

The pNFS Page Sync State Model

Benny Halevy <bhalevy@panasas.com>


Connectathon 2009

Good Old NFS

Client



Good Old NFS



Good Old NFS

Client

Dirty

WRITE

Server

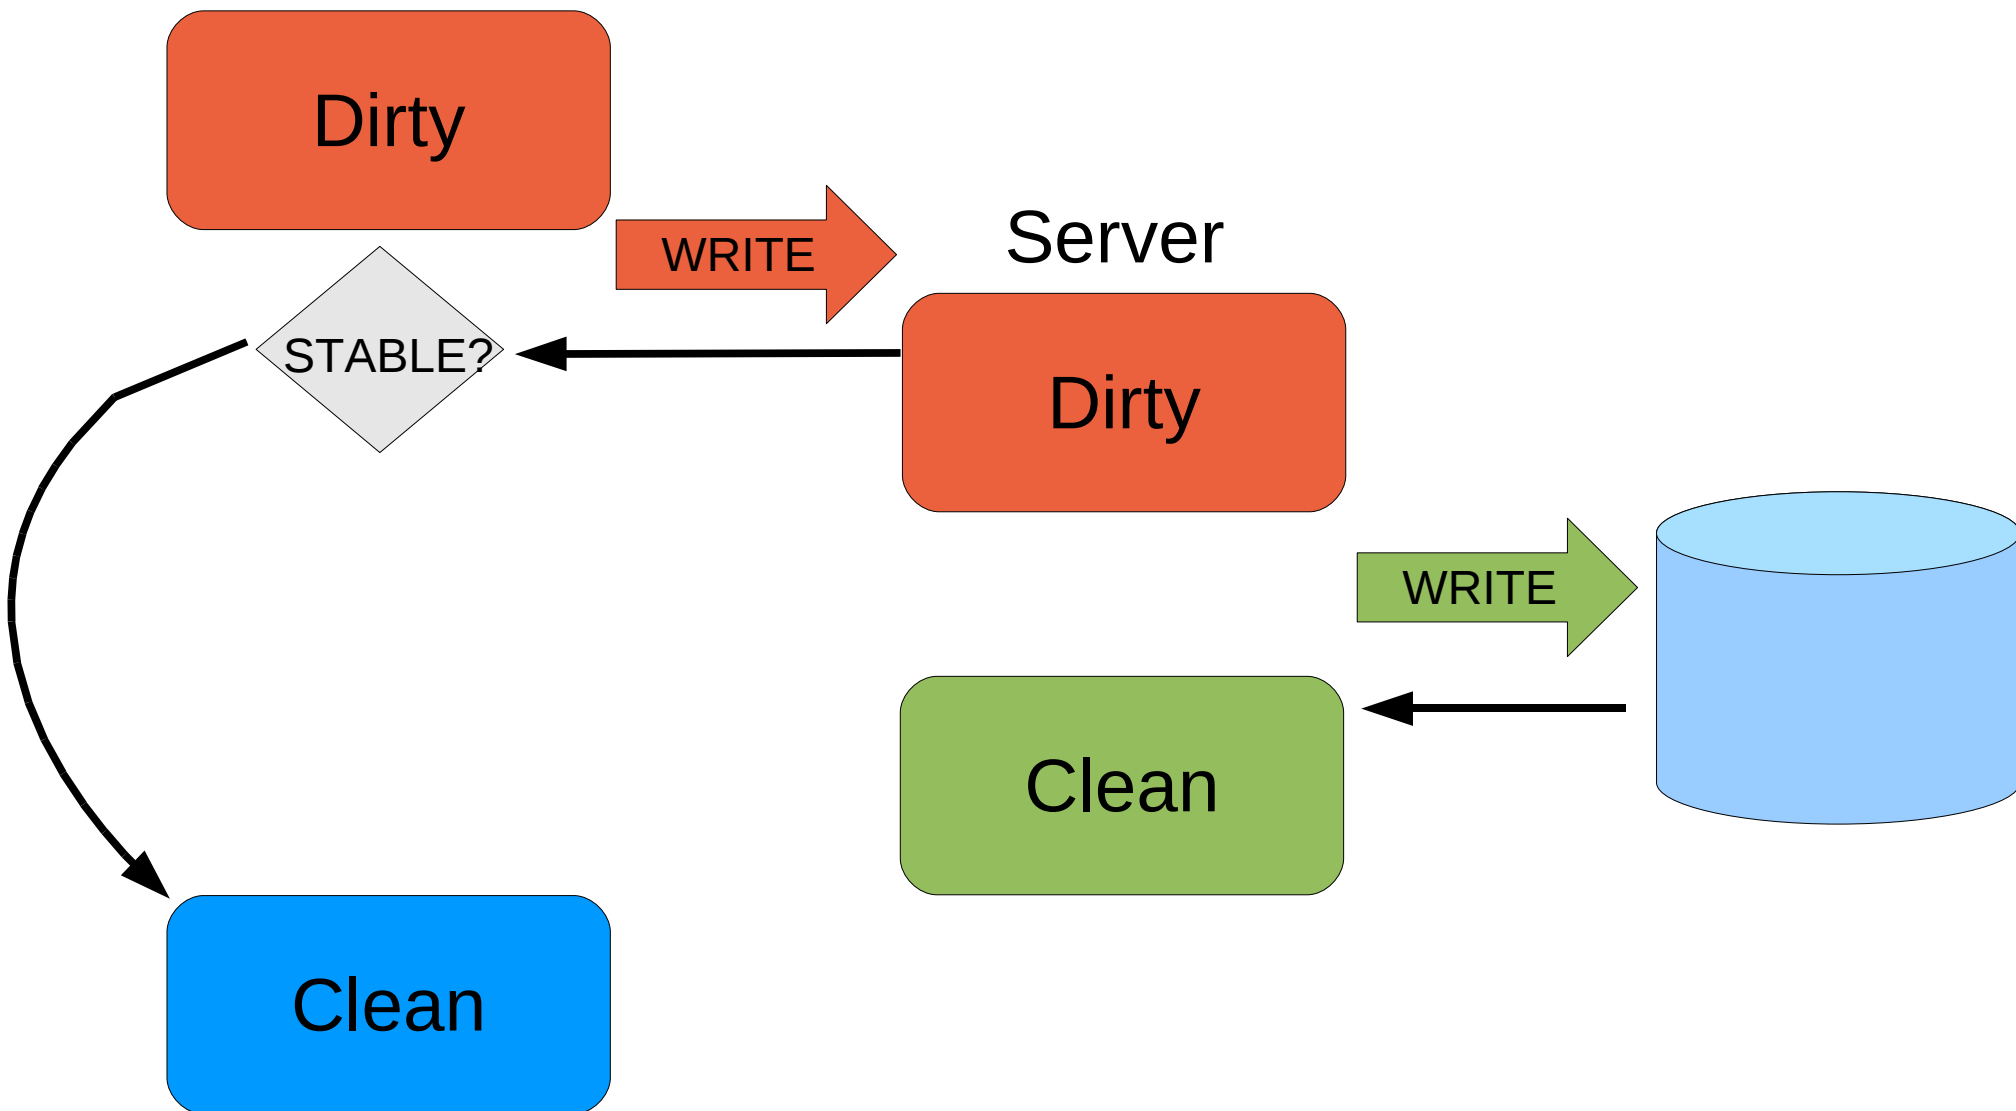
Dirty

STABLE?

