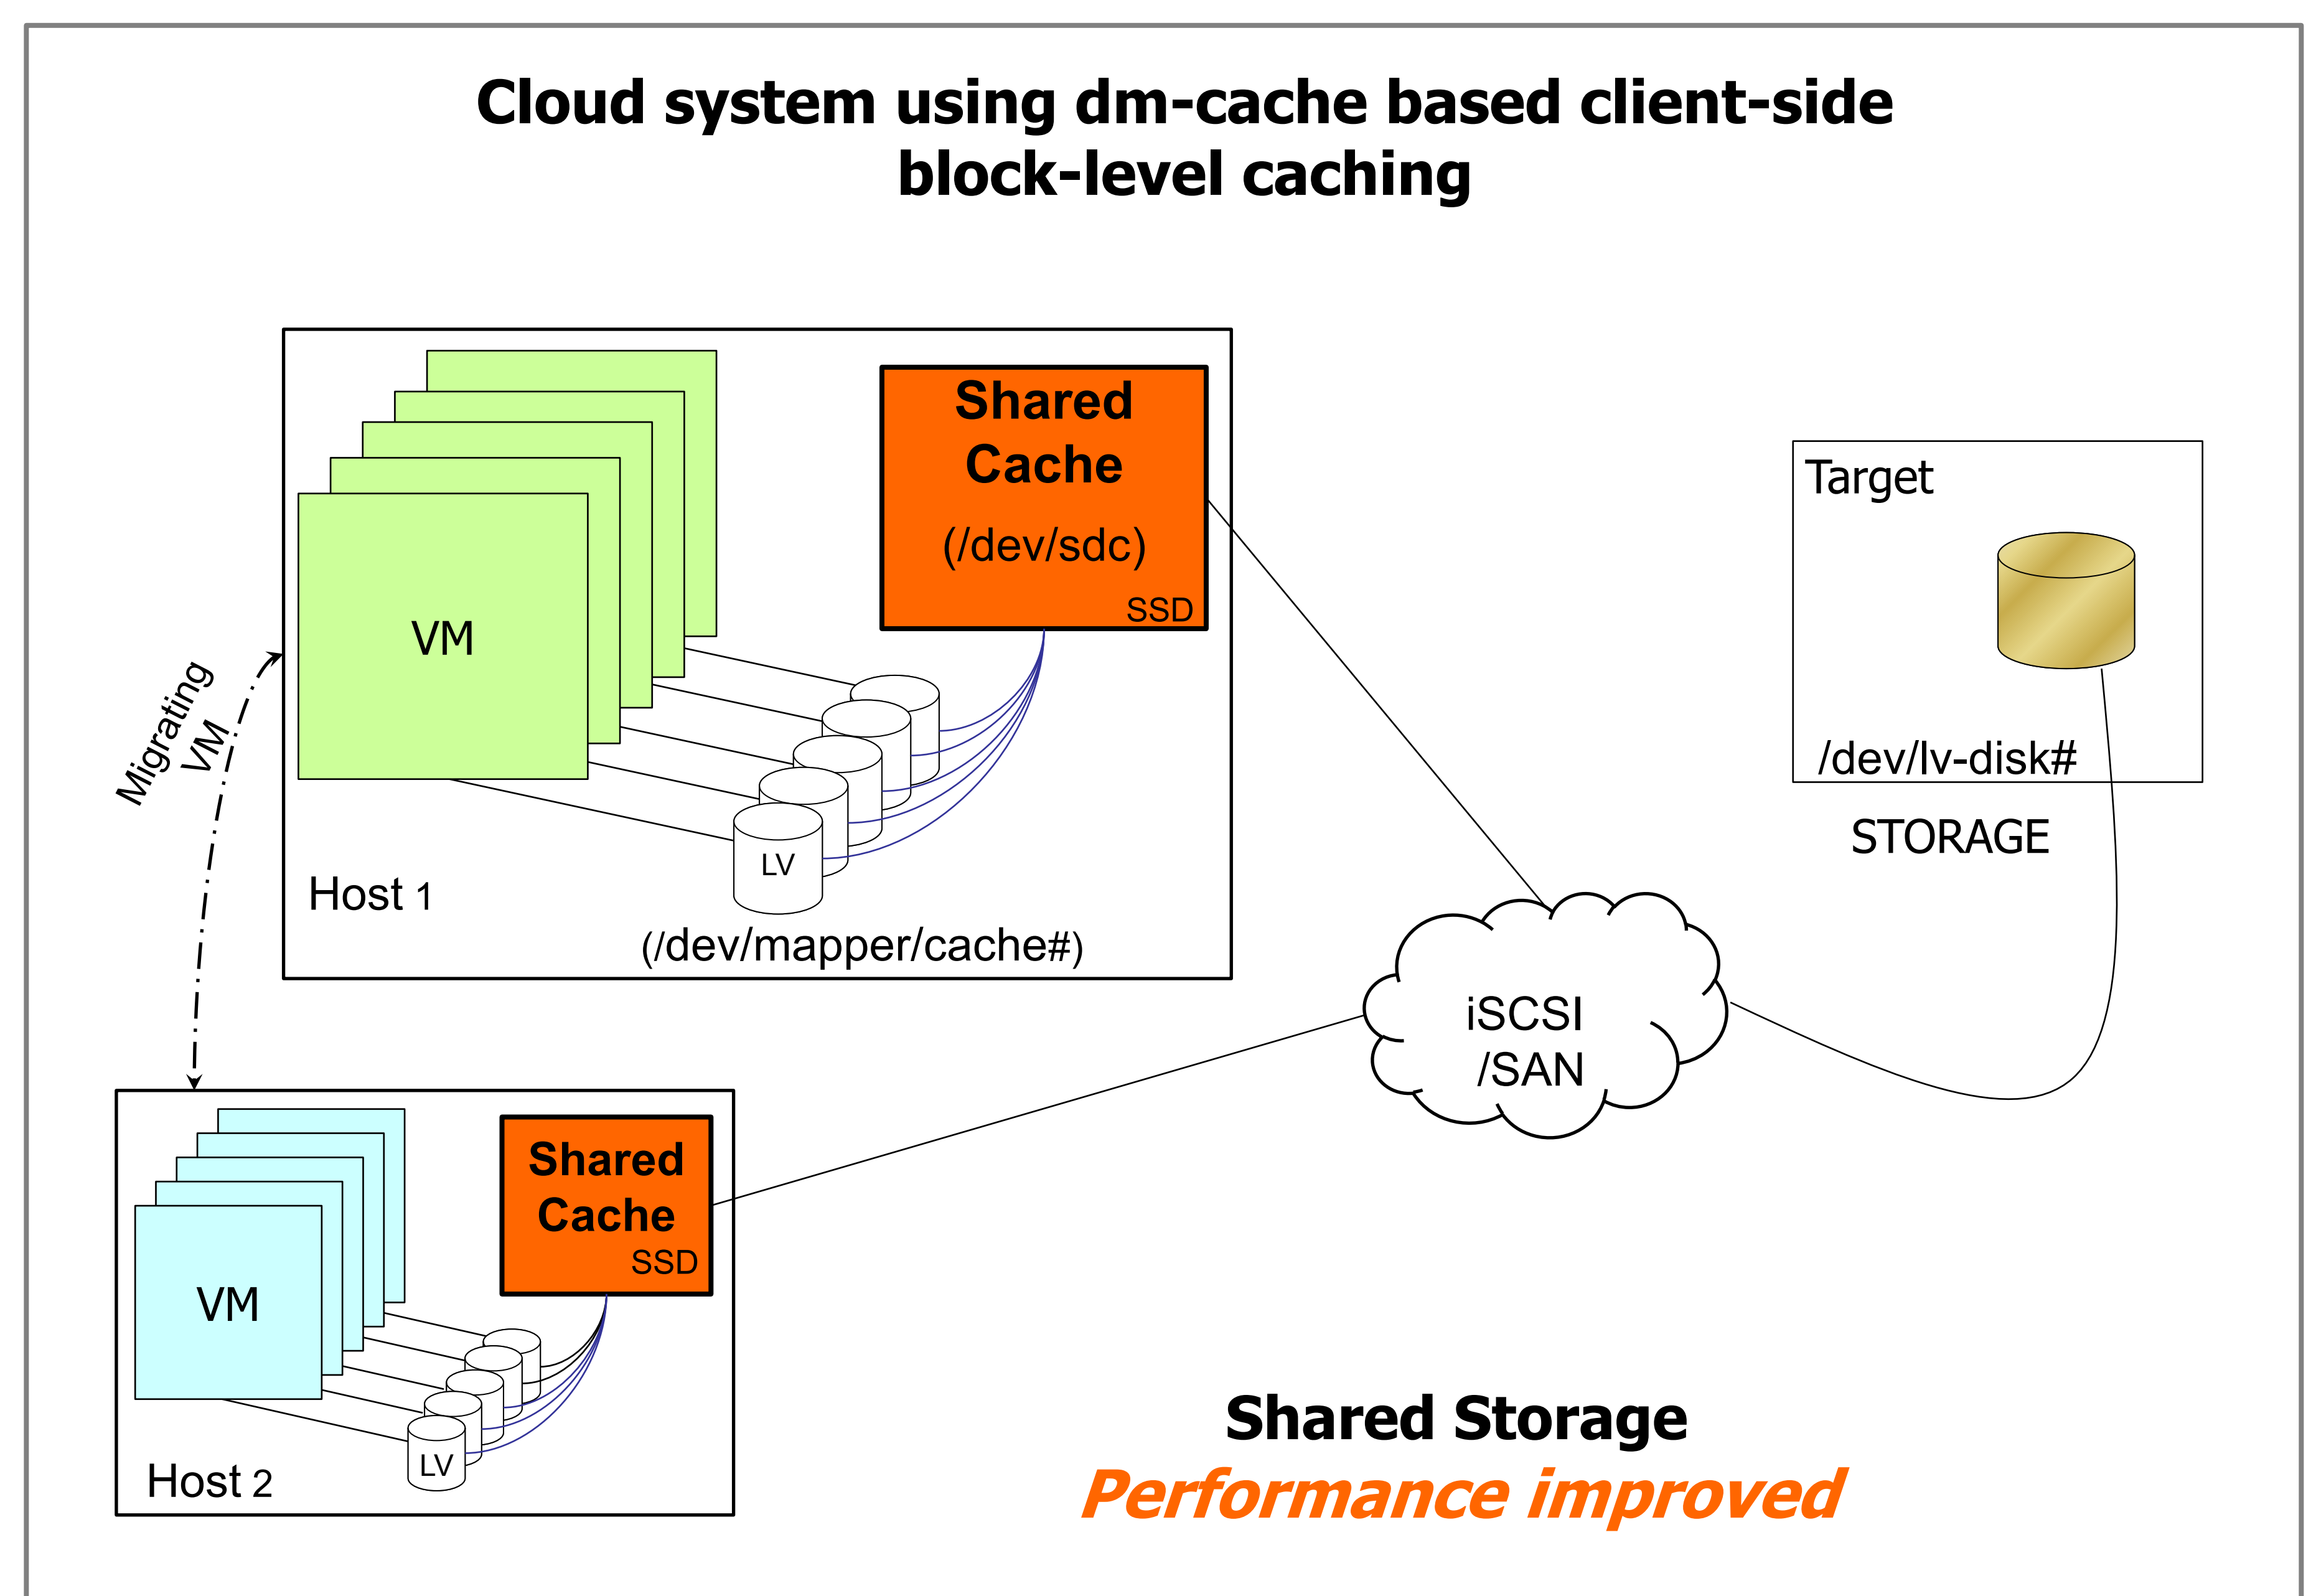
Unified cache for co-hosted VMs

- Create per-VM virtual block device through dm-cache in order to differentiate block-level I/Os from different VMs
- Map the different virtual block devices to the same cache device in order to maximize the cache utilization



