

Virtuozzo

QEMU BACKUP STATUS

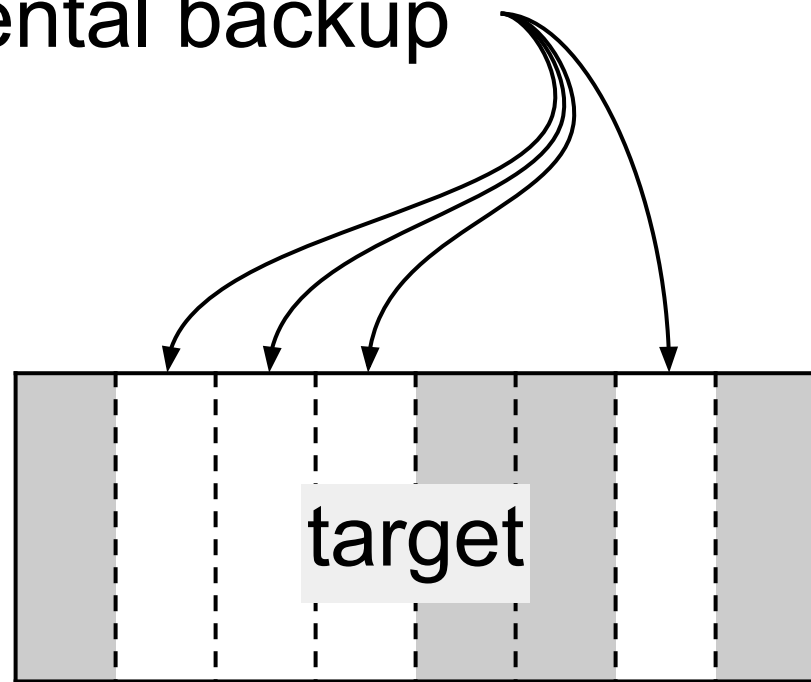
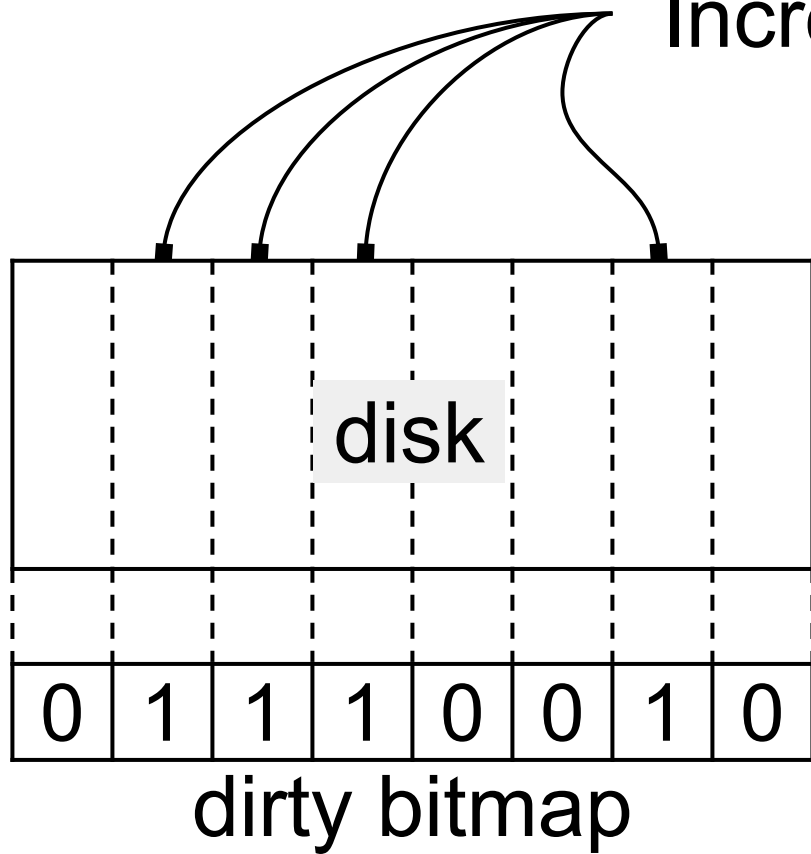
Vladimir Sementsov-Ogievskiy