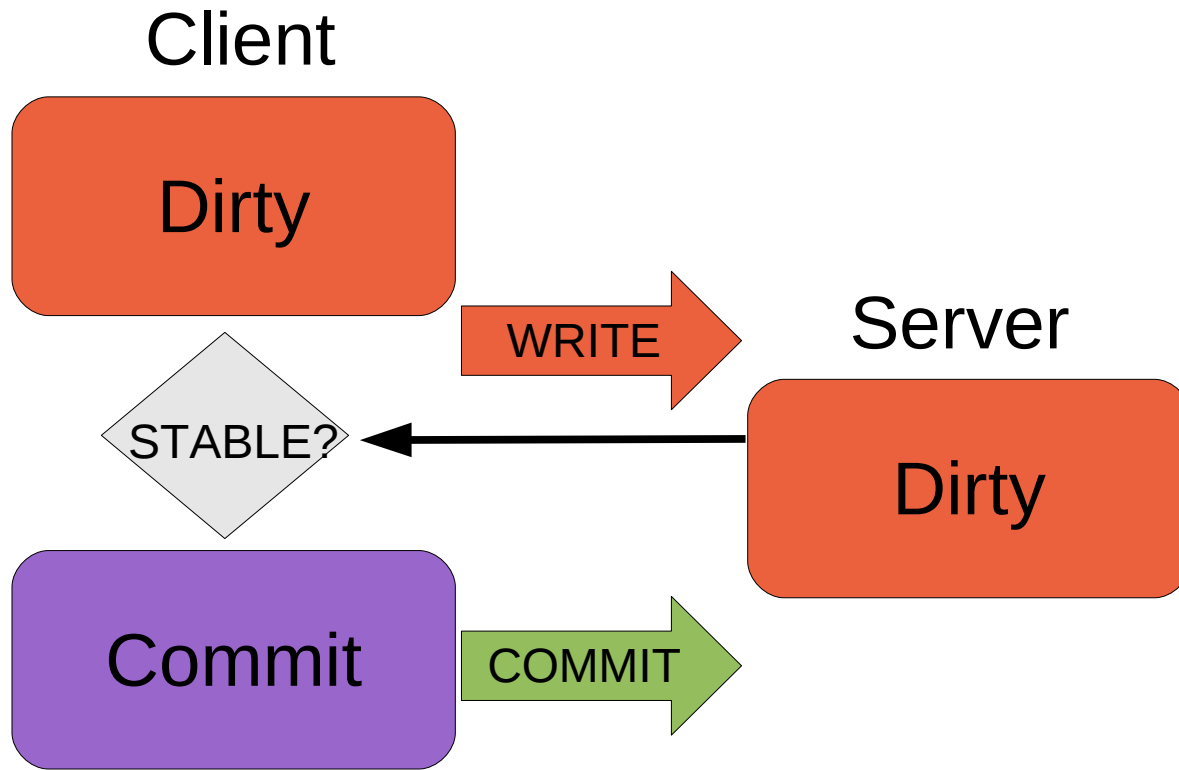
WRITE

Clean

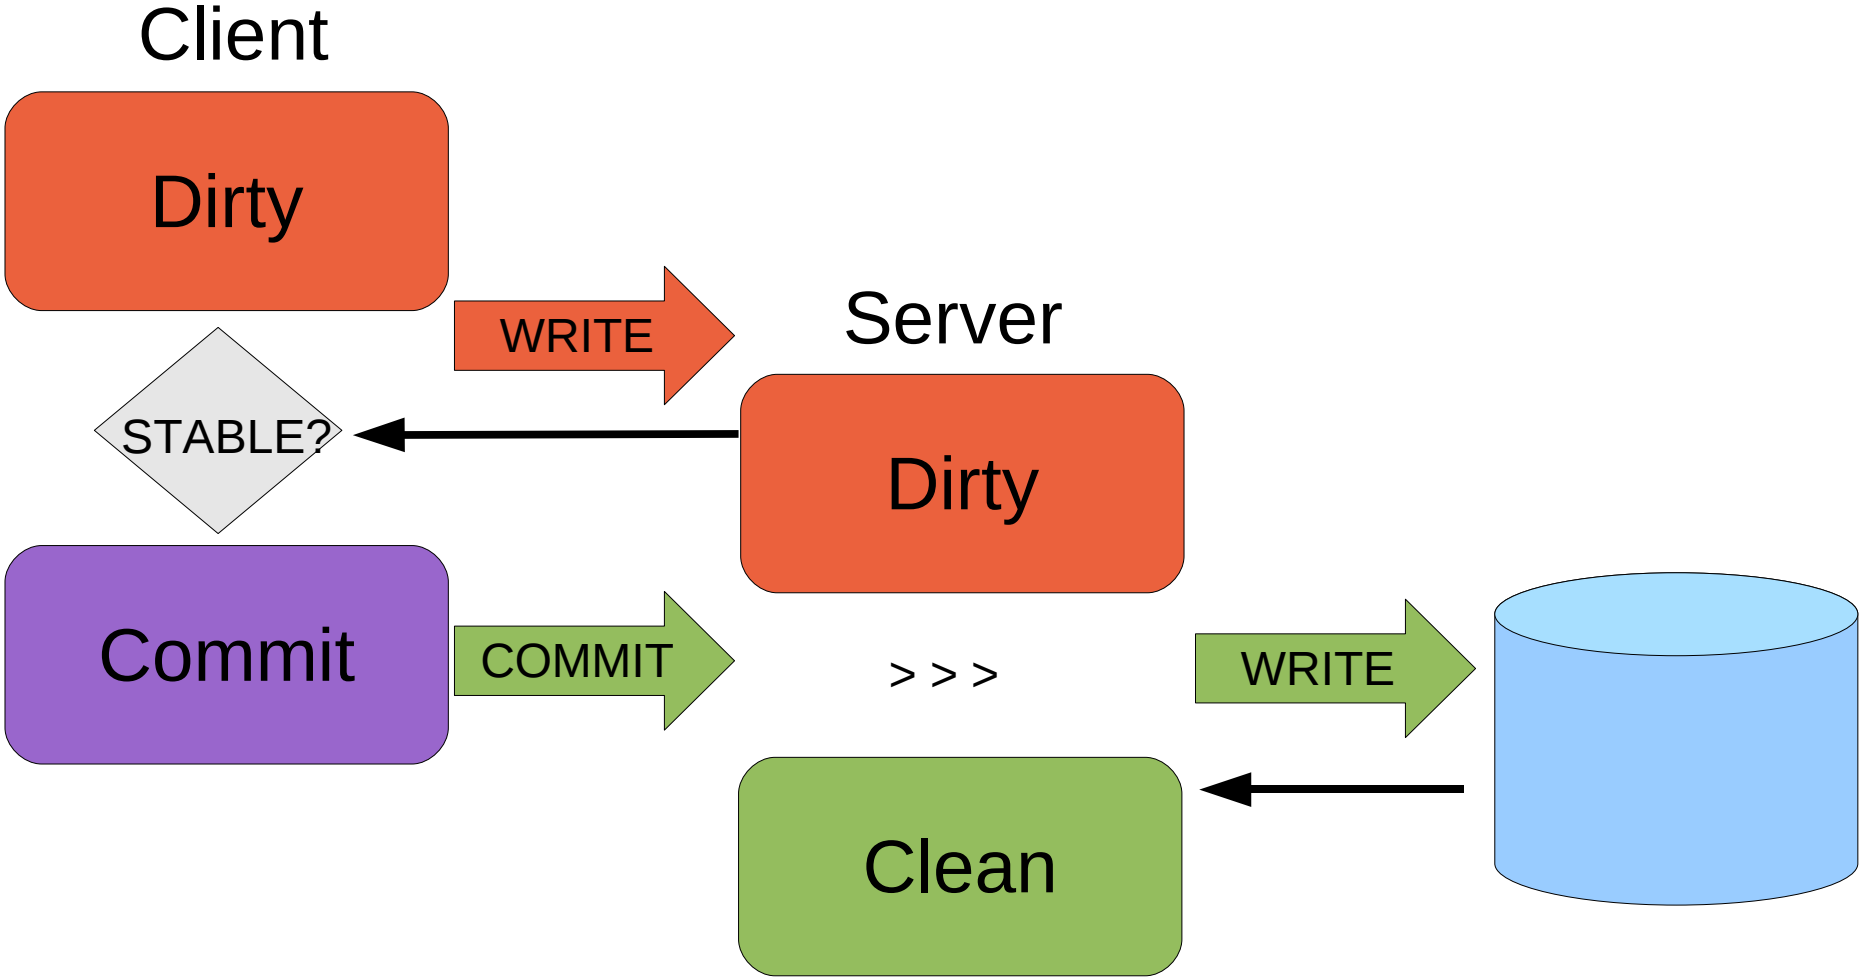
Clean



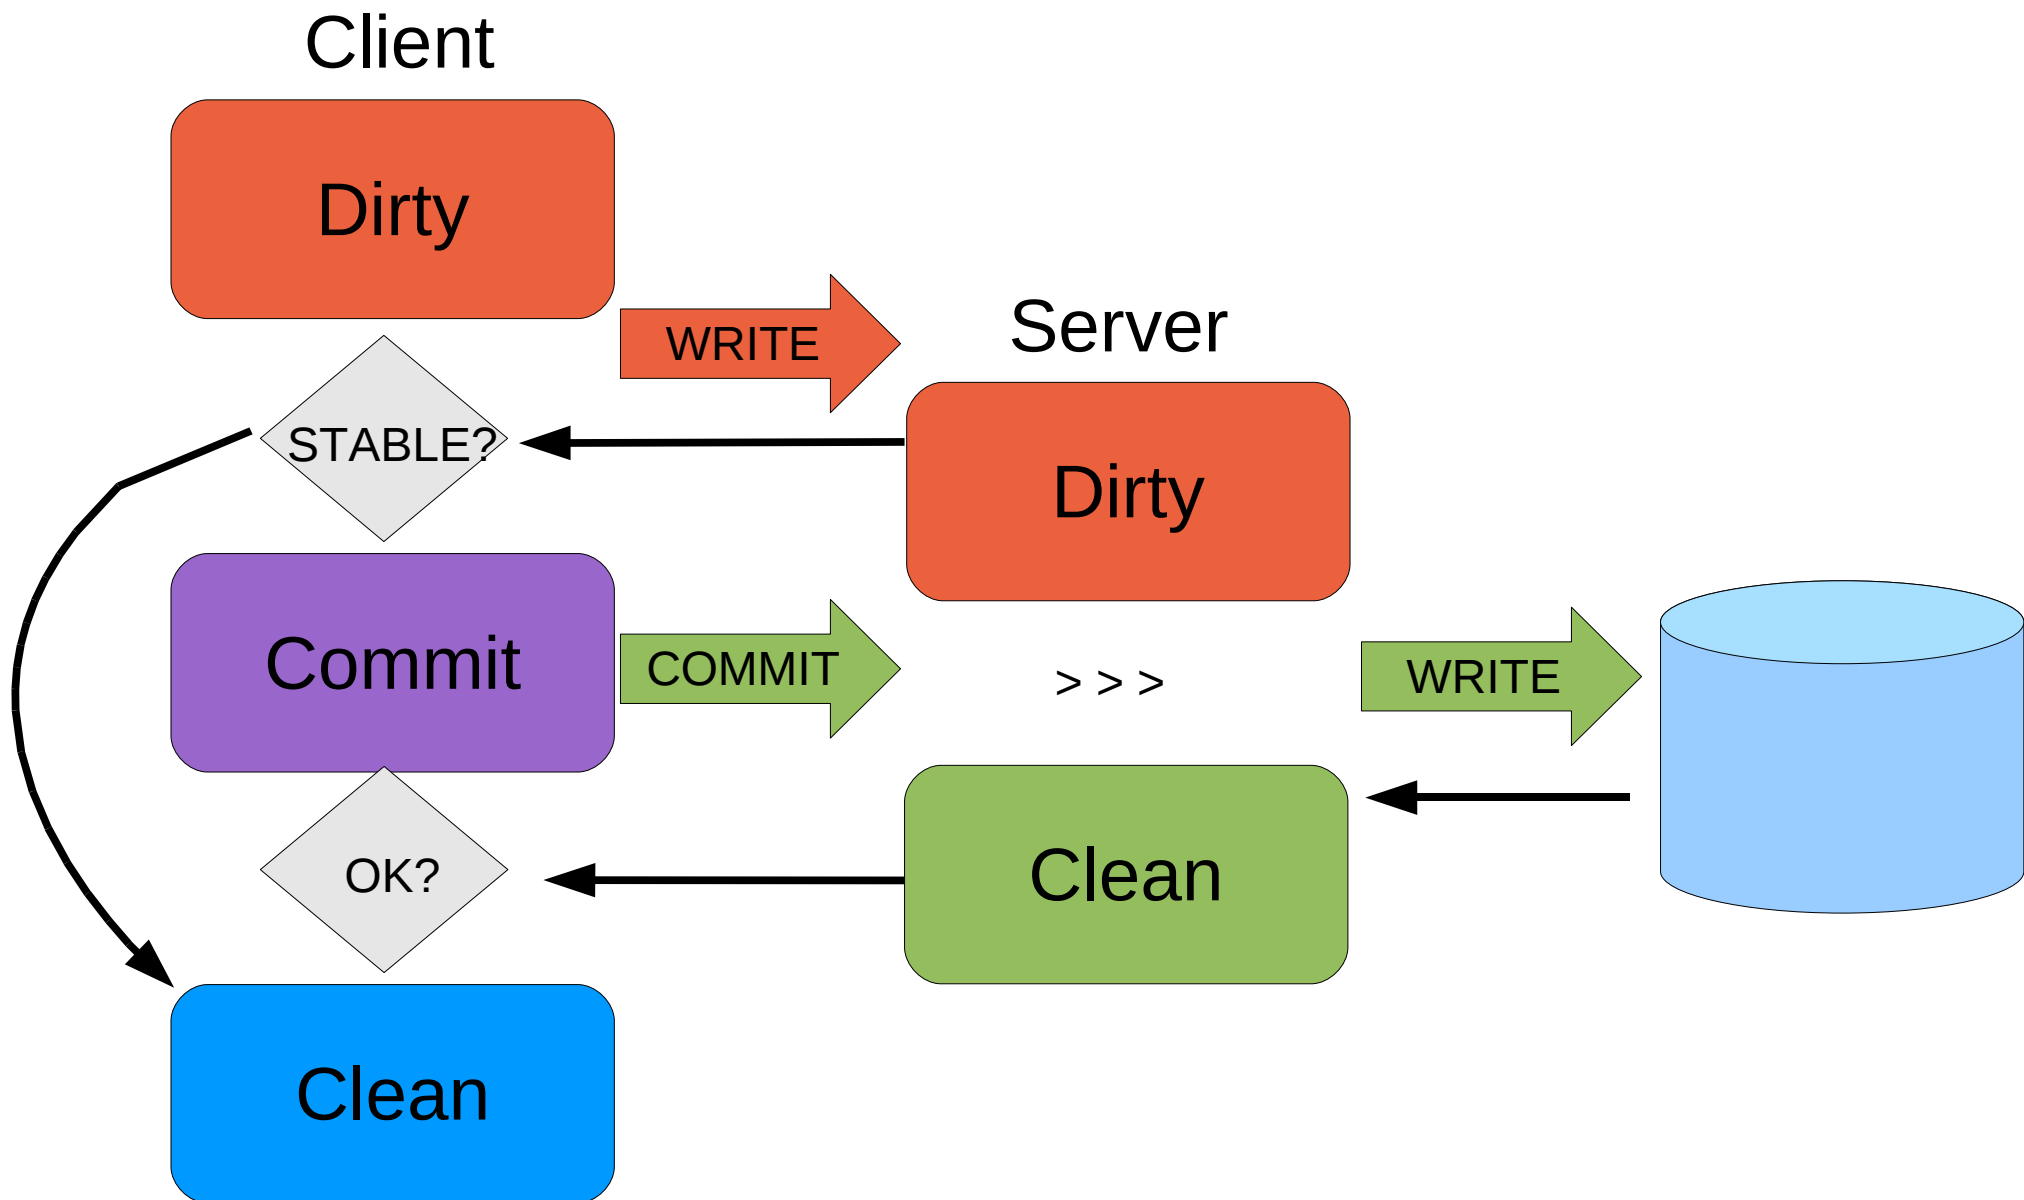
Good Old NFS



Good Old NFS




Good Old NFS




pNFS / NFS

Client



pNFS / NFS

Client



pNFS / NFS

Client

Dirty

Use pNFS?