Proposed Approach

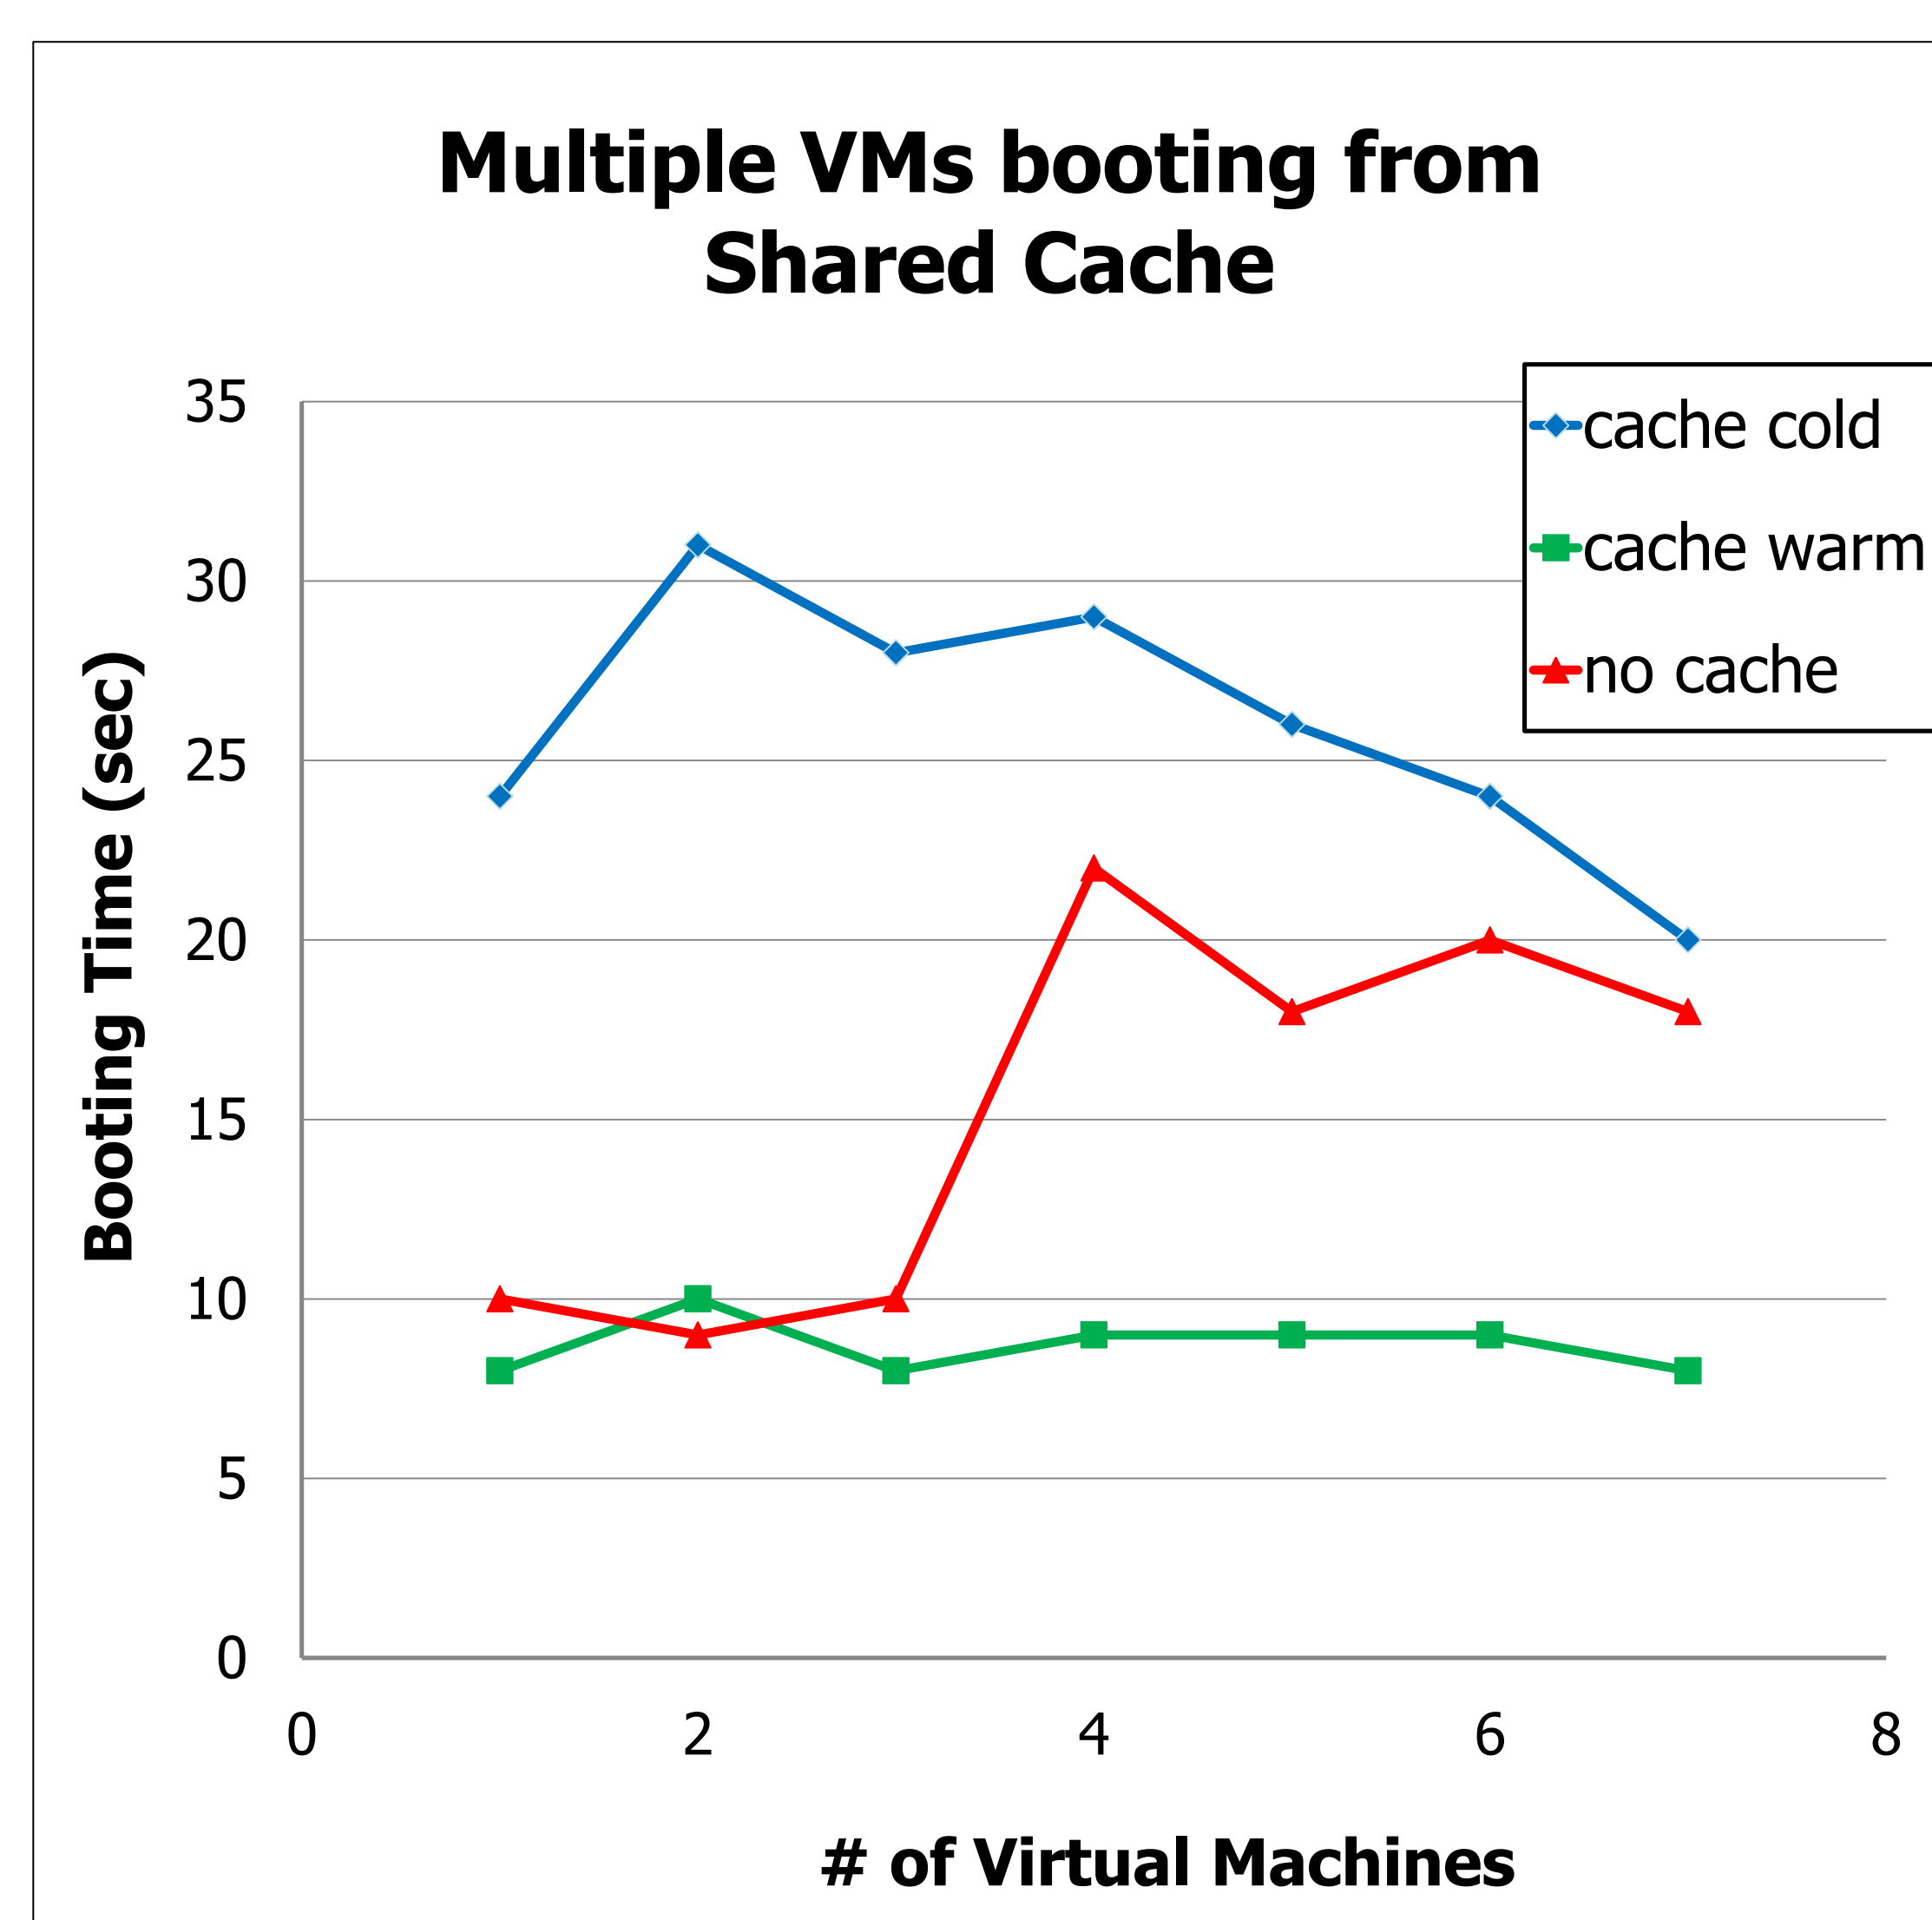

Cloud system using dm-cache based client-side block-level caching



Experimental Evaluation

Experiment setup:

- SSD devices for local cache
- iSCSI for network storage
- Benchmarks:
 - IOzone read/re-read
 - Boot multiple VMs using share cache



Conclusion and Future Work

Conclusions

- Dm-cache effectively uses client-side storage to exploit locality for multiple VMs running on the same physical host
- SSD-based results show a performance improvement of 43% when booting VMs concurrently

Future Work

- Implement more intelligent algorithm for shared cache partitioning while guaranteeing fairness across all VMs
- Consider the unique characteristics of SSDs devices and design optimized caching policies
- Consider cross-client cooperative caching to further improve caching efficiency



Virtualized Infrastructure, Systems & Applications (VISA) Research Laboratory (<http://visa.cis.fiu.edu>)

This work is partly sponsored by a gift donation of XLS Hosting, a cloud provider company in Netherlands.