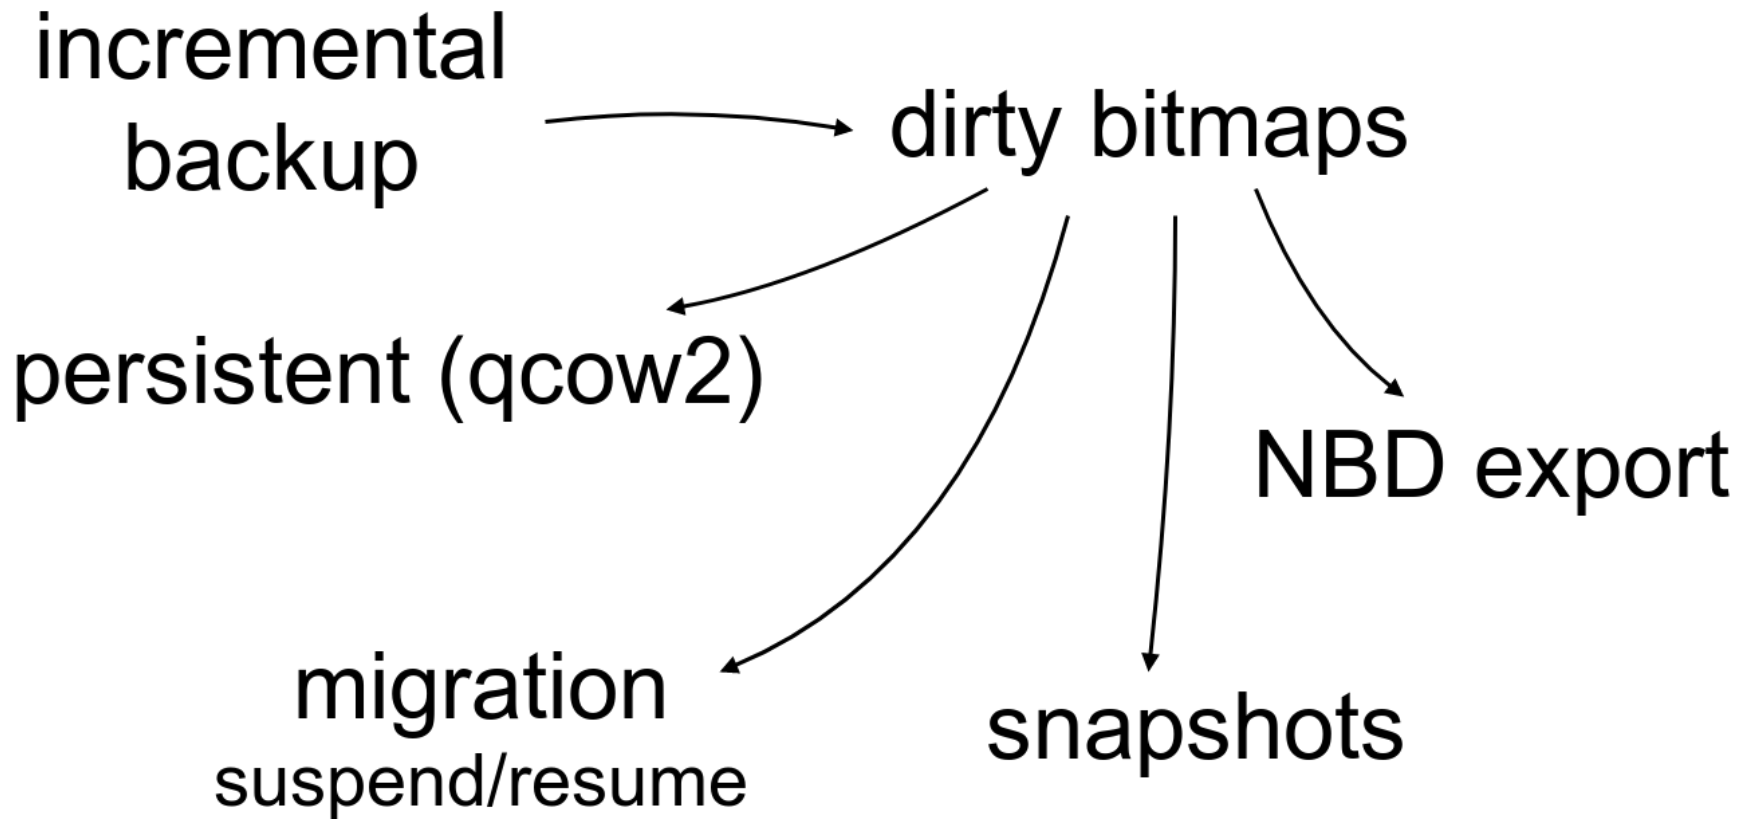
October 2018

QEMU Backup

- Incremental
- External API
- Architecture & performance