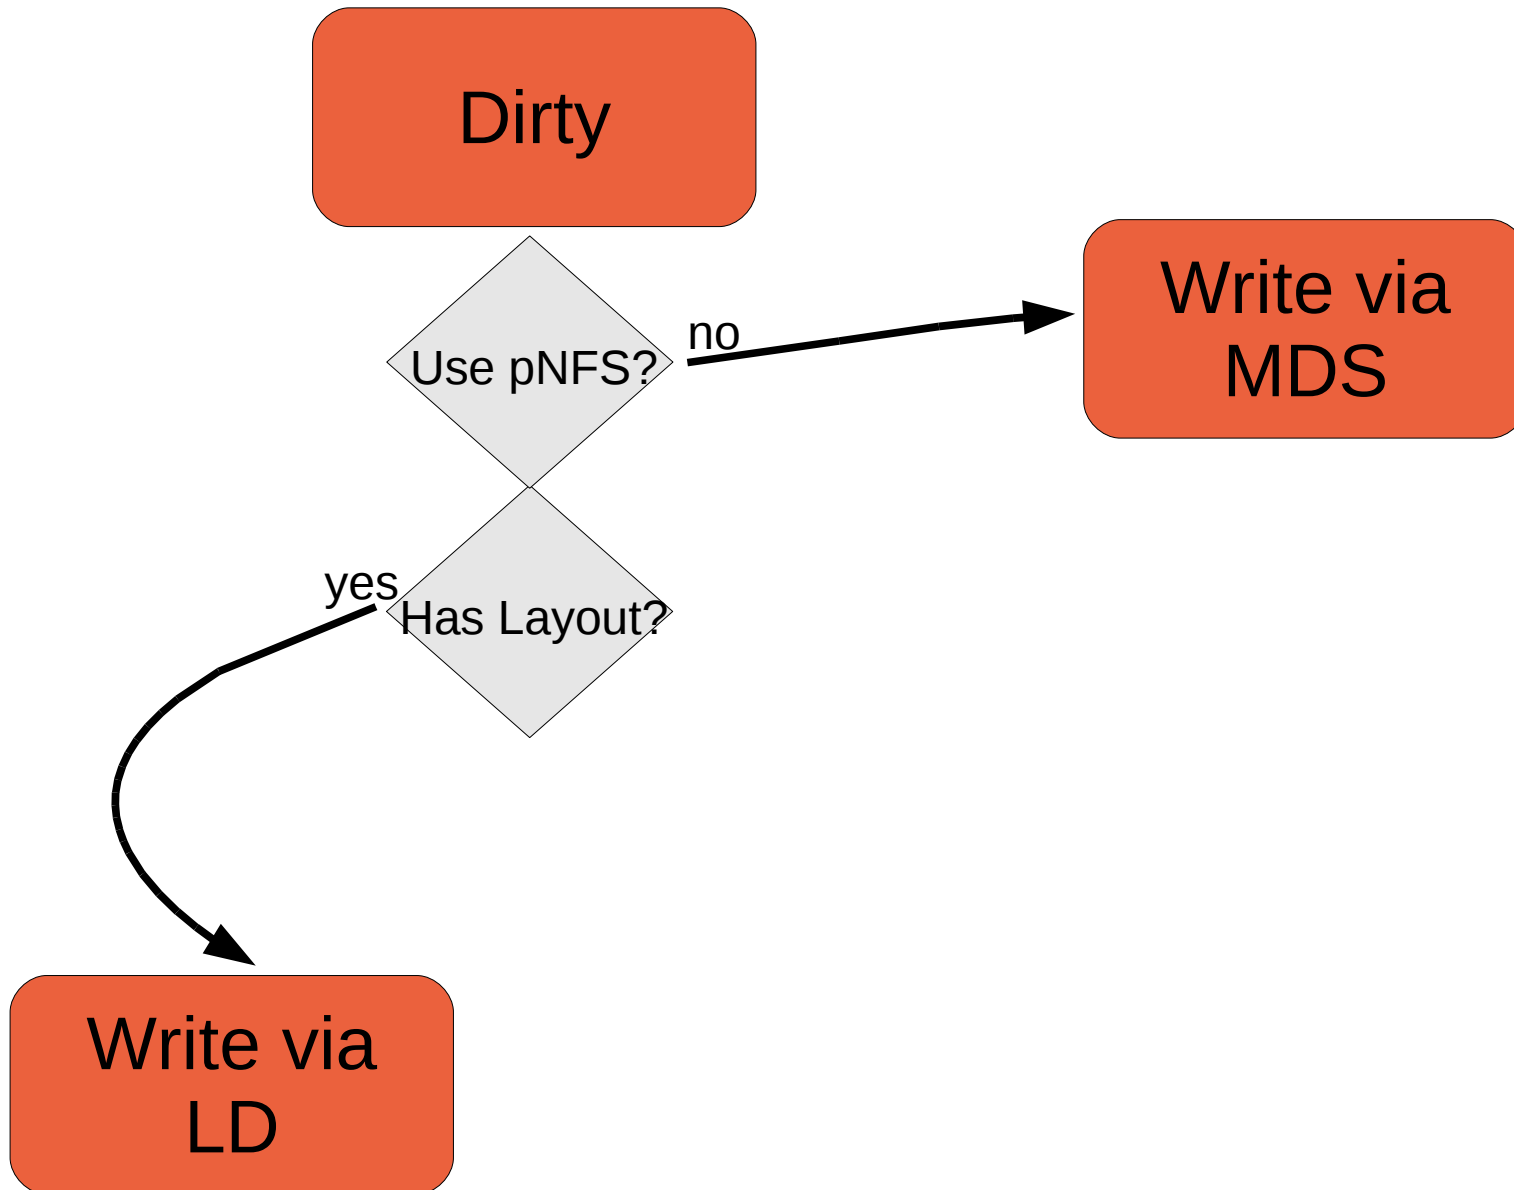
no

Write via
MDS

yes

Has Layout?

Write via
LD



pNFS / NFS

Client

Dirty

Use pNFS?

no

Write via
MDS

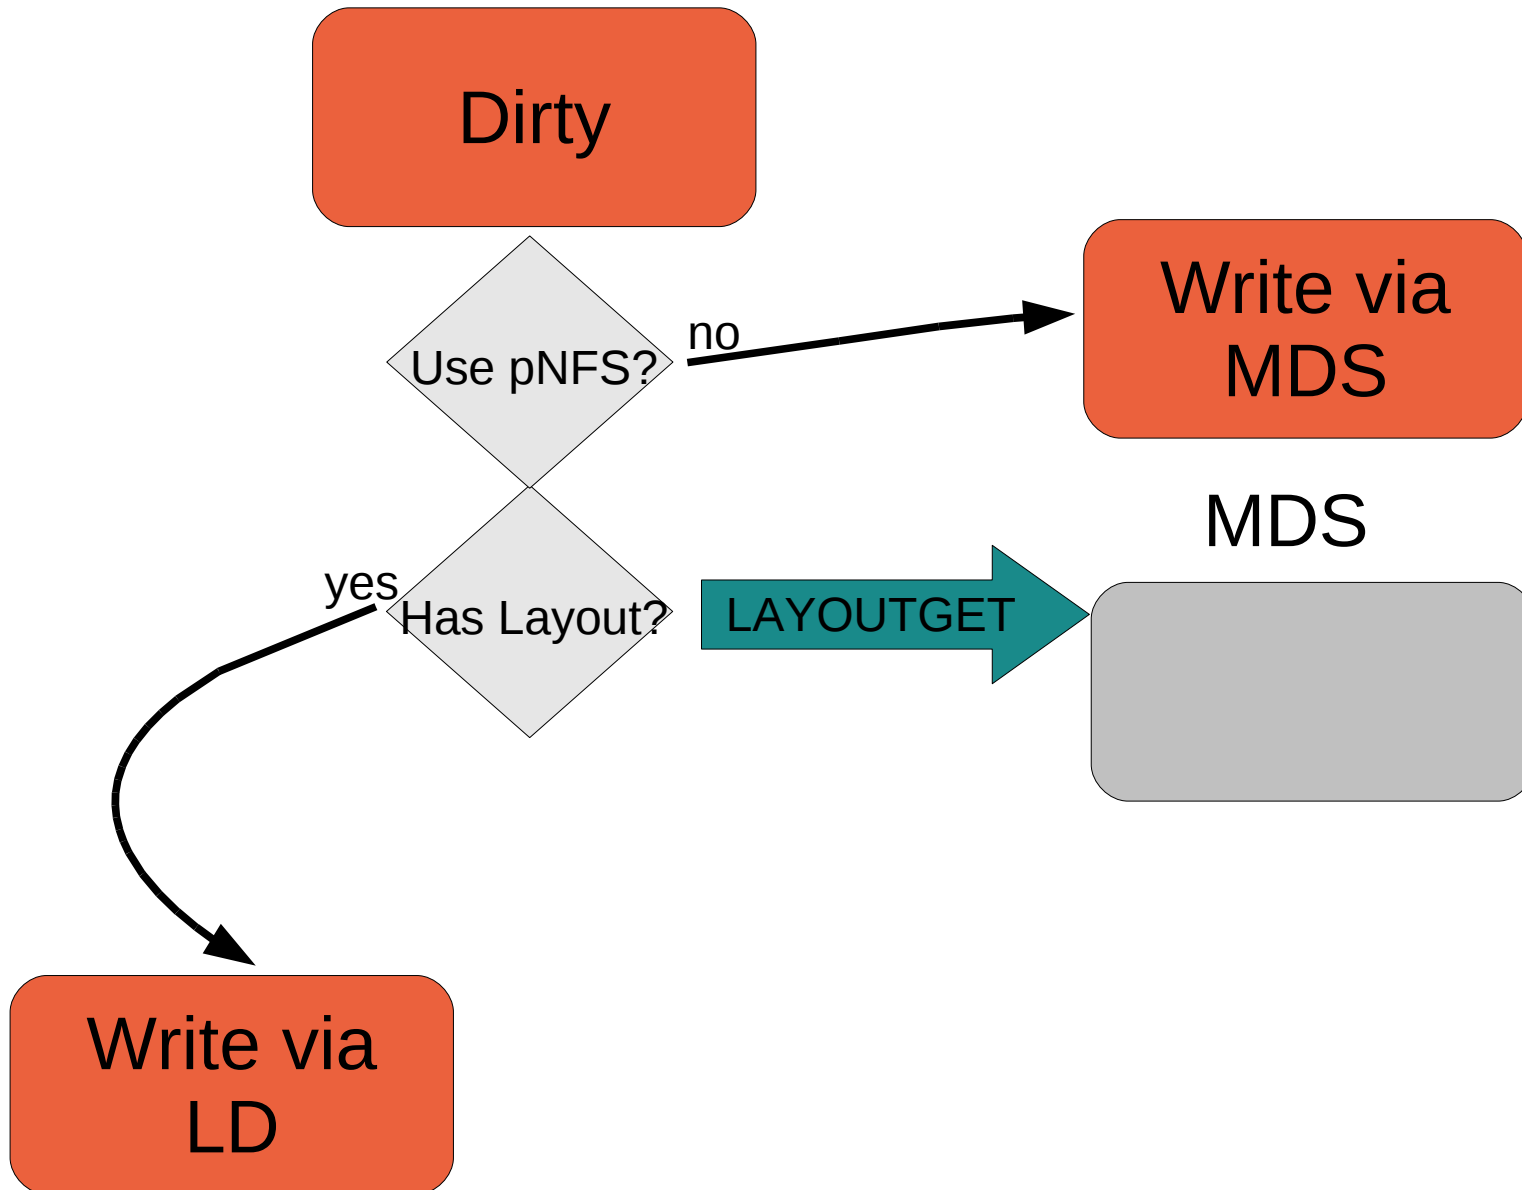
MDS

Has Layout?

