Flexible scheme of space management in ZNS SSD / SMR HDD storage pool for multiple virtual machines

Viacheslav Dubeyko

Content

1. Storage pool concept
2. Architecture vision
3. Storage pool space management problem
4. Page table like architecture
5. Open questions

Storage pool concept

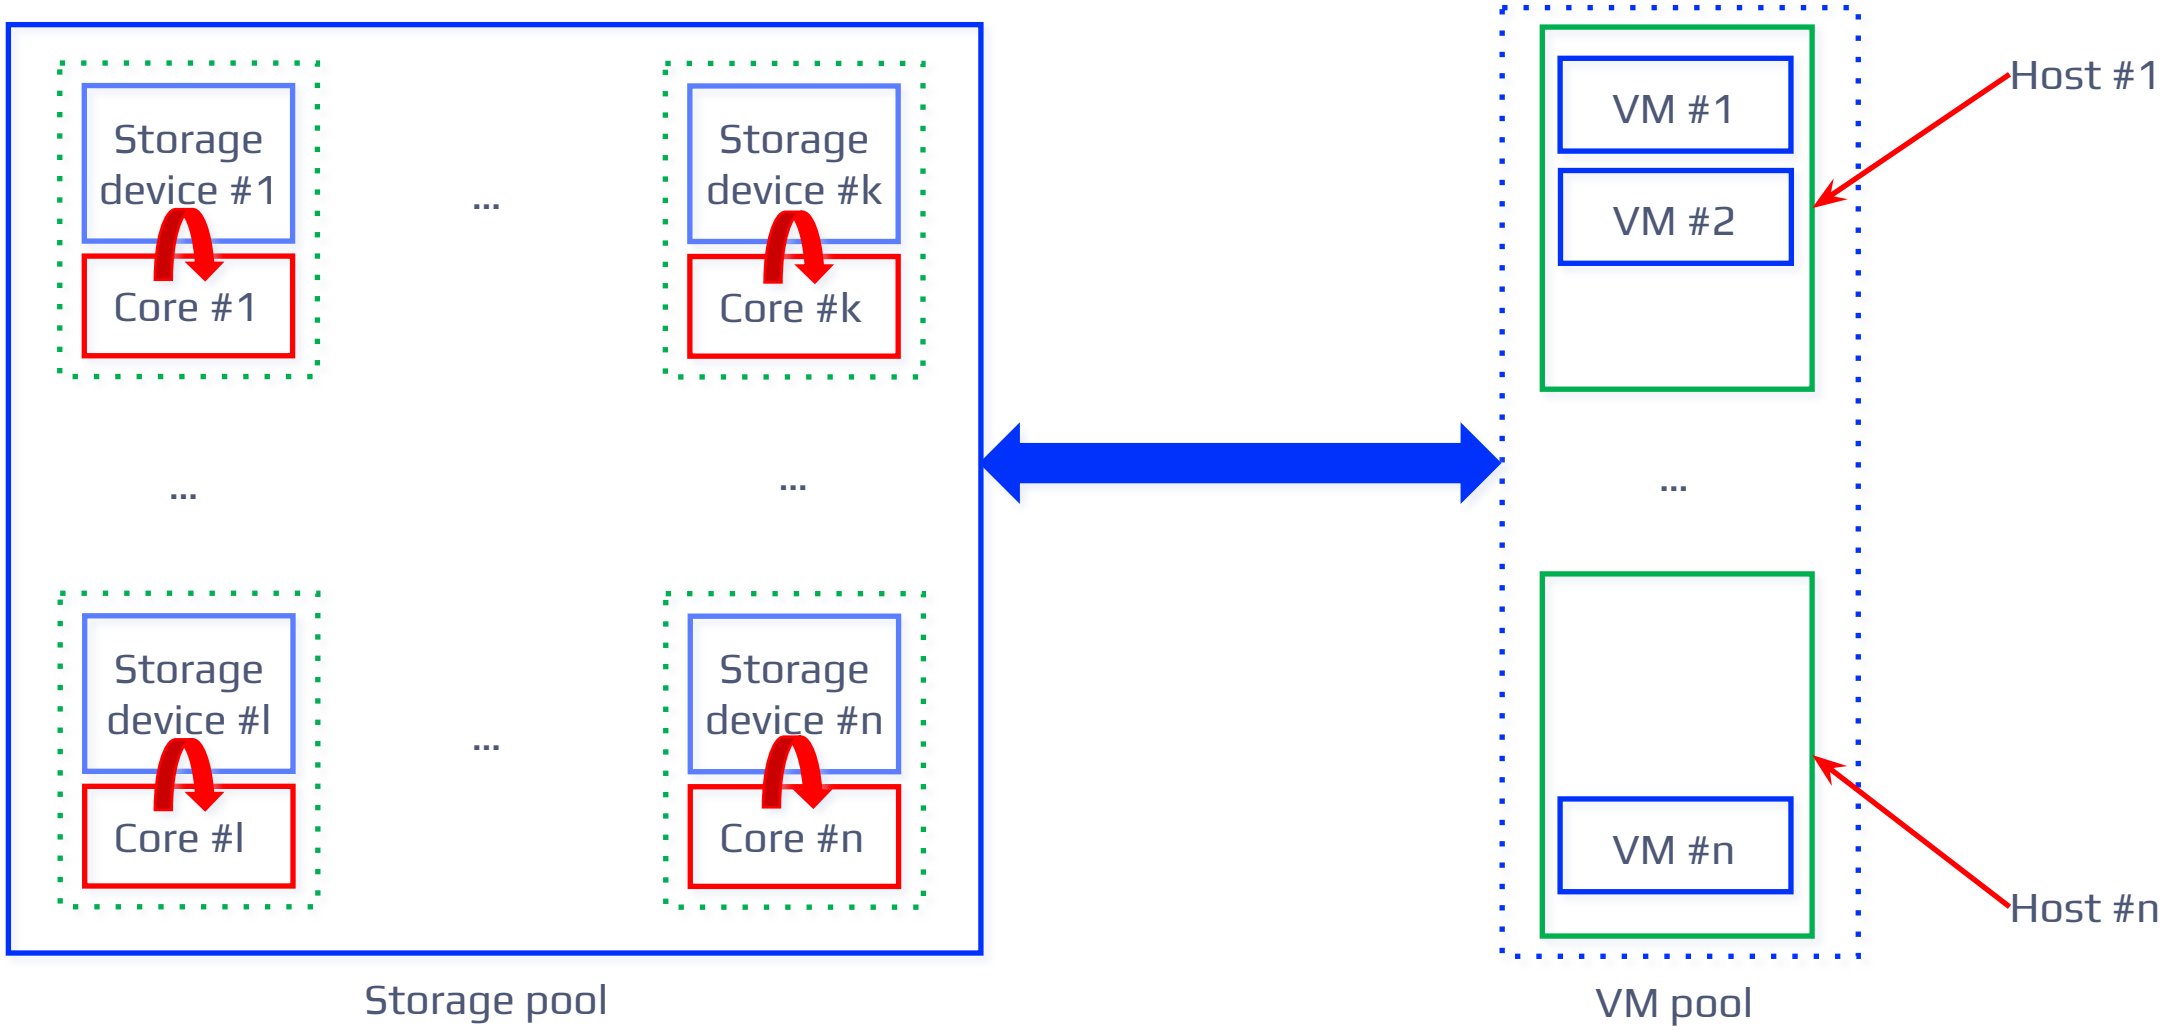

Architecture vision