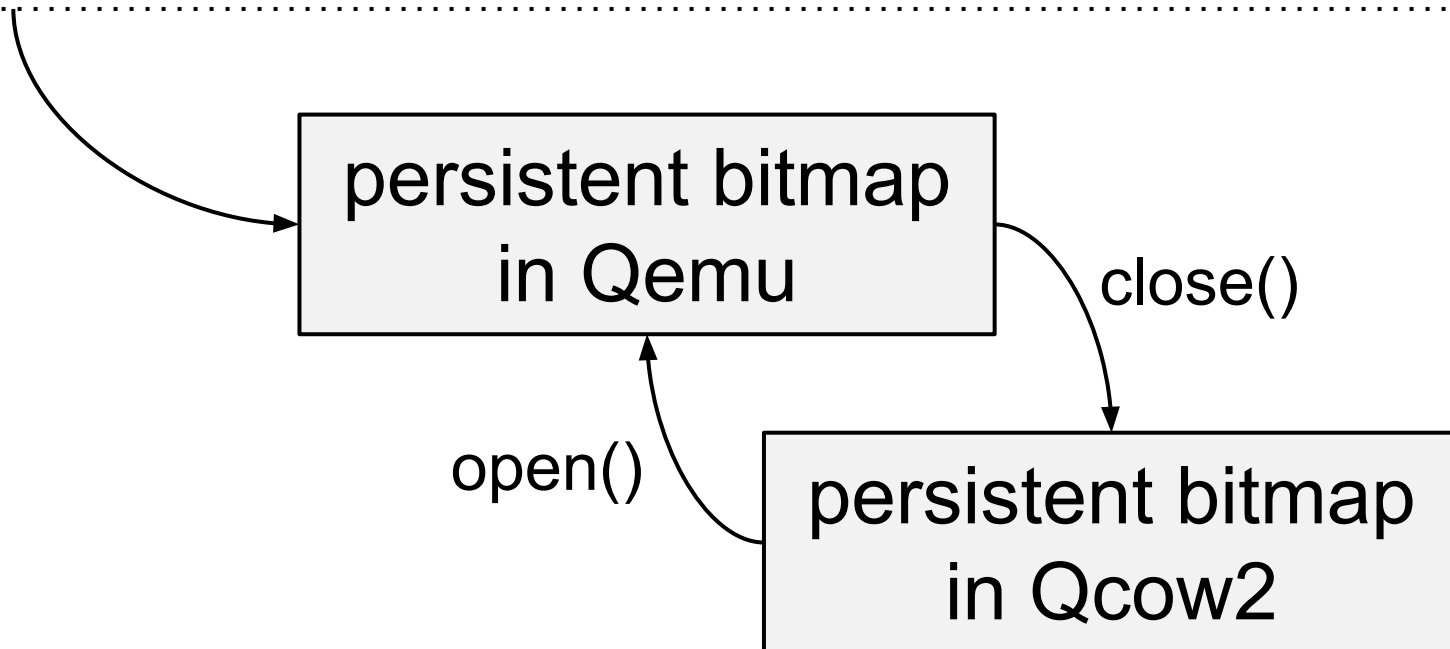
Incremental backup





Bitmaps: persistent (2.10)

qmp: block-dirty-bitmap-add (persistent=true)