LAYOUTGET

yes

Write via
LD



pNFS / NFS

Client

Dirty

Use pNFS?

no

Write via
MDS

Has Layout?

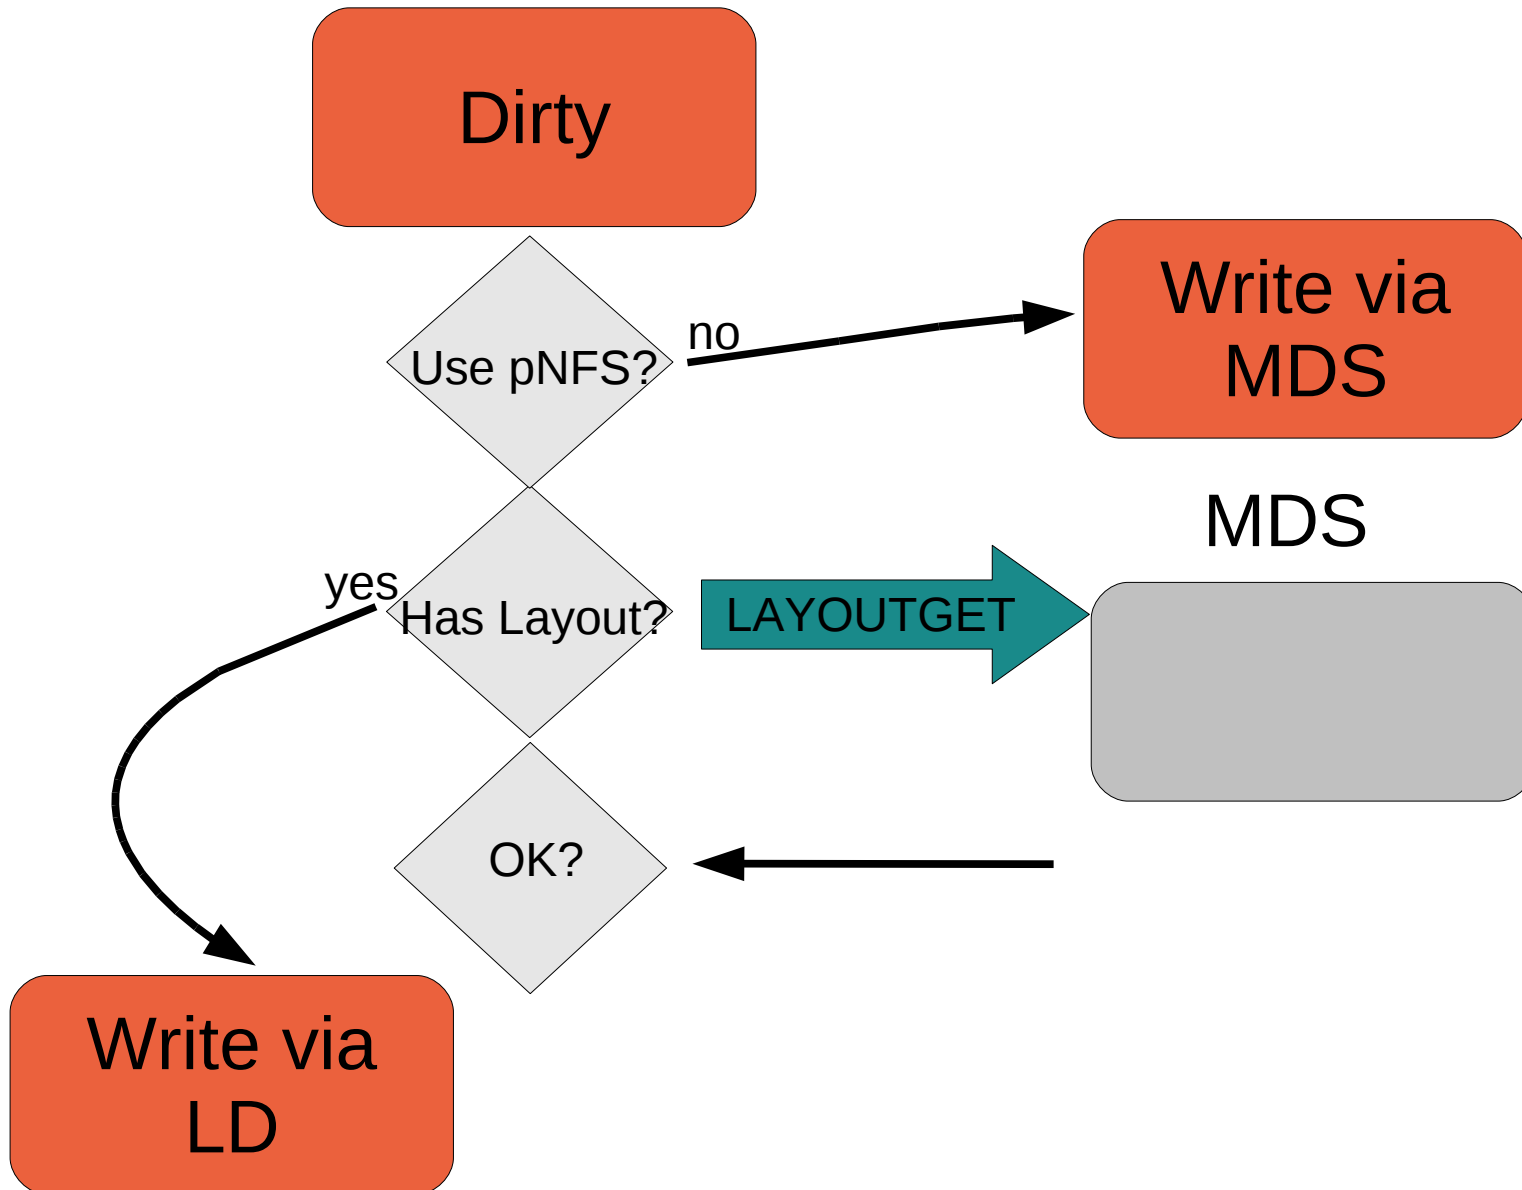
yes

LAYOUTGET

MDS

OK?

Write via
LD



pNFS / NFS

Client

Dirty

Use pNFS?

no

Write via
MDS

Has Layout?