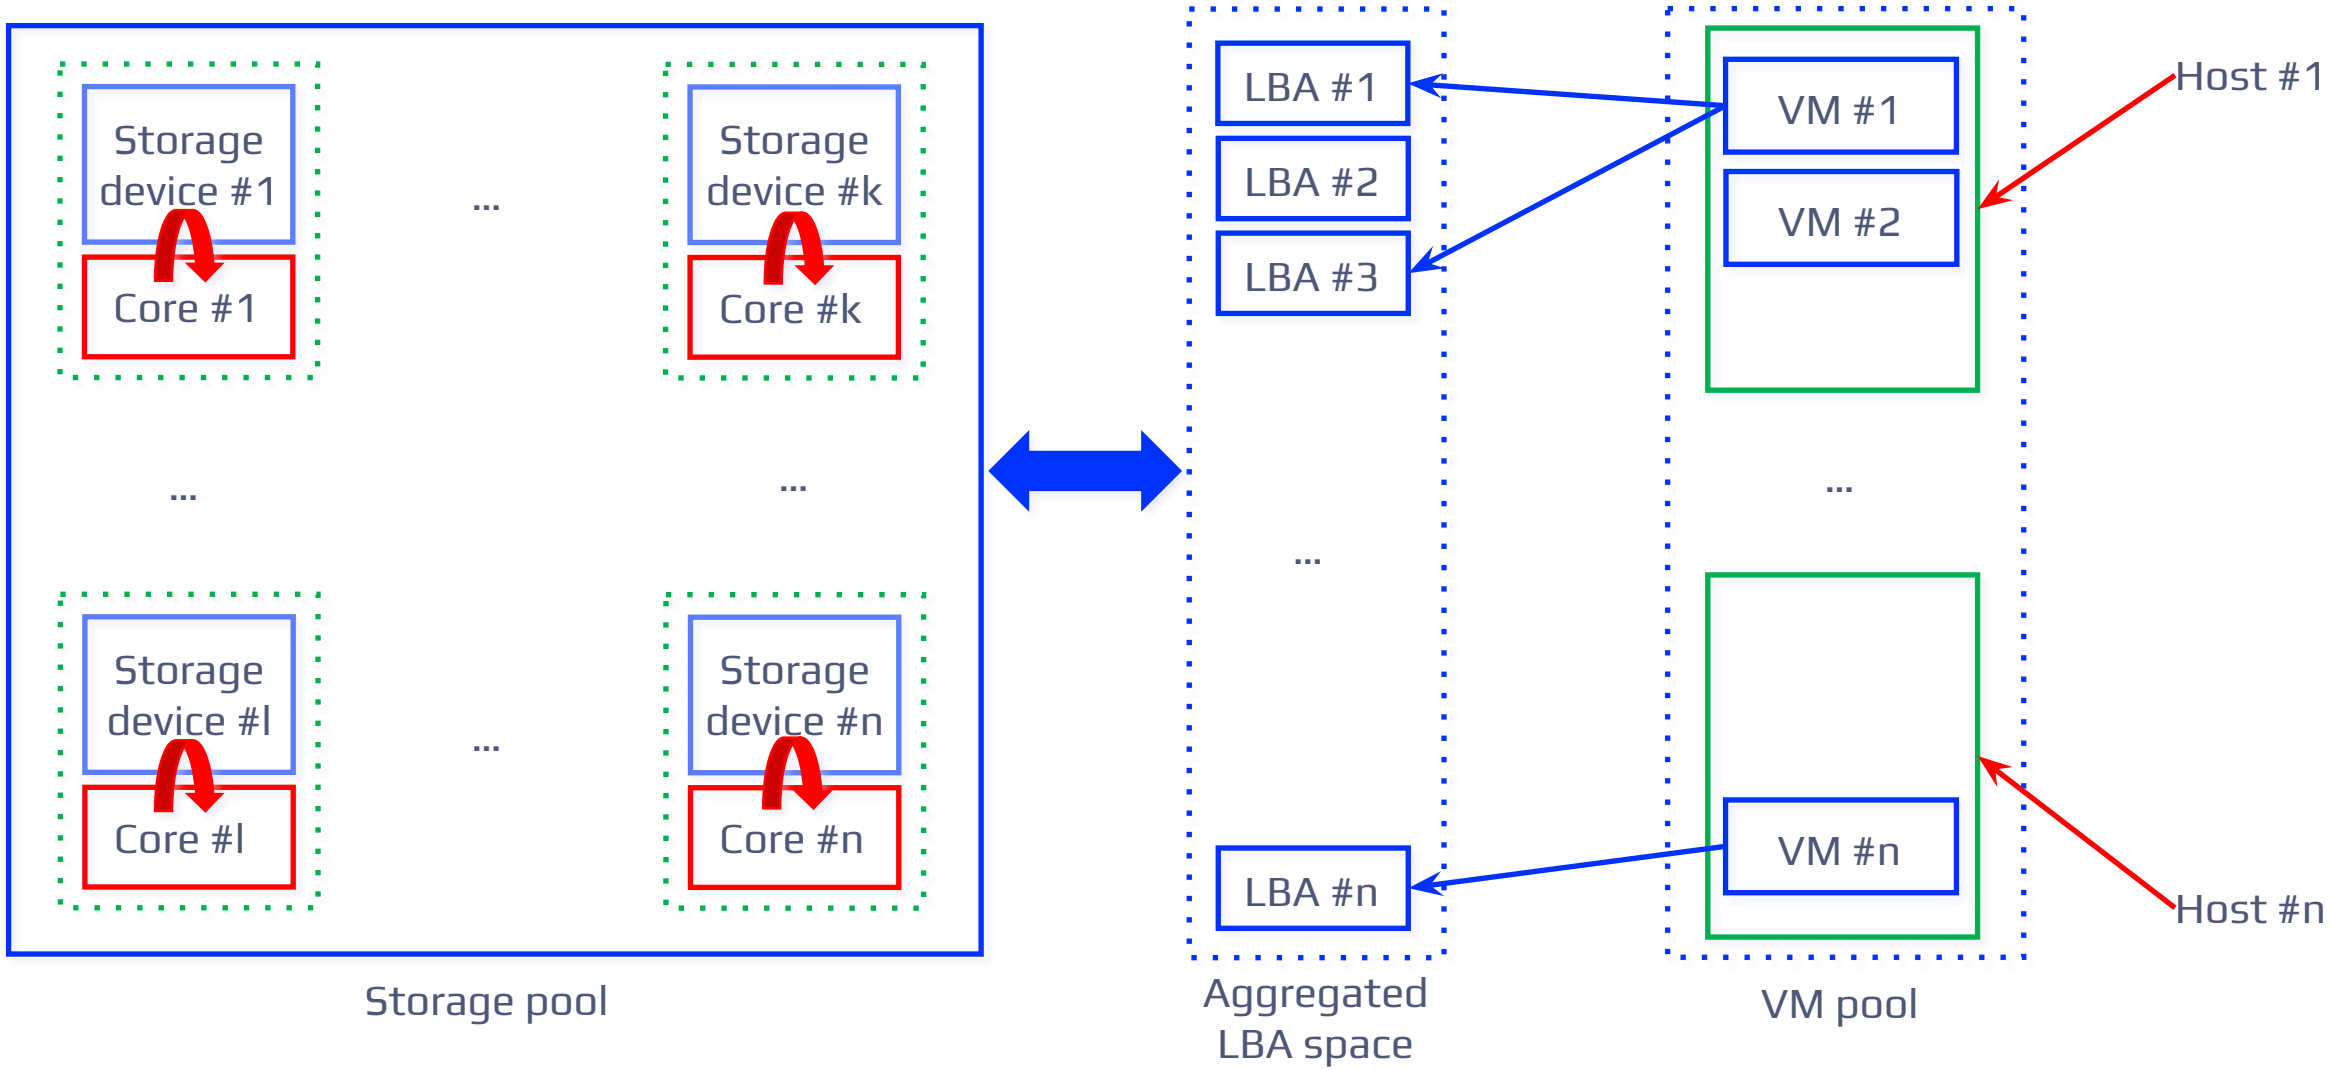

Storage pool space management problem

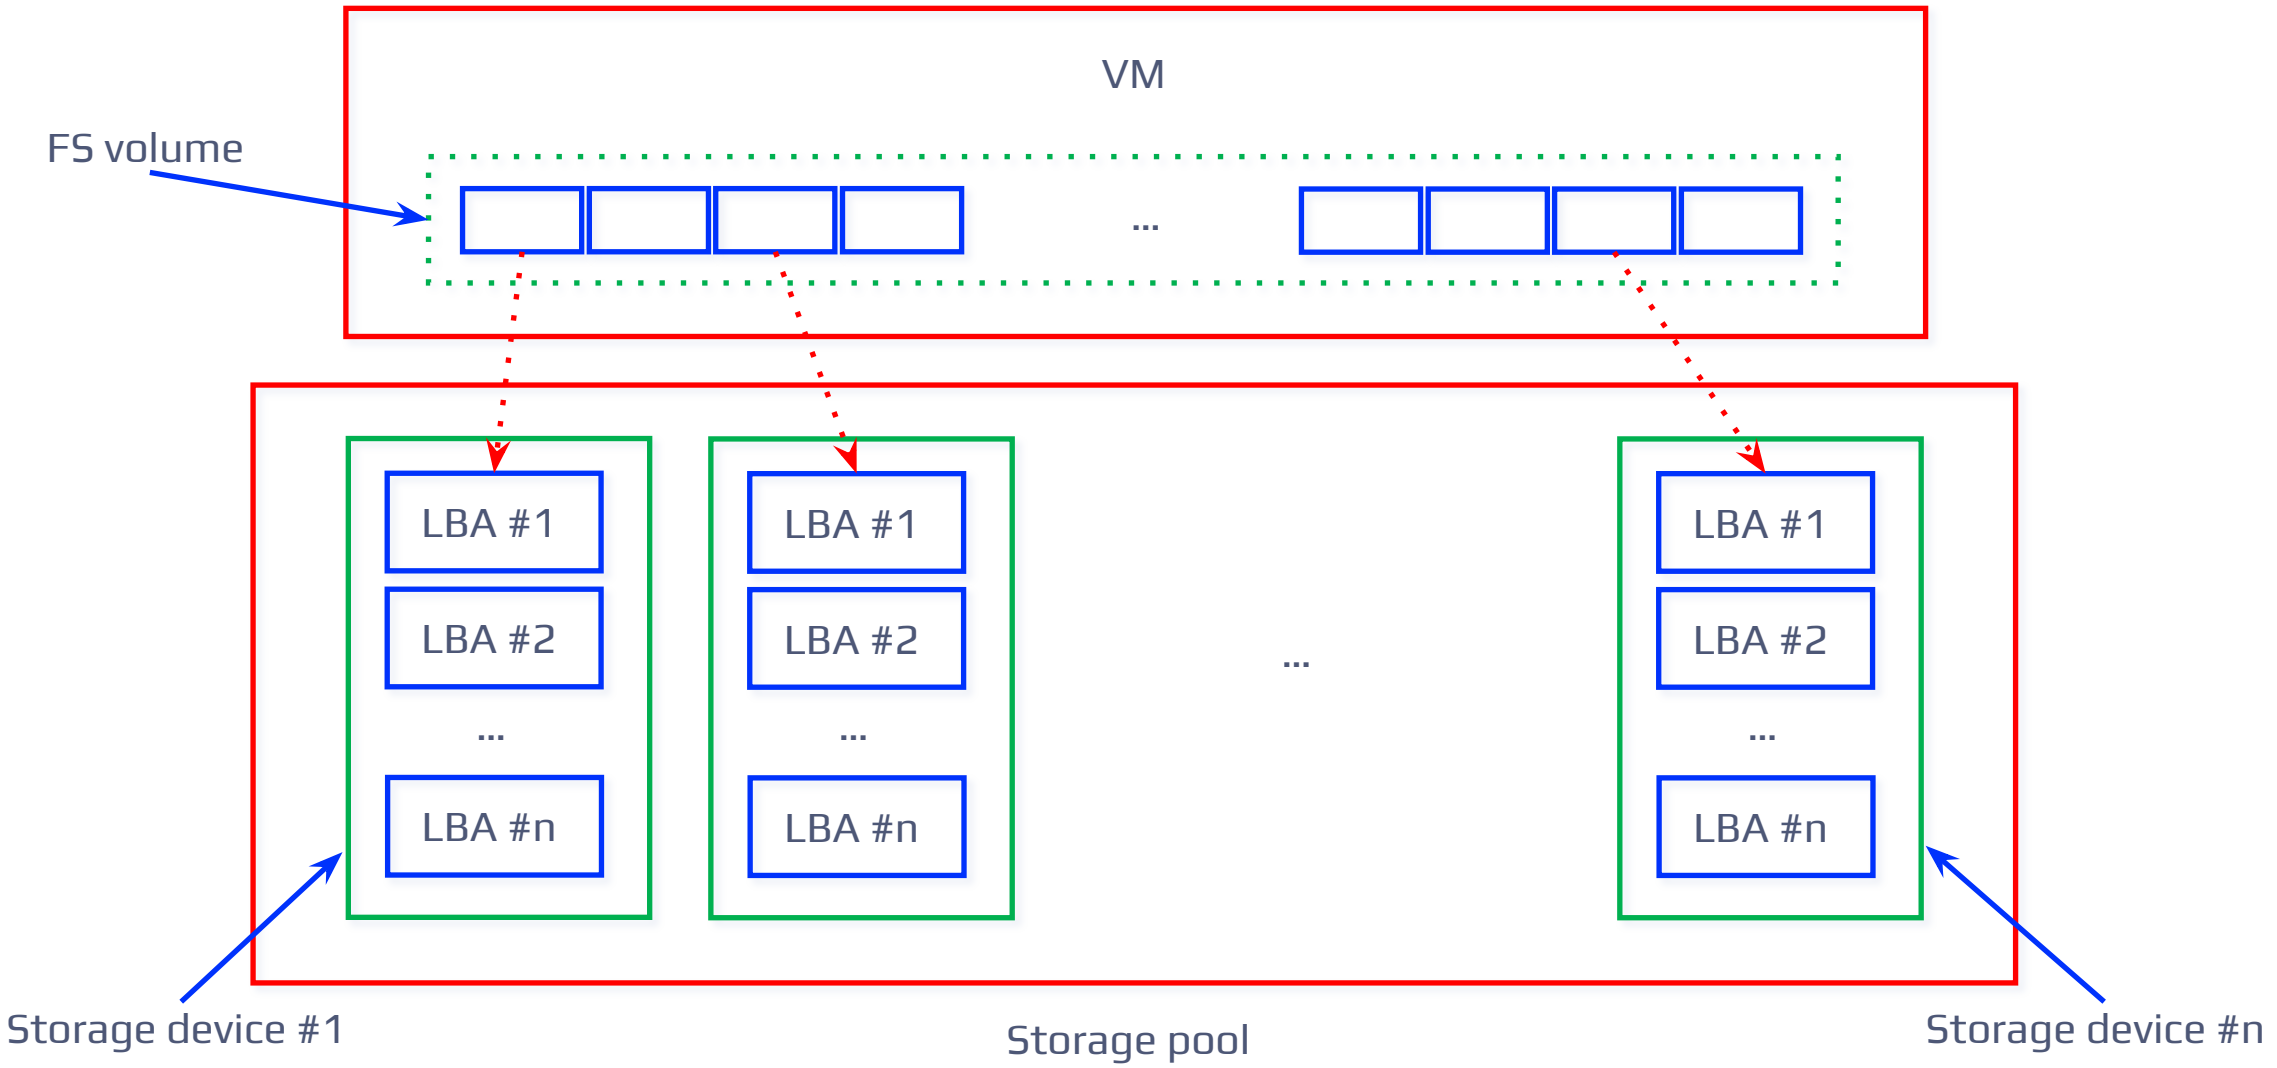

Page-table like architecture

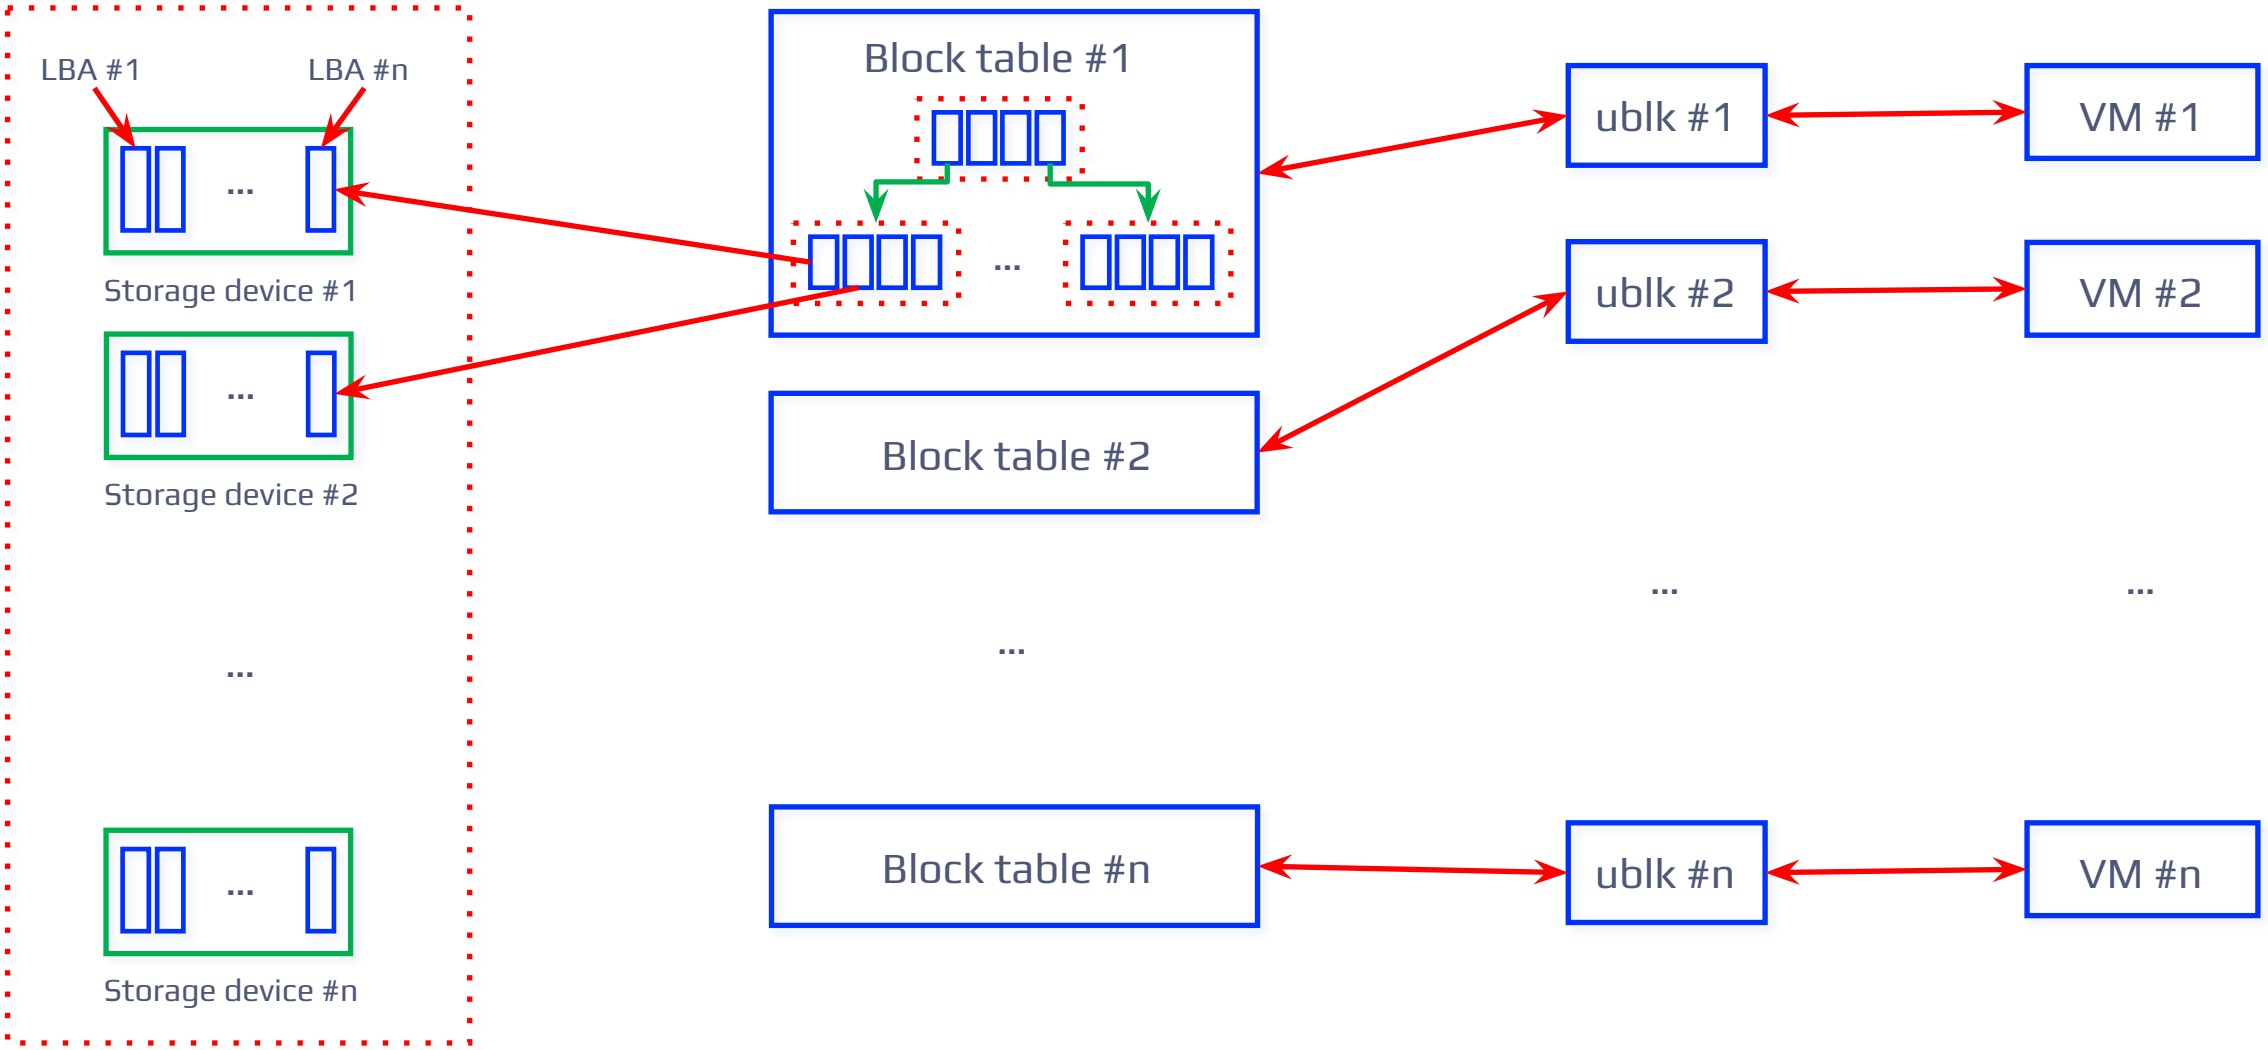