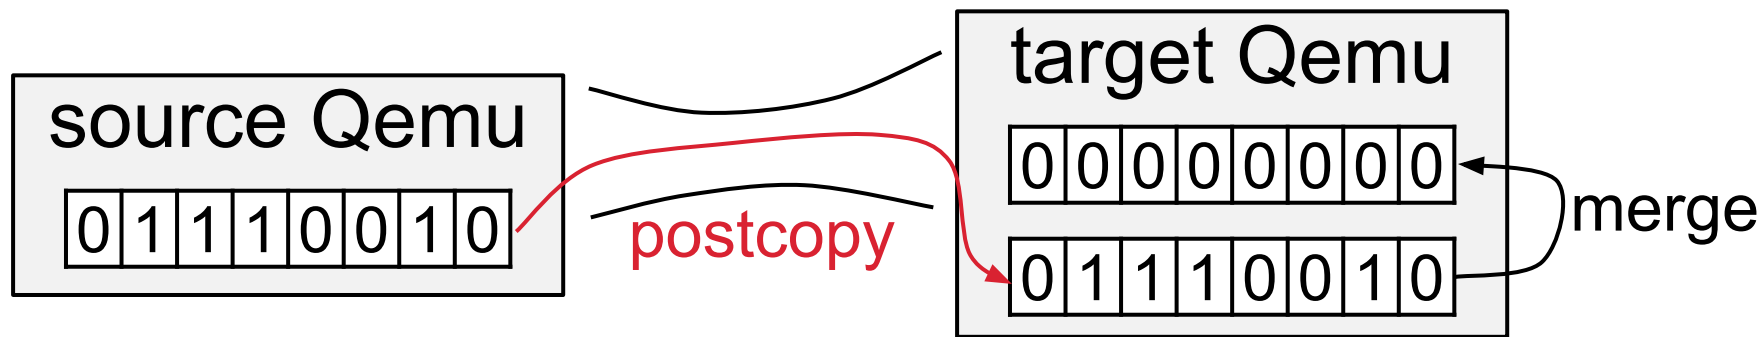


Bitmaps: migration (2.12)

qmp: migrate-set-capabilities (capabilities=
[capability: **dirty-bitmaps**, state: True])