yes

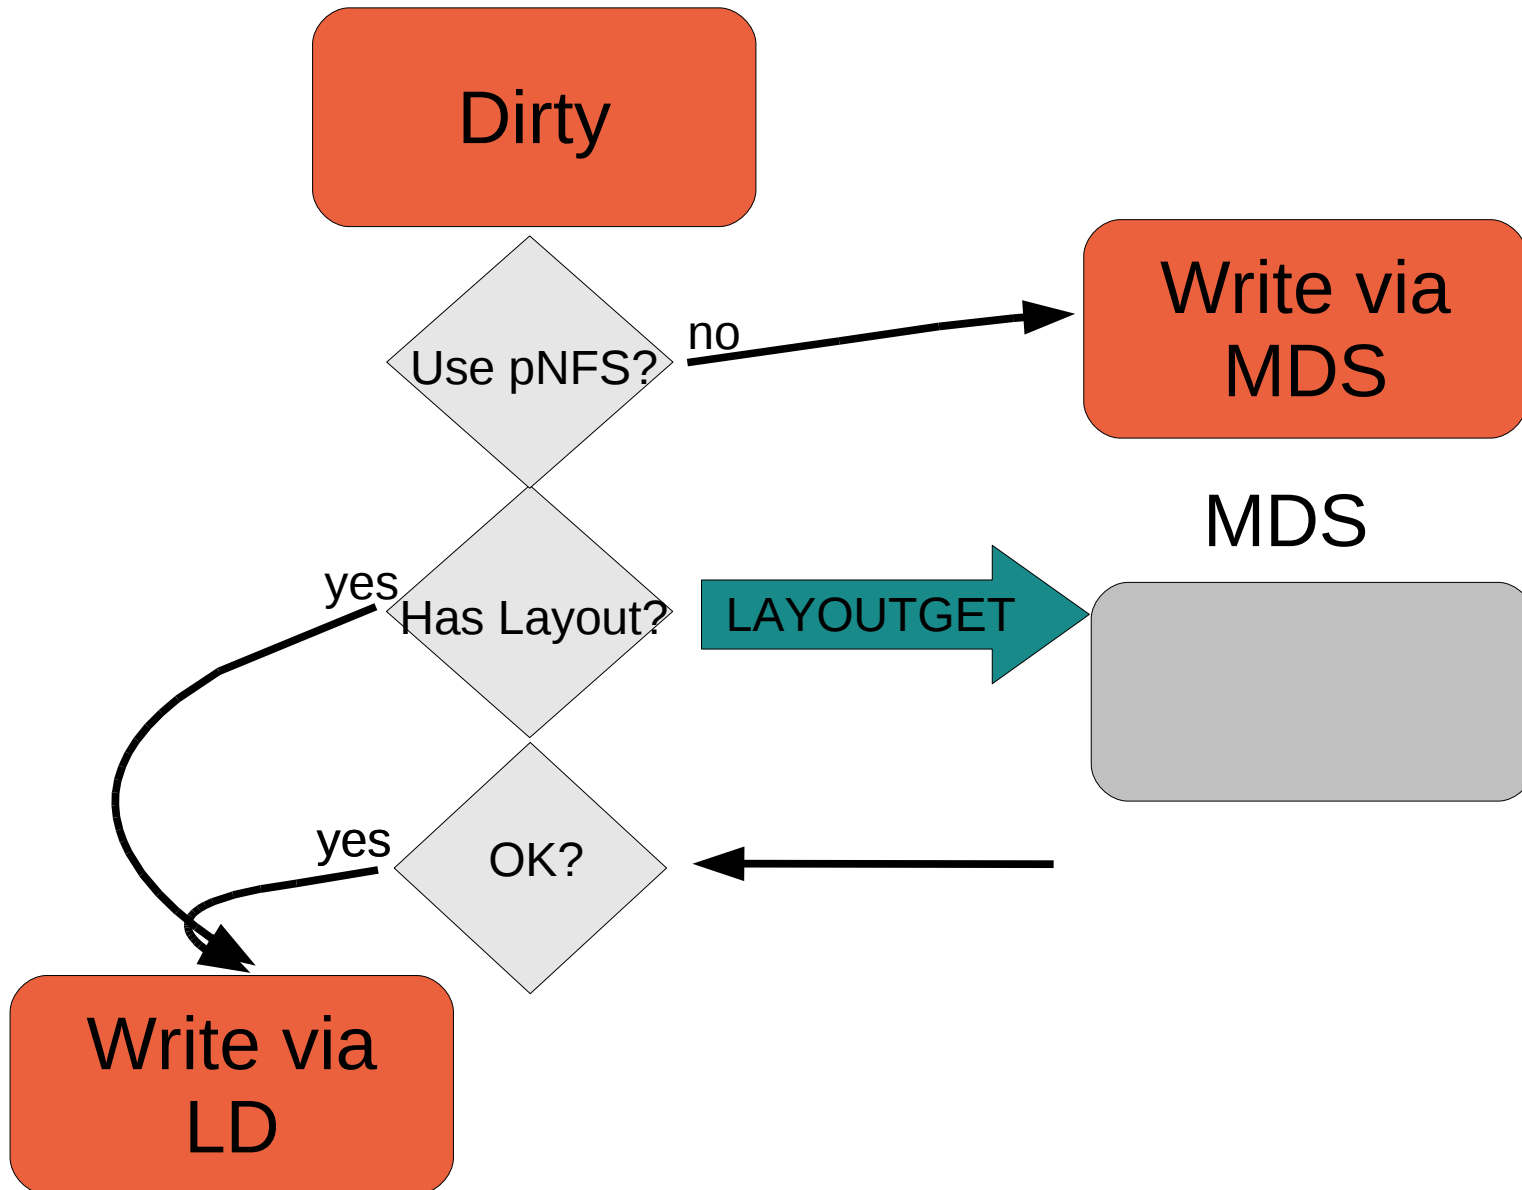
LAYOUTGET

MDS

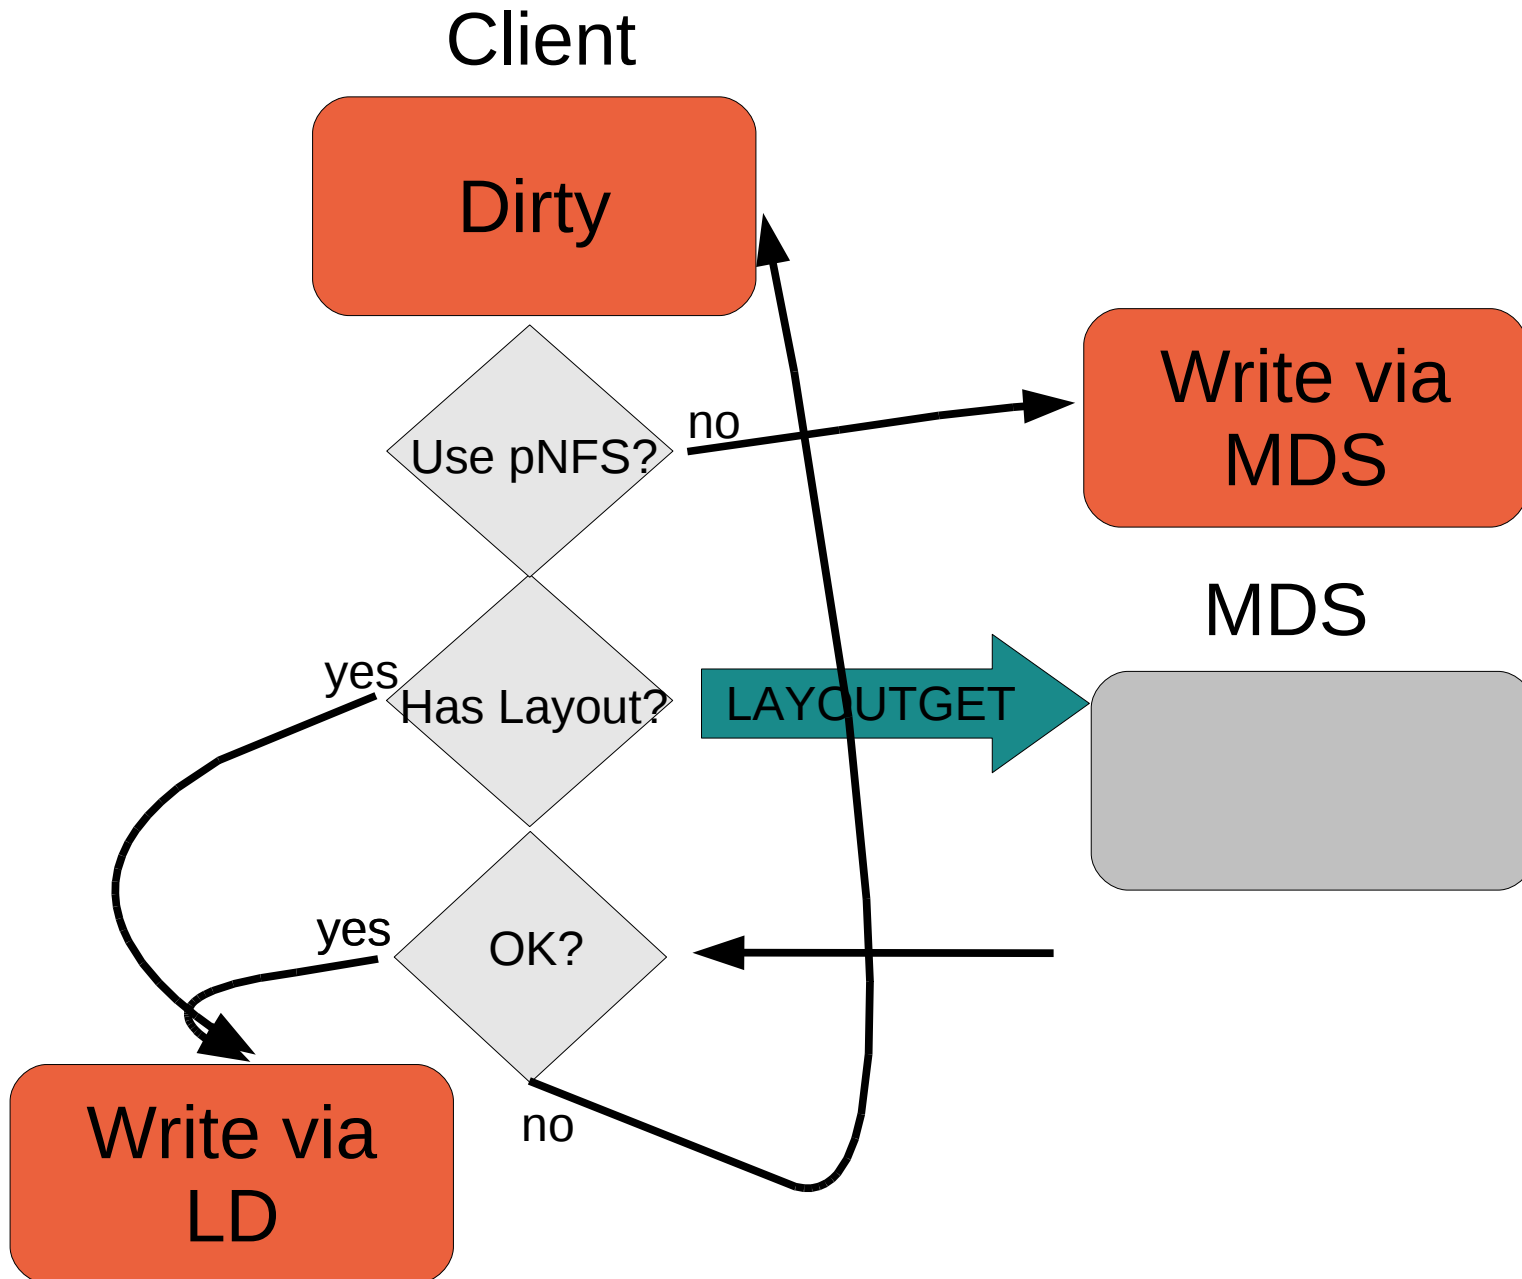
OK?