Flexible resize scheme

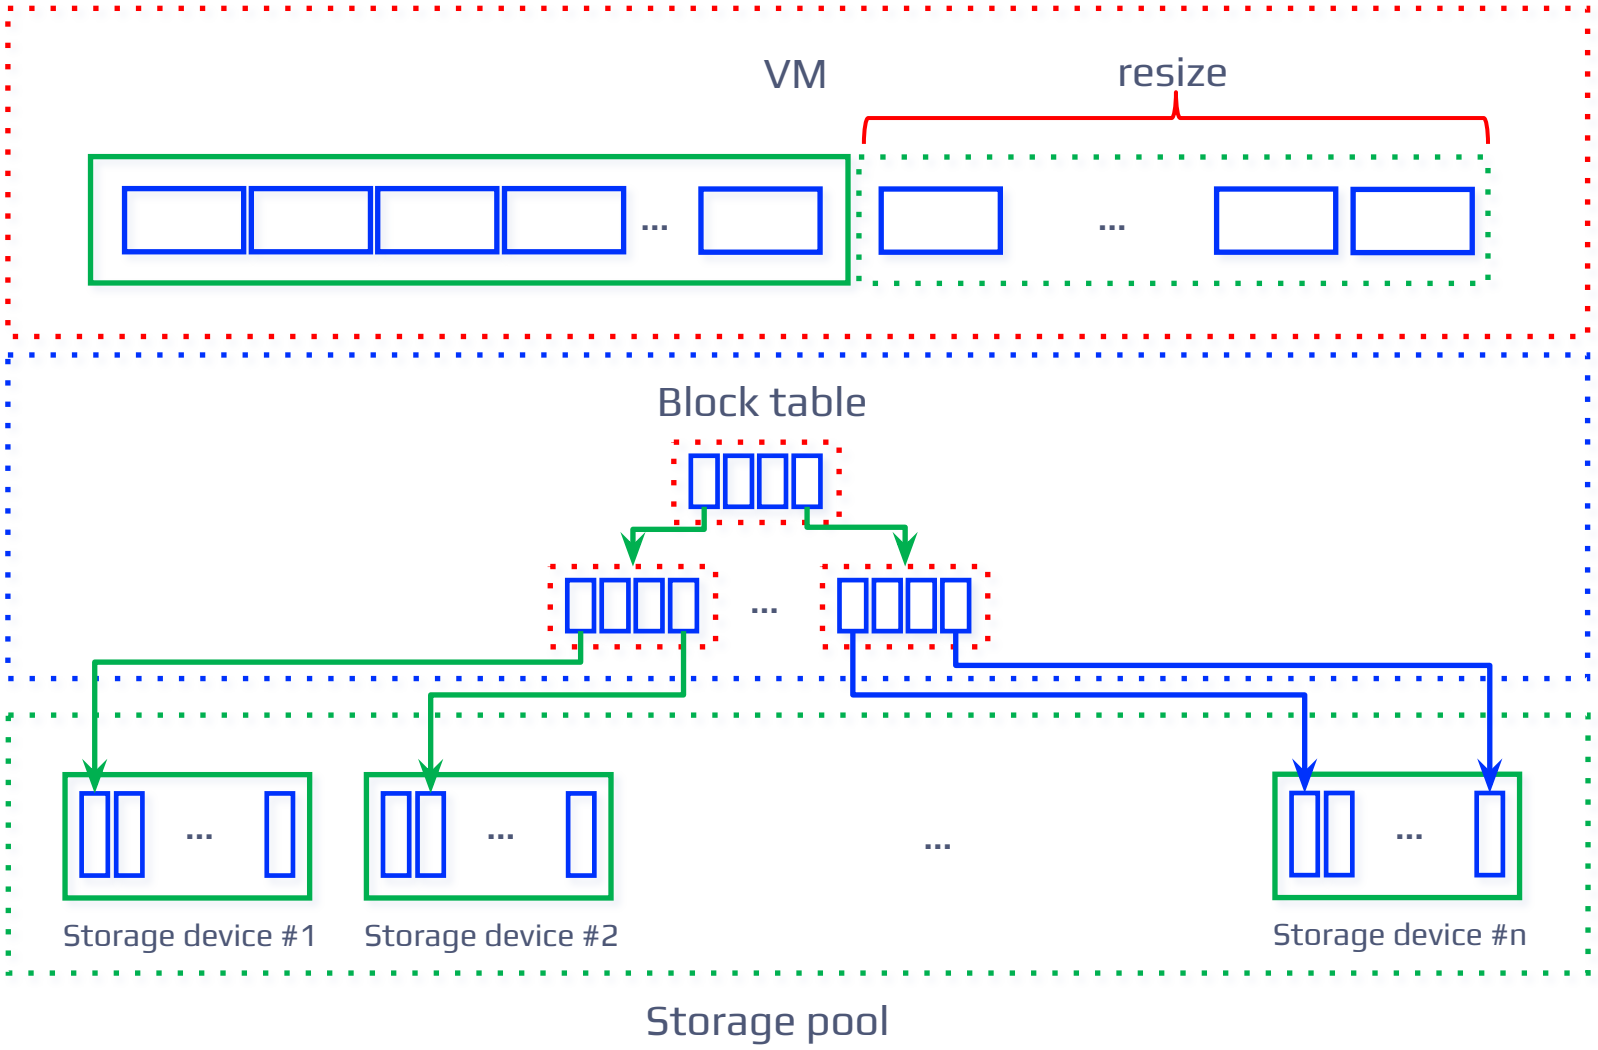

On-demand allocation scheme

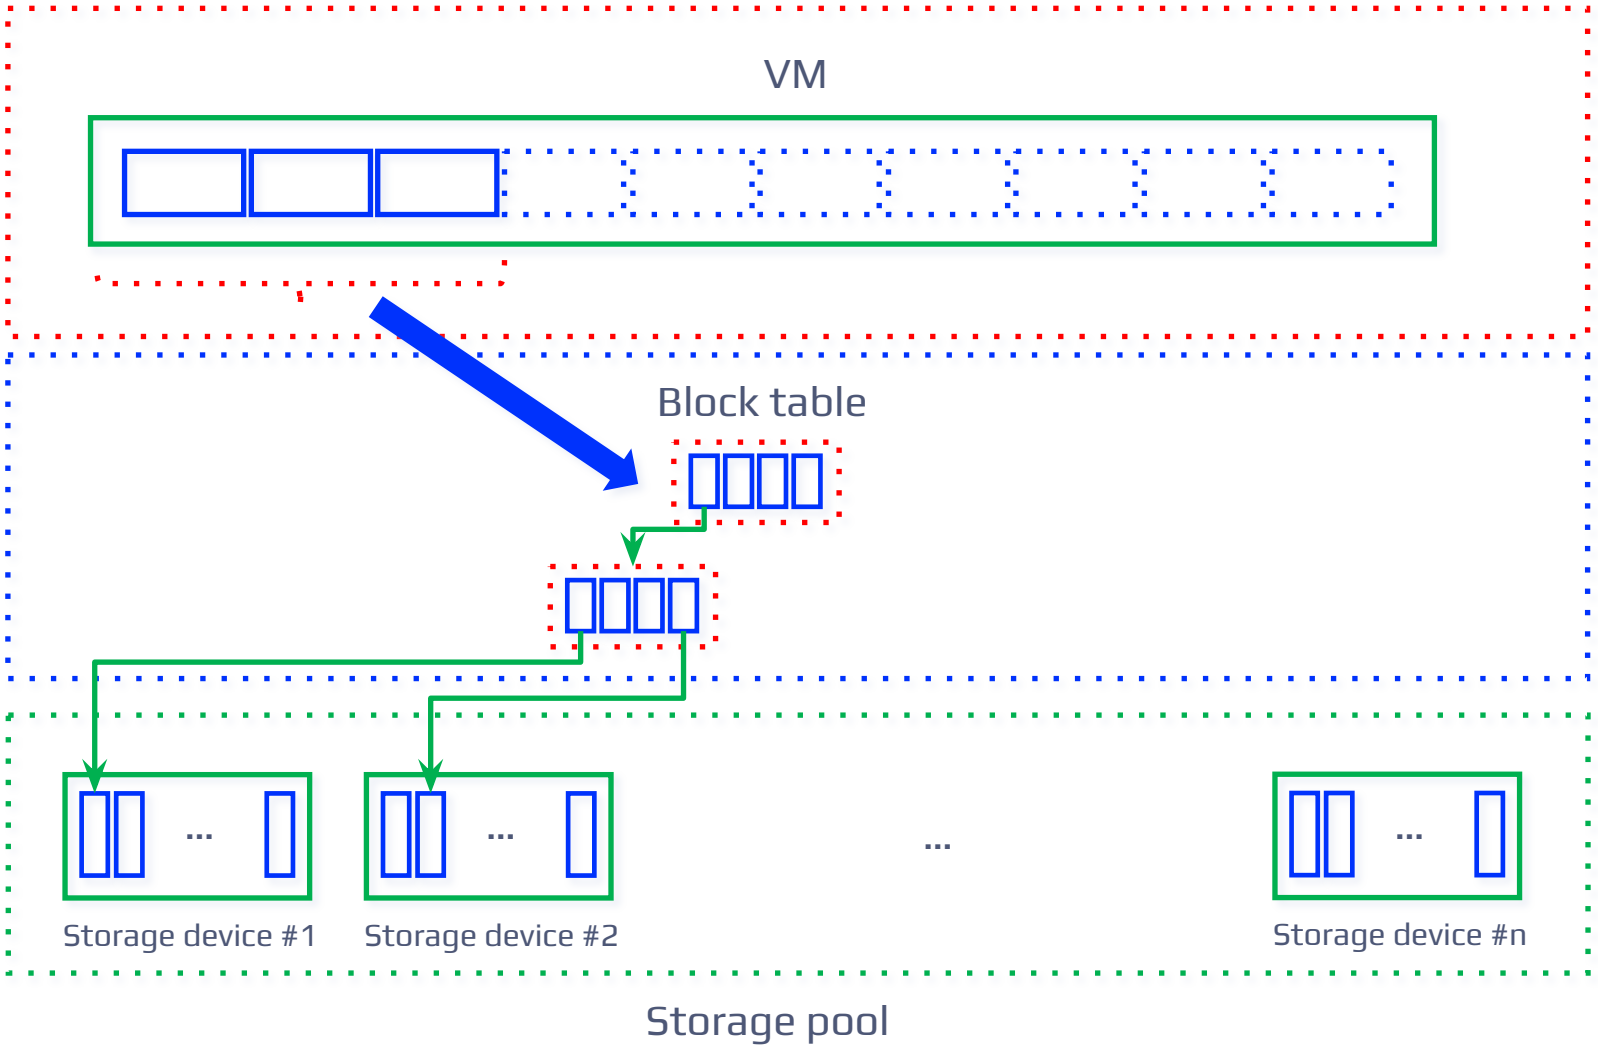

Easy migration scheme

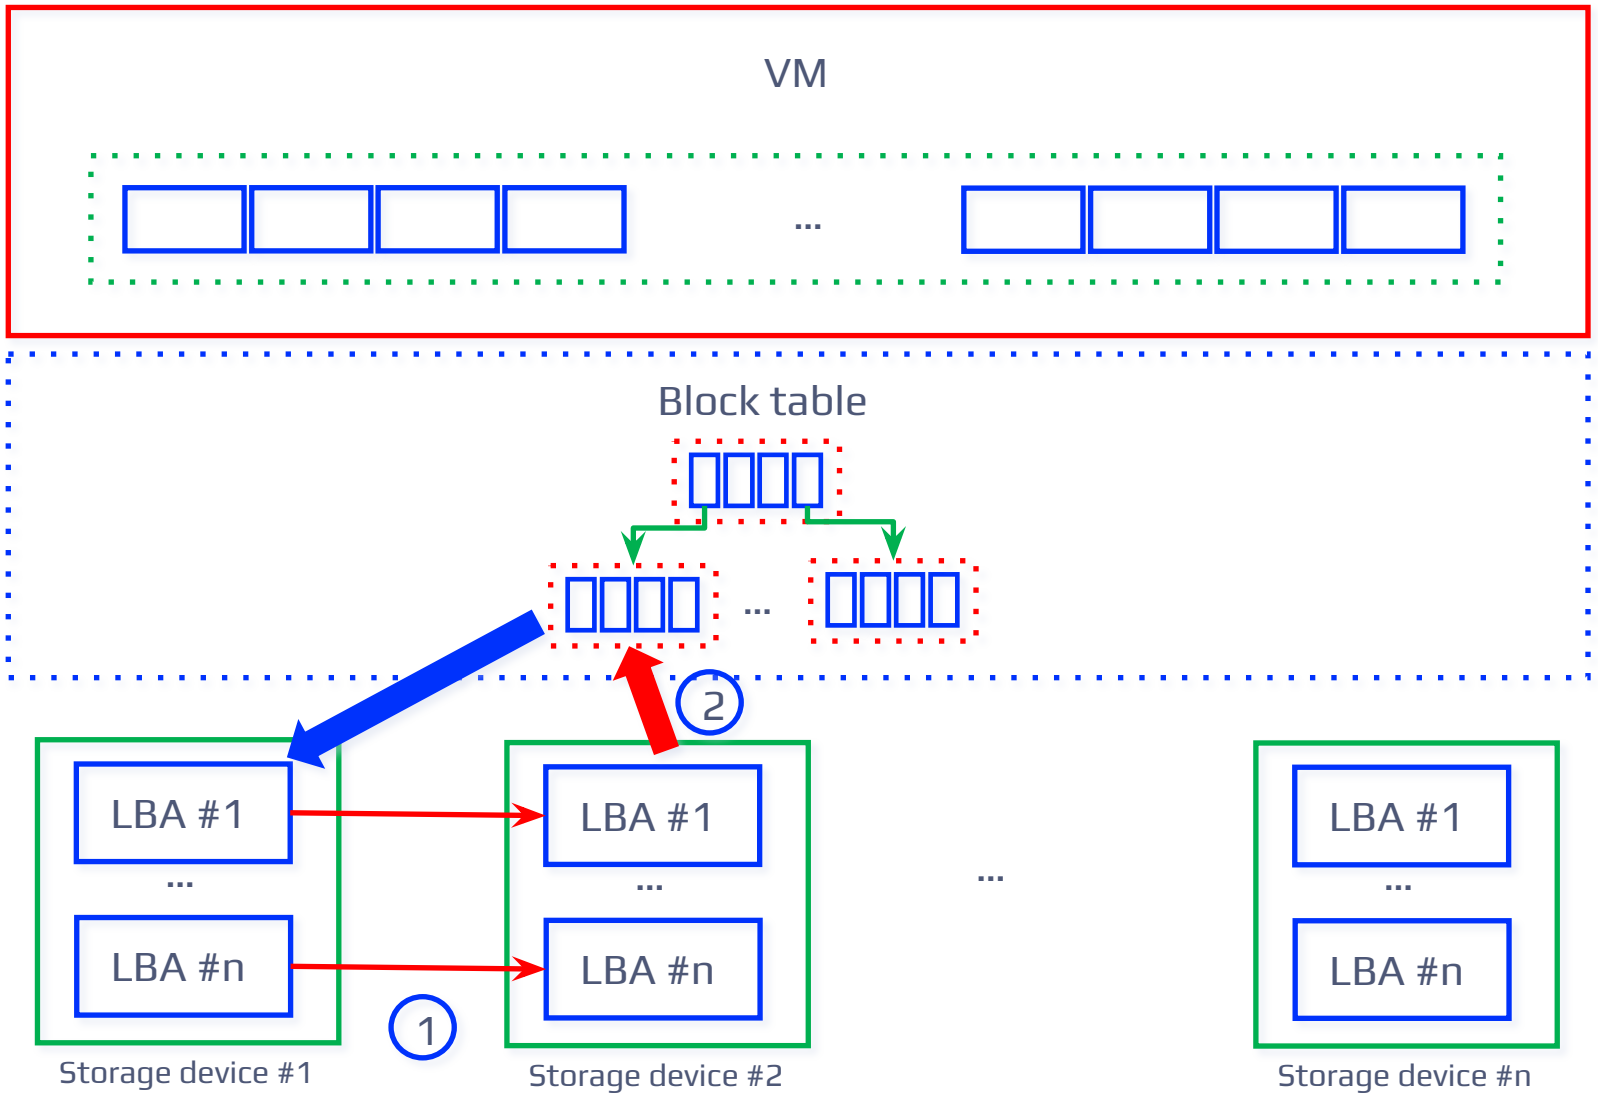