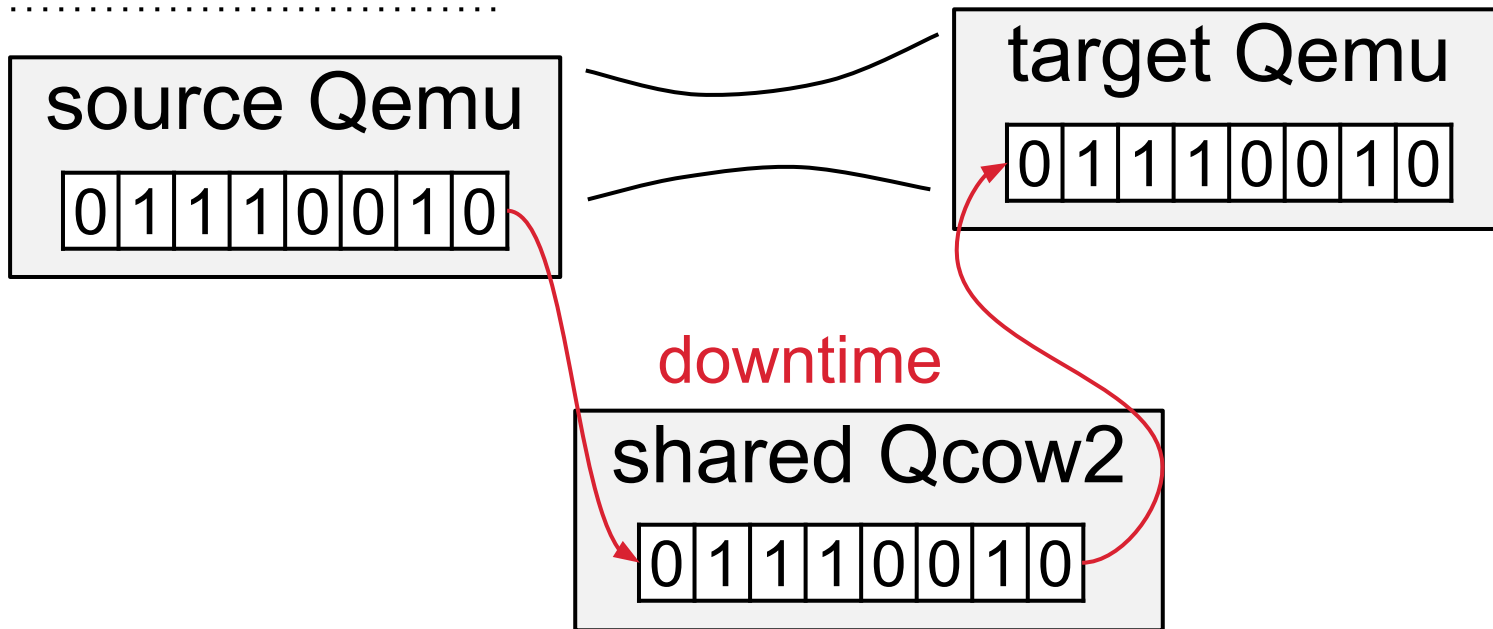
qmp: migrate

qmp: migrate-start-postcopy



Bitmaps: migration through storage (2.10)

qmp: migrate



Bitmaps: snapshots

make a snapshot	switch to the snapshot
do nothing	all bitmaps are invalid <ul style="list-style-type: none">• mark them all dirty• drop all the bitmaps• <u>error</u>

Bitmaps: snapshots


make a snapshot	switch to the snapshot
do nothing	all bitmaps are invalid <ul style="list-style-type: none">• mark them all dirty• drop all the bitmaps• <u>error</u>
enable dirty-bitmaps migration capability	drop all the bitmaps and load them from vmstate

Bitmaps: snapshots

Legend:

snapshot

active state




Bitmaps: snapshots

Legend:

snapshot

active state

incremental backup



Bitmaps: snapshots


Legend:

snapshot

active state

incremental backup

0	1	1	1	1	0	0	1	0
---	---	---	---	---	---	---	---	---



Bitmaps: snapshots


Legend:

snapshot

active state

incremental backup

0	1	1	1	1	0	0	1	0
---	---	---	---	---	---	---	---	---




Bitmaps: snapshots

Legend:

snapshot

active state

incremental backup



Bitmaps: NBD Export (3.0)

Negotiation phase: select metadata contexts

- Commands:
 - NBD_OPT_LIST_META_CONTEXT
 - NBD_OPT_SET_META_CONTEXT
- base:allocation - common block status (2.12)
- qemu:dirty-bitmap:<bitmap-export-name> (3.0)

Transmission phase: get metadata

- NBD_CMD_BLOCK_STATUS

Bitmaps: QMP API


basic (2.4)

- block-dirty-bitmap-add
 - persistent (2.10)
 - x-disabled (3.0)
- block-dirty-bitmap-remove
- block-dirty-bitmap-clear


new (3.0)

- x-block-dirty-bitmap-enable
- x-block-dirty-bitmap-disable
- x-block-dirty-bitmap-merge
 - node
 - dst_name
 - src_name
- x-nbd-server-add-bitmap
 - name
 - bitmap
 - bitmap-export-name


External backup API: image fleecing (3.0)




External backup API: image fleecing (next)



External backup API: image fleecing (next)



External backup API: image fleecing (next)



- extra CBW operations
- extra save/load to temp

Backup architecture and performance

Guest write hangs

- don't wait for write to far target (fleecing to temp)
- don't wait for write at all if possible
- priority of CBW