yes

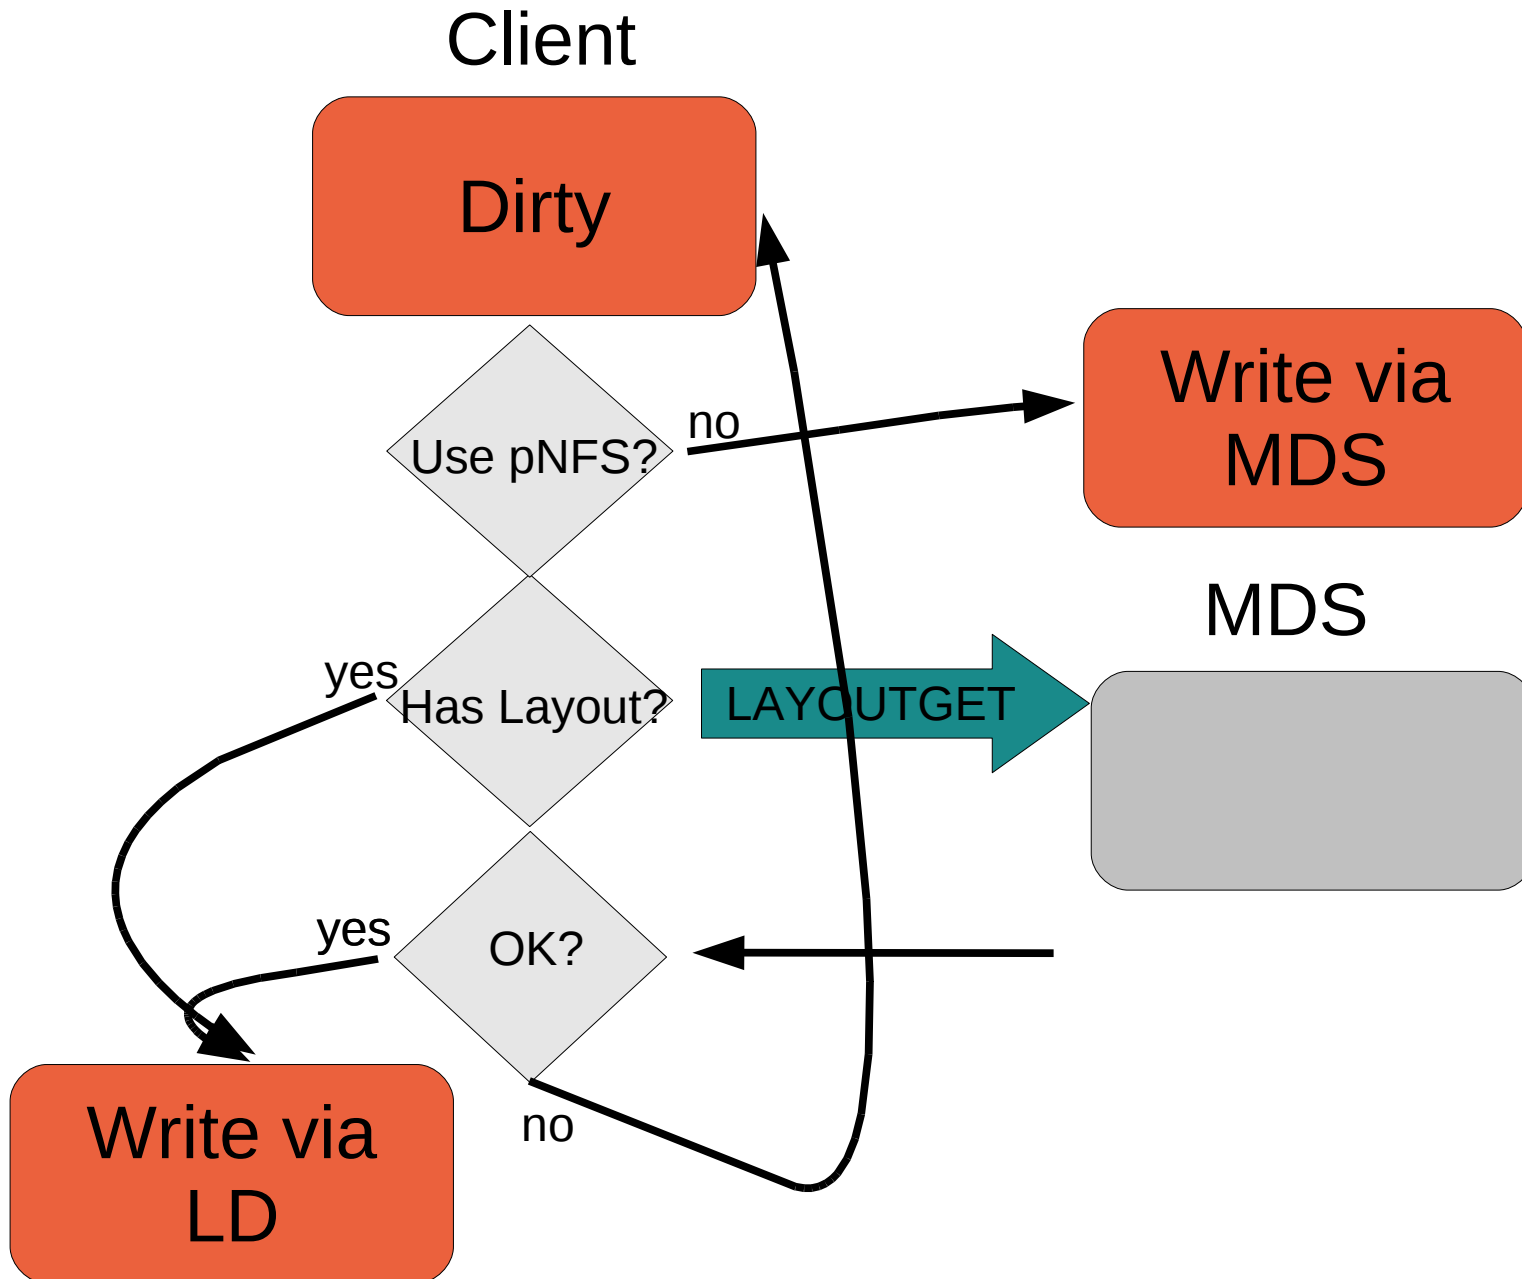
Write via
LD



pNFS / NFS



pNFS / NFS

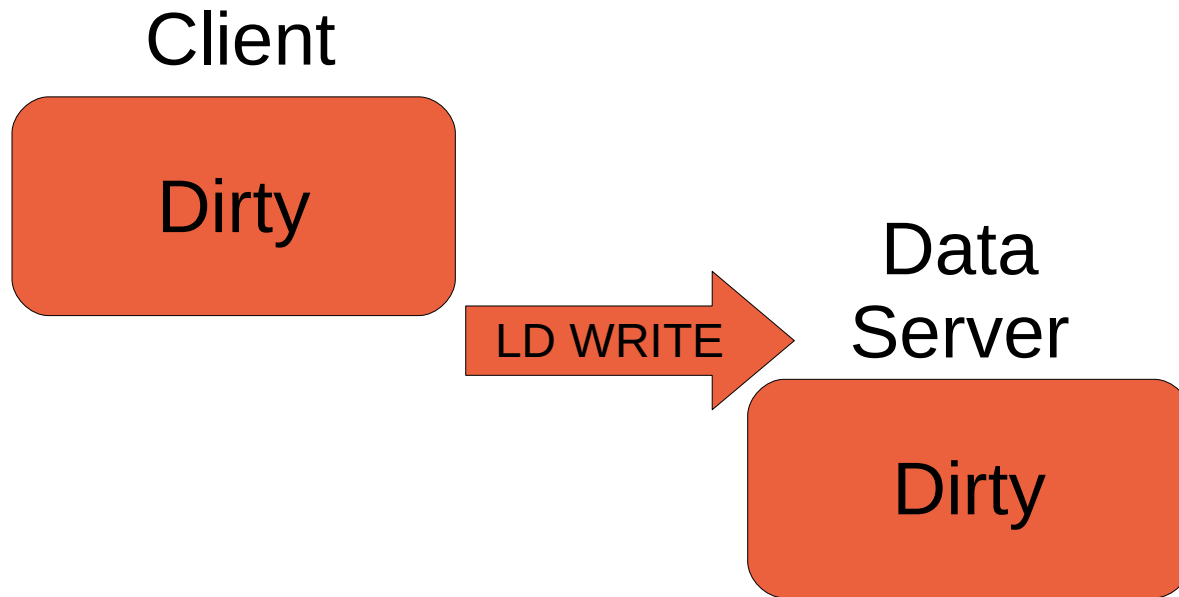


PNFS: Data Server

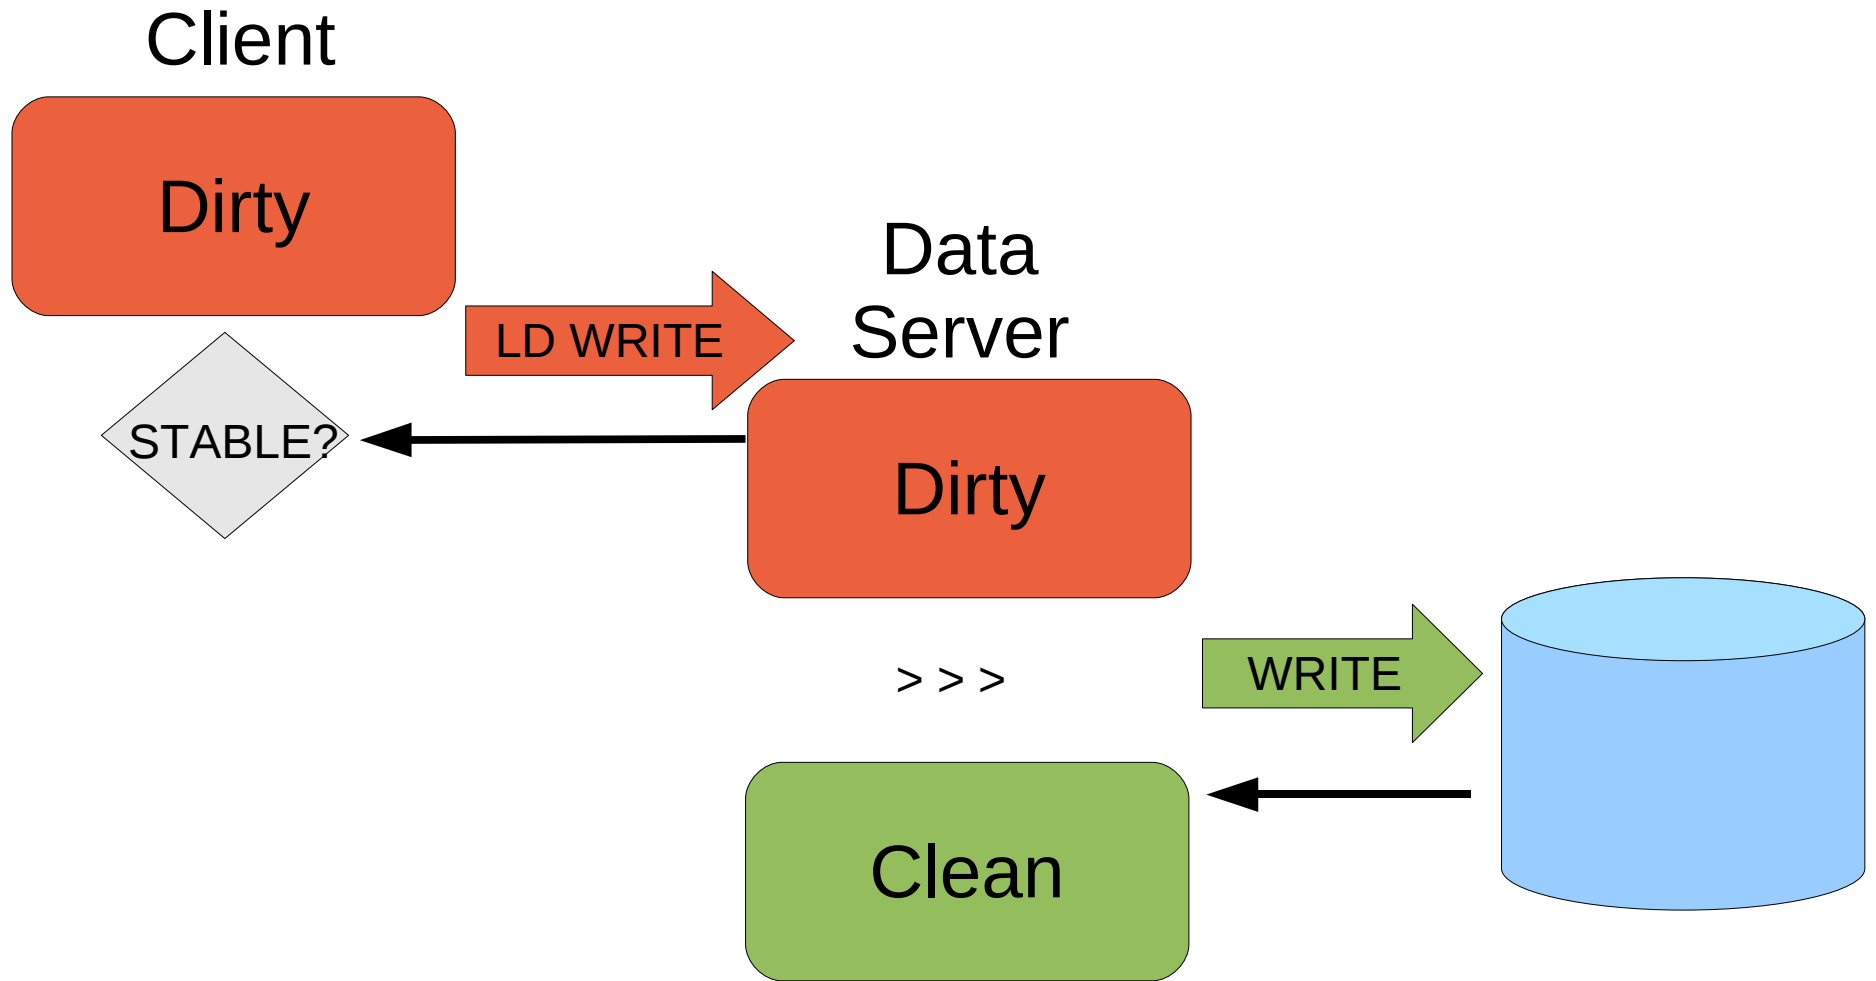
Client



PNFS: Data Server



PNFS: Data Server



PNFS: Data Server

Client

Dirty

LD WRITE

Data Server

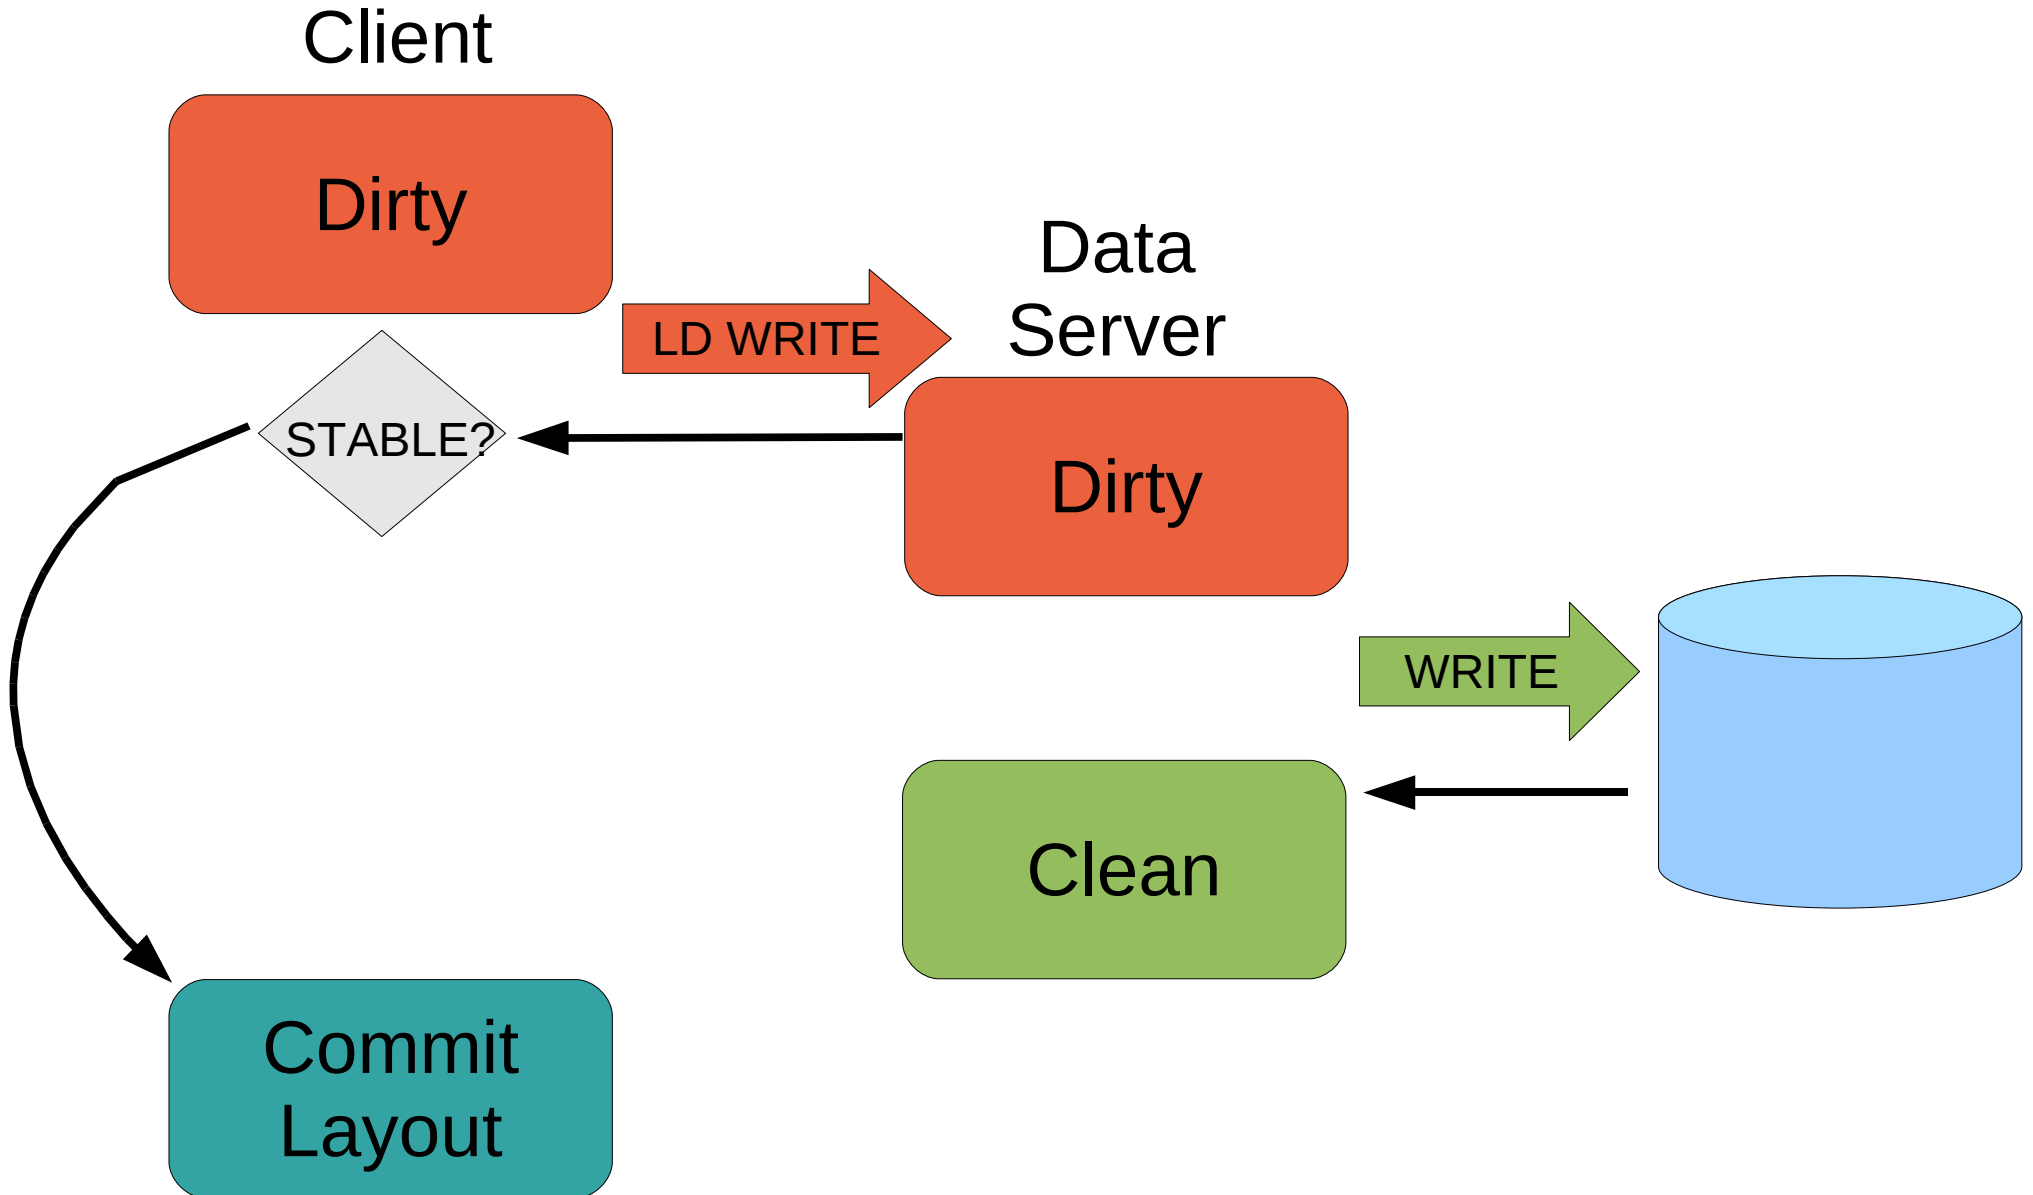