(1) Move
(2) Update block table

- COW policy
- ZNS SSD
- SMR HDD
- GC

Reliability scheme

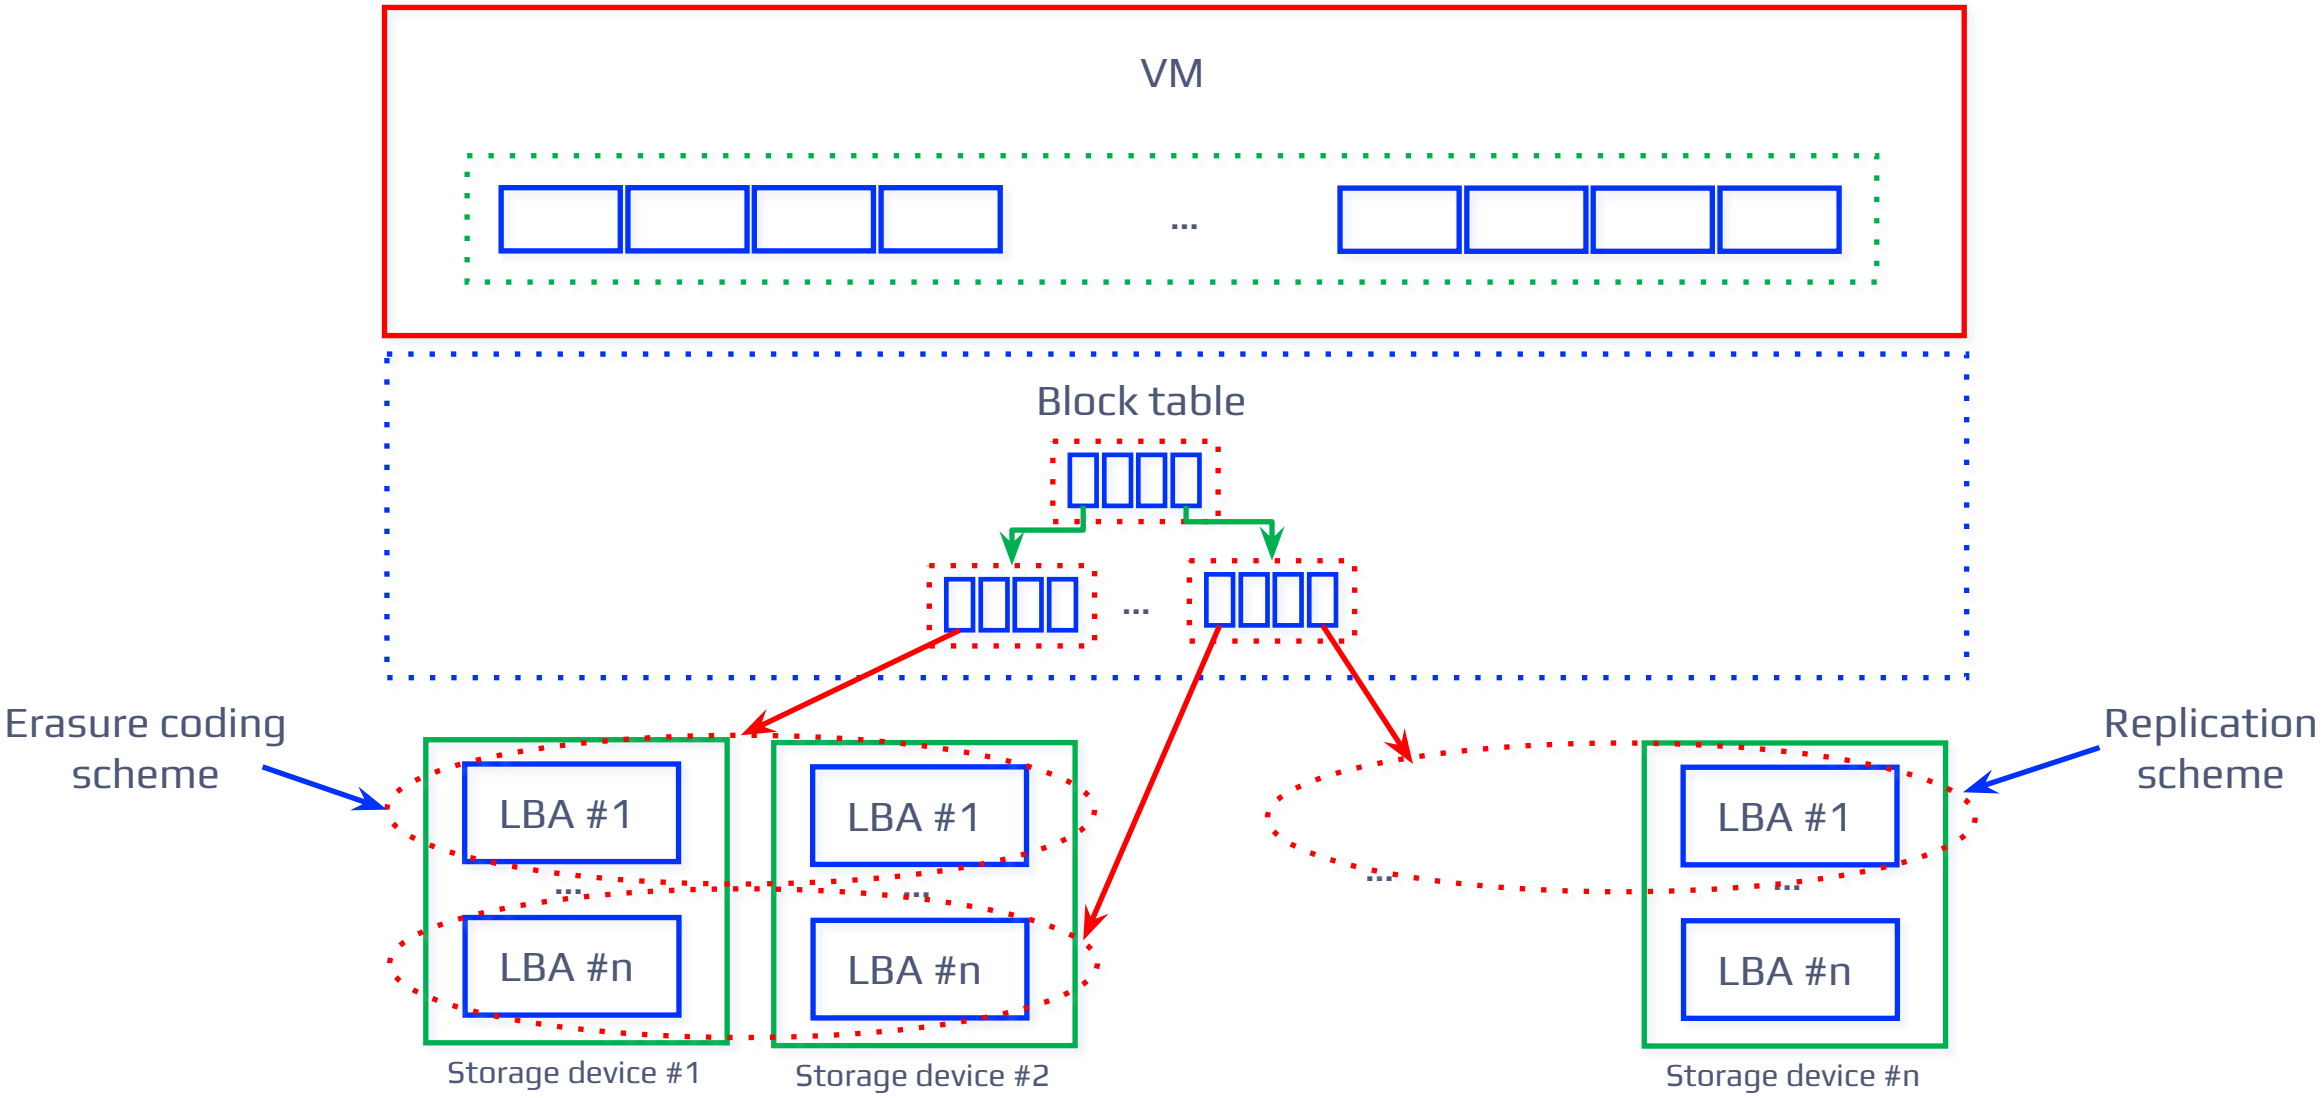

Computation offloading