Total number of iops

- don't write to temp if possible
- share reads

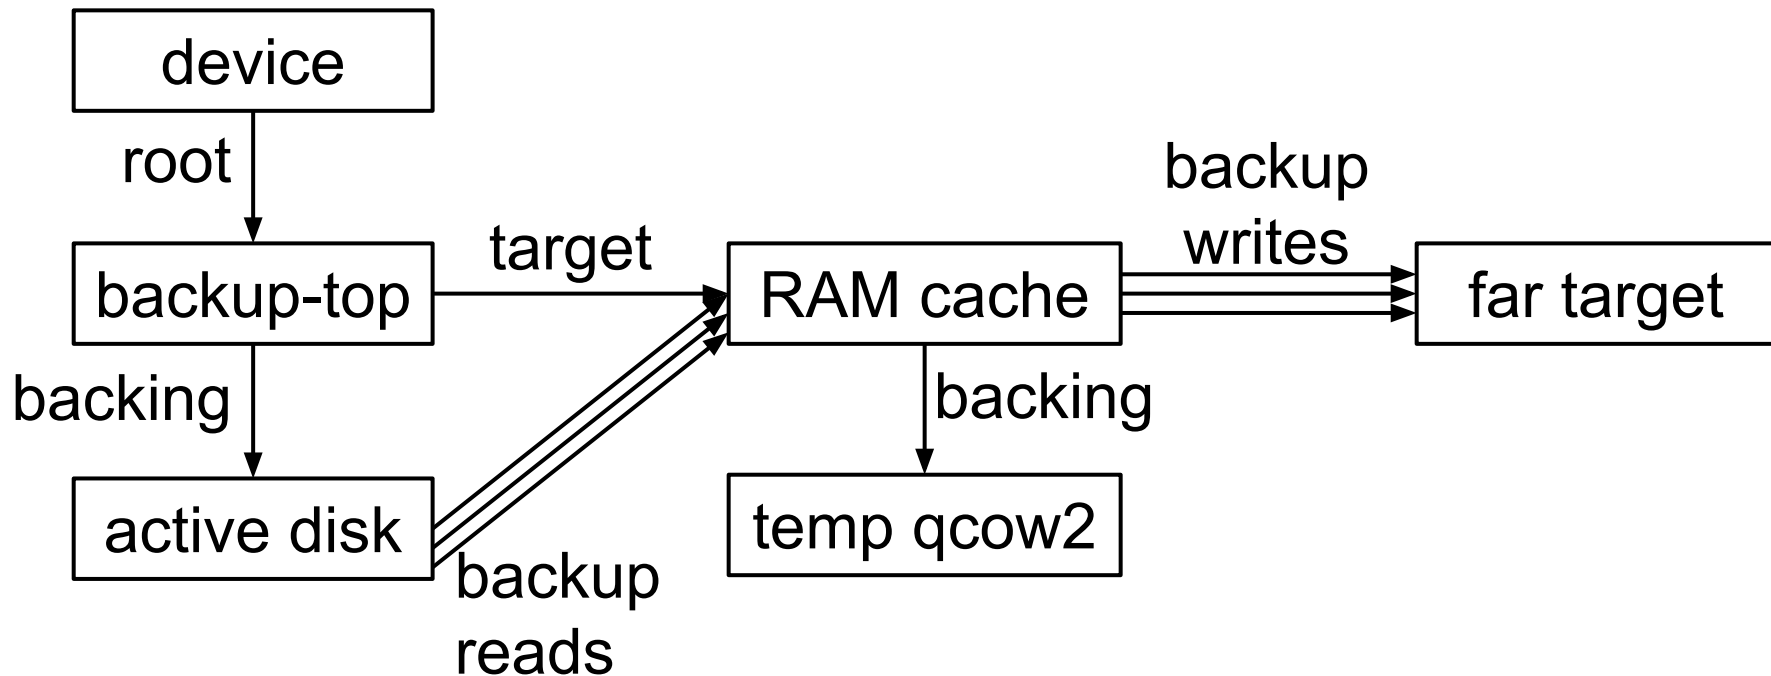
Decrease number of CBW

- lazy discards
- copy from temp at last (or at first, optionally?)
- predict guest writes

Optimize copying loop

- asynchronous copying like in mirror instead of sequential

Backup architecture and performance



Qemu backup status

Merged

- Qcow2 bitmaps (2.10)
- Migration (2.12) ⇒ suspend/resume & snapshots
- Fleecing (3.0)
- NBD BLOCK_STATUS (3.0)
- QMP API (3.0, still with x- prefix)

Plans

- Backup architecture

Qemu backup status

Merged

- Qcow2 bitmaps (2.10)
- Migration (2.12) \Rightarrow suspend/resume & snapshots
- Fleecing (3.0)
- NBD BLOCK_STATUS (3.0)
- QMP API (3.0, still with x- prefix)

Plans

- Backup architecture

Vladimir Sementsov-Ogievskiy
vsementsov@virtuozzo.com