Dirty

STABLE?

WRITE

Clean

Commit Layout



PNFS: Data Server

Client

Dirty

LD WRITE

Data Server

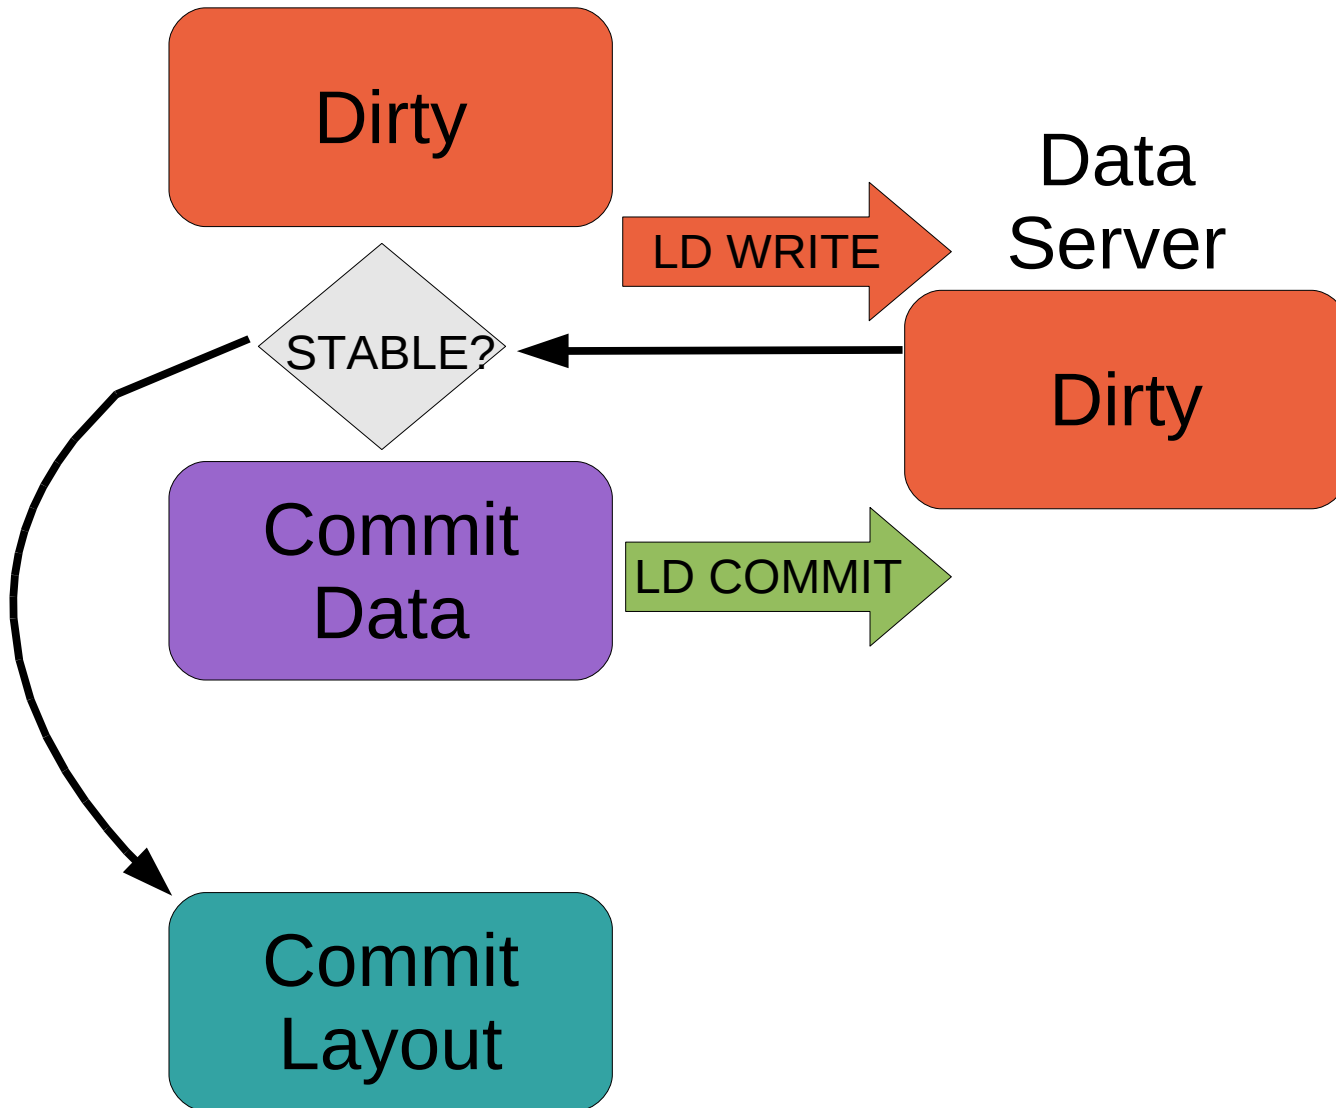
Dirty

STABLE?

Commit Data

LD COMMIT

Commit Layout



PNFS: Data Server

Client

Dirty

LD WRITE

Data Server

Dirty

STABLE?

Commit Data

LD COMMIT


>>>

WRITE

OK?