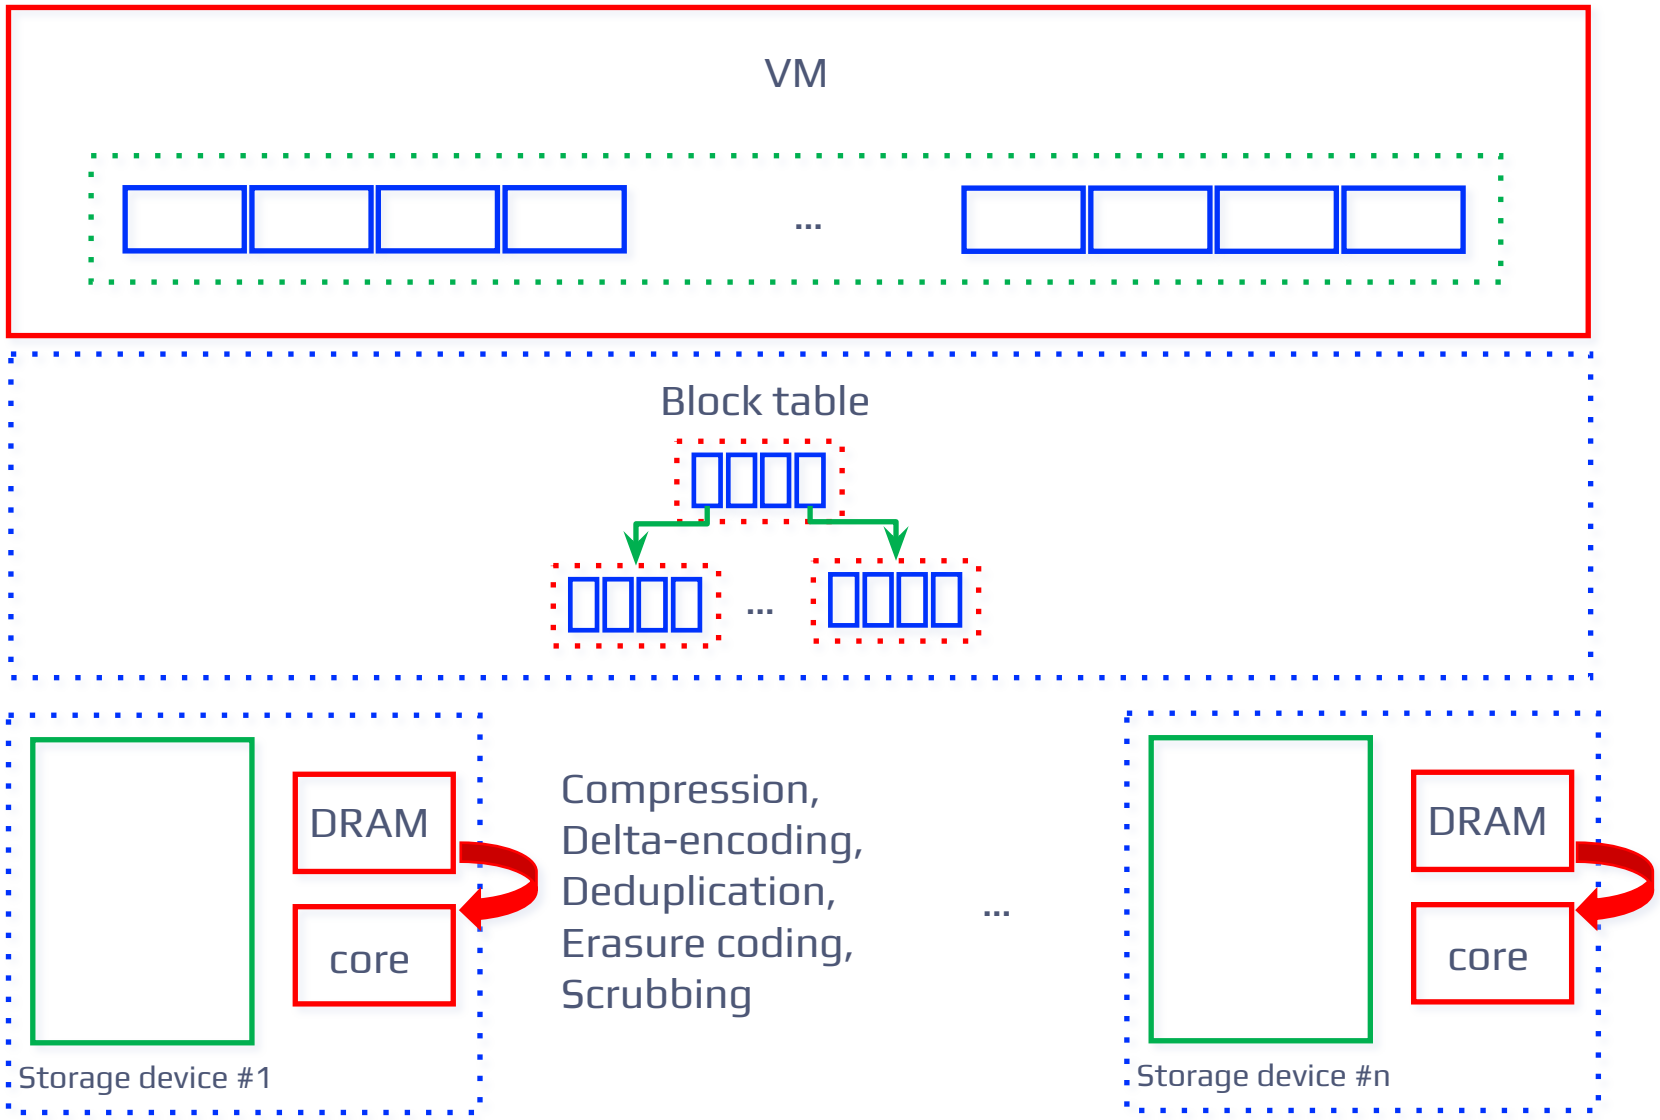

Aggregated LBA space management

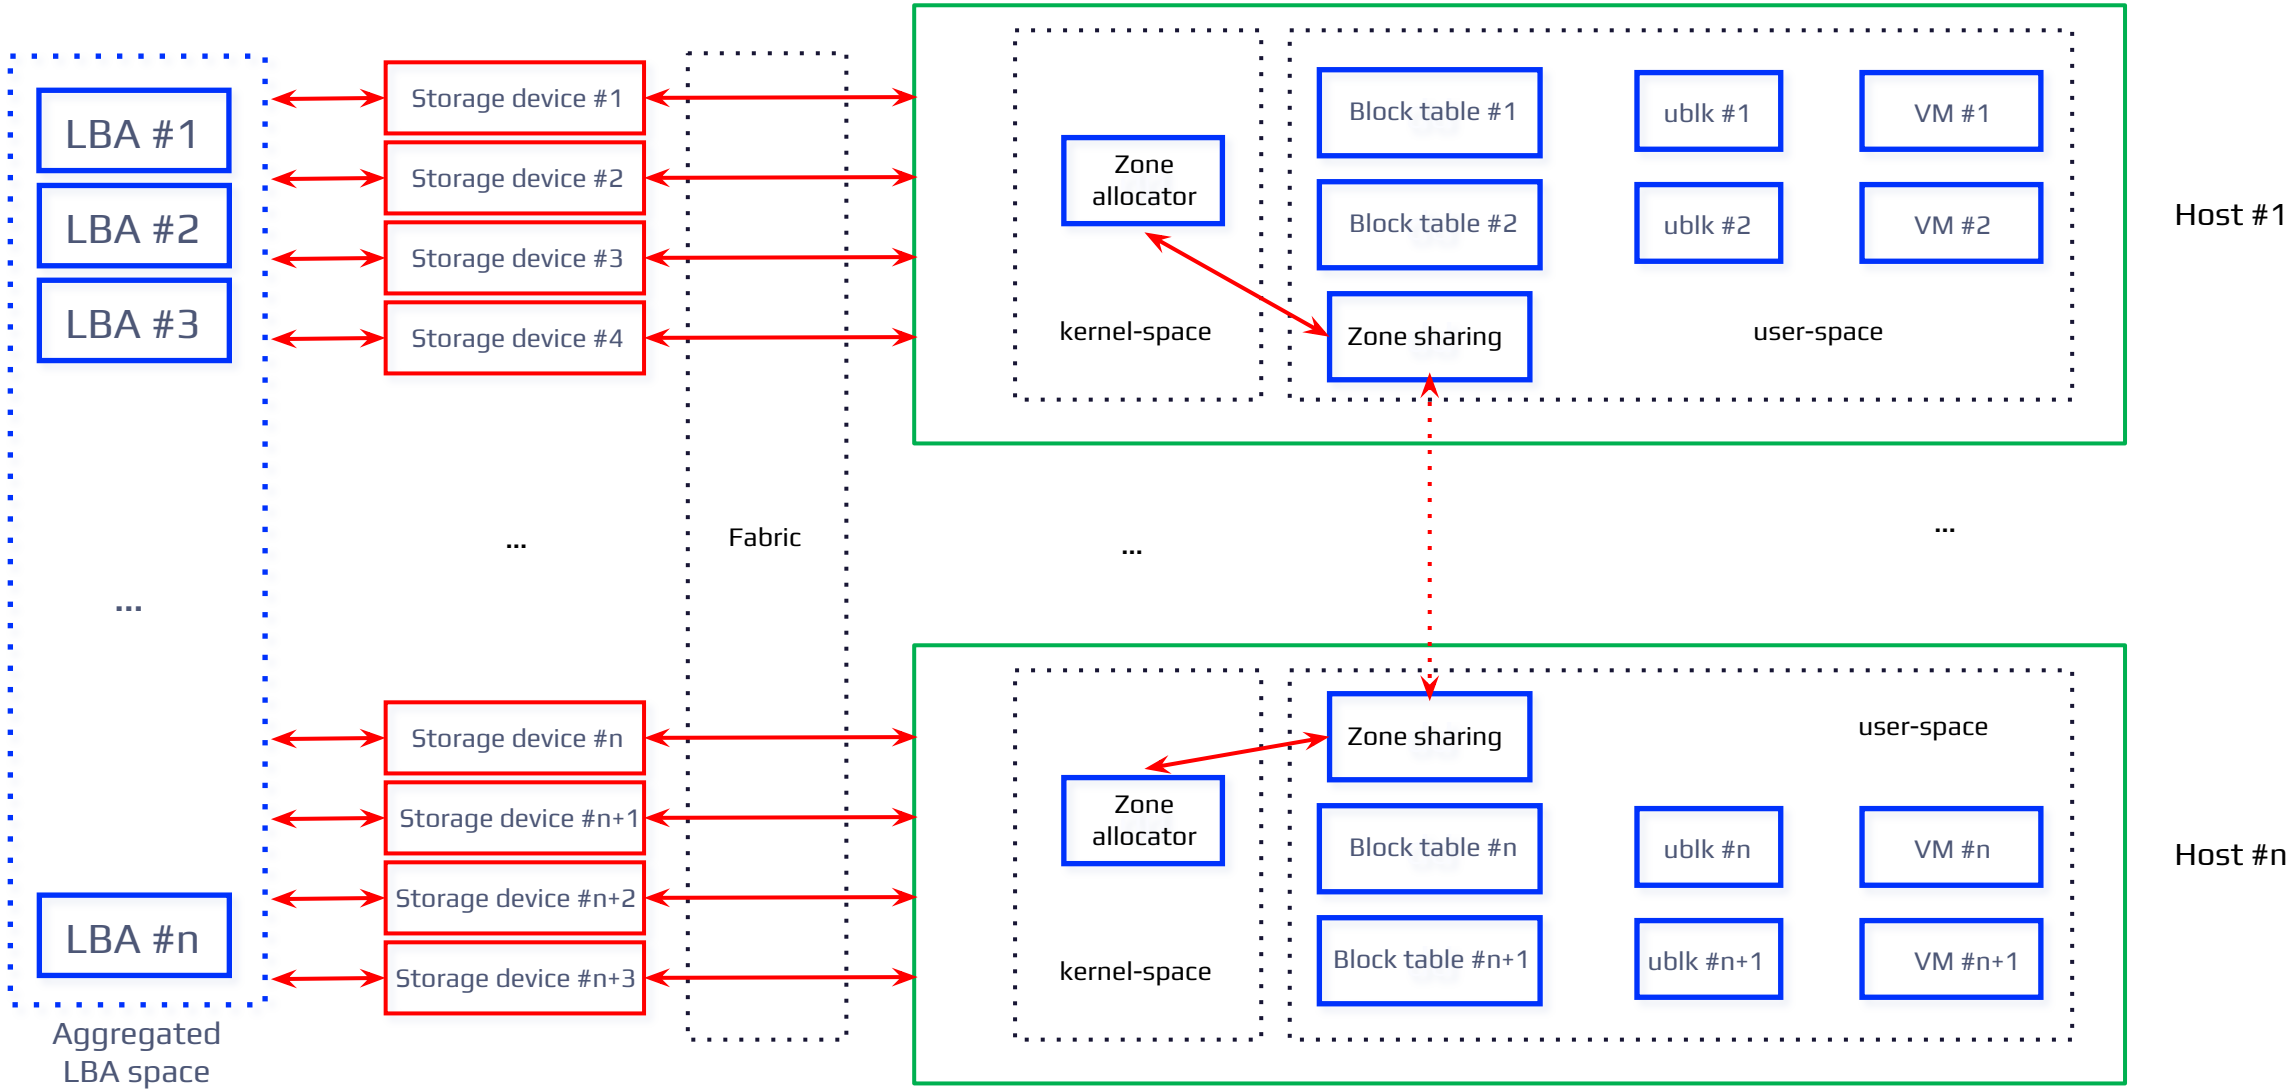

Open questions

- How to distribute the aggregated LBA space among virtual machines without the knowledge about each other?
- How to store block table for every VM?
- Should it be one block table for all VMs?
- Where does it make sense to keep the block table (in user-space or kernel-space)?
- Does it make sense to keep the aggregated LBA space allocator in kernel-space and block table for each VM in user-space?


THANK YOU
QUESTIONS???
slava@dubeyko.com