


Zumastor: Enterprise NAS for Linux

Daniel Phillips

phillips@google.com



or: It is high time Tux
arrived in the NAS space...

Zumastor Linux Storage Server

- Multiple live volume snapshots
- User accessible snapshots
- Remote volume replication
- Online volume backup
- Kerberized NFS and Samba
- CIFS previous versions API
- Easy administration interface
- NVRAM acceleration (work in progress)

ddsnap virtual block device

- ddsnap is the engine of zumastor
- Originally designed for cluster snapshots
- Small kernel driver coupled to biggish user space server
- rpc-like interface between kernel and user space
 - but not big and fat like rpc
- Copy-before-write snapshot strategy

ddsnap remote replication

- Which chunks are different between two snapshots?
- ddsnap server peeks into metadata
- Then read snapshot data to build a volume delta
- Get delta as a file or stream it
- Various kinds of delta compression

Zumastor Volume Monitor

- Hides the details of ddsnap, dmsetup, virtual device names, mountpoints
- Scheduled snapshot rotations
- Implements complex replication topology
- All driven by a filesystem based database
- Easy to use database editing interface

Snapshot write performance