Clean

Commit Layout




pNFS: Layout Commit

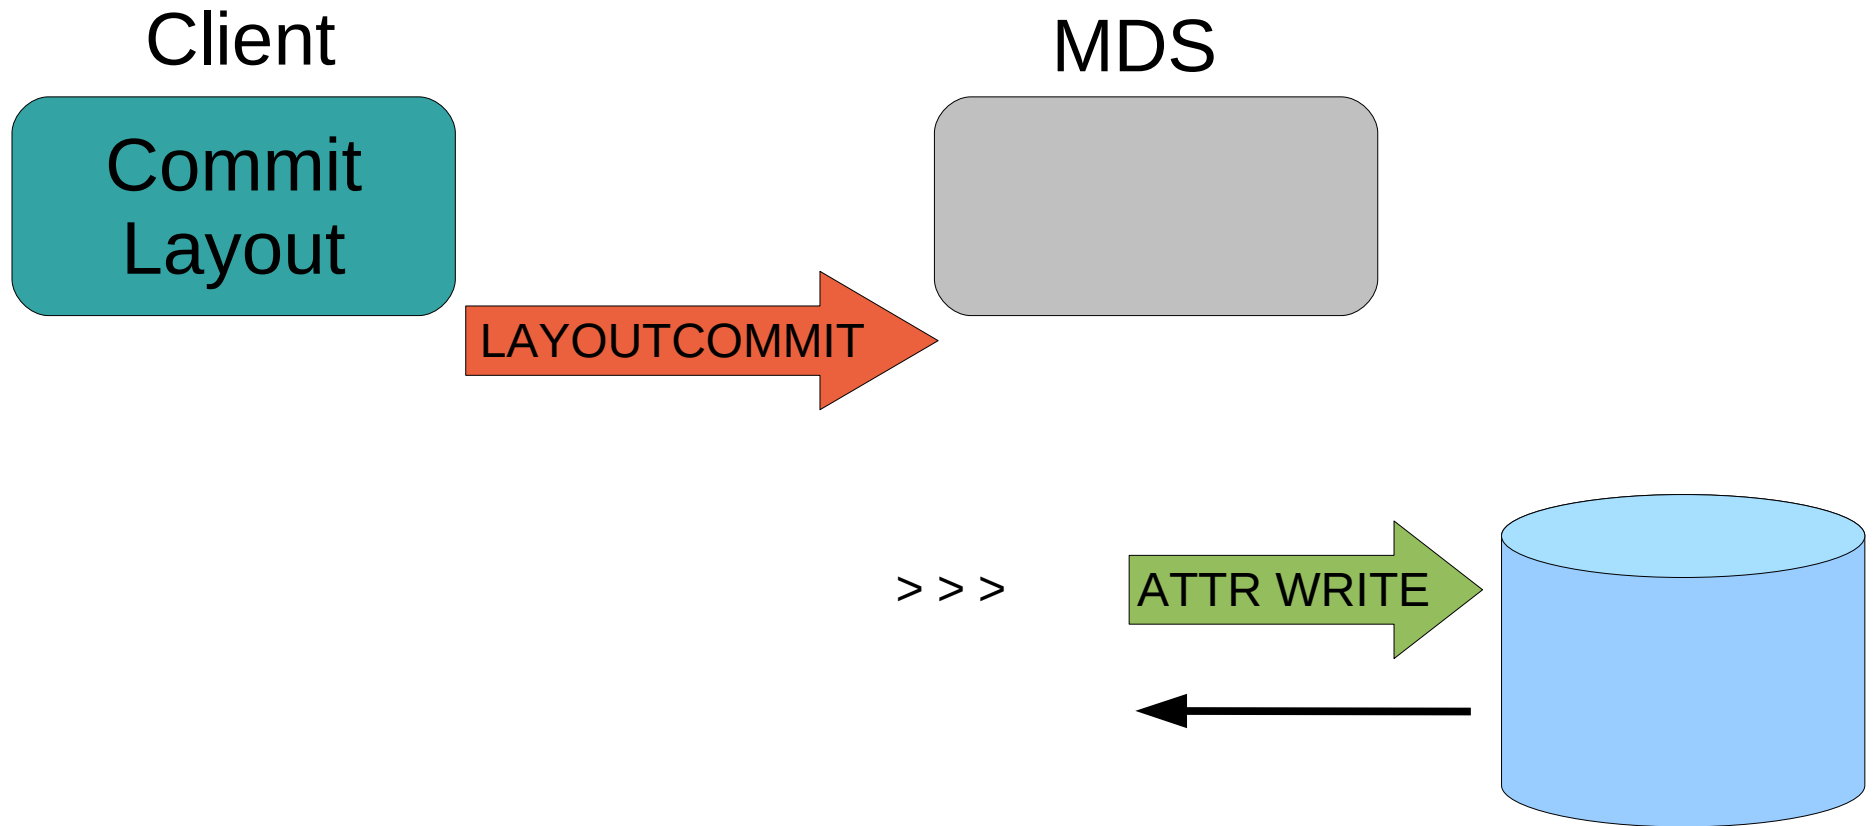
Client

Commit
Layout

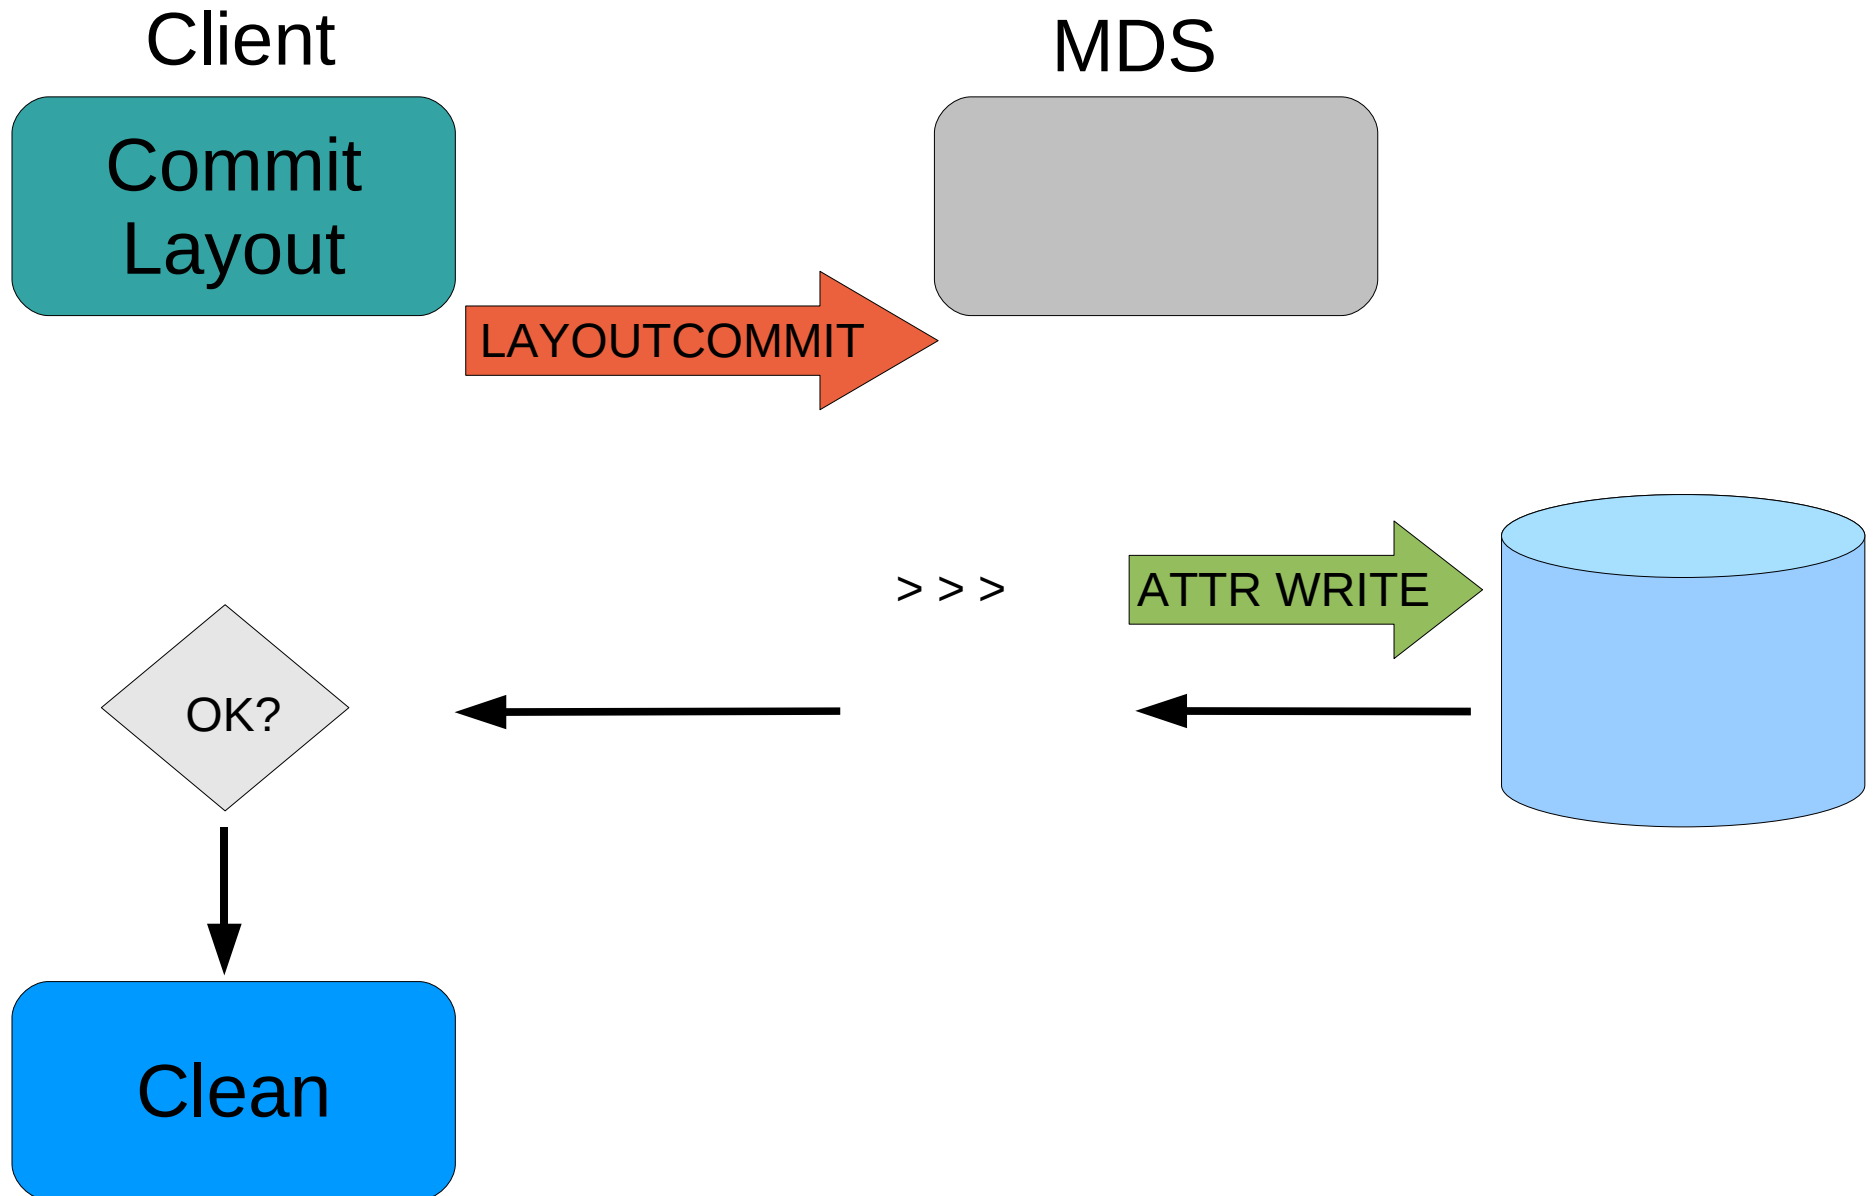
pNFS: Layout Commit



pNFS: Layout Commit



pNFS: Layout Commit



pNFS Page Sync State Model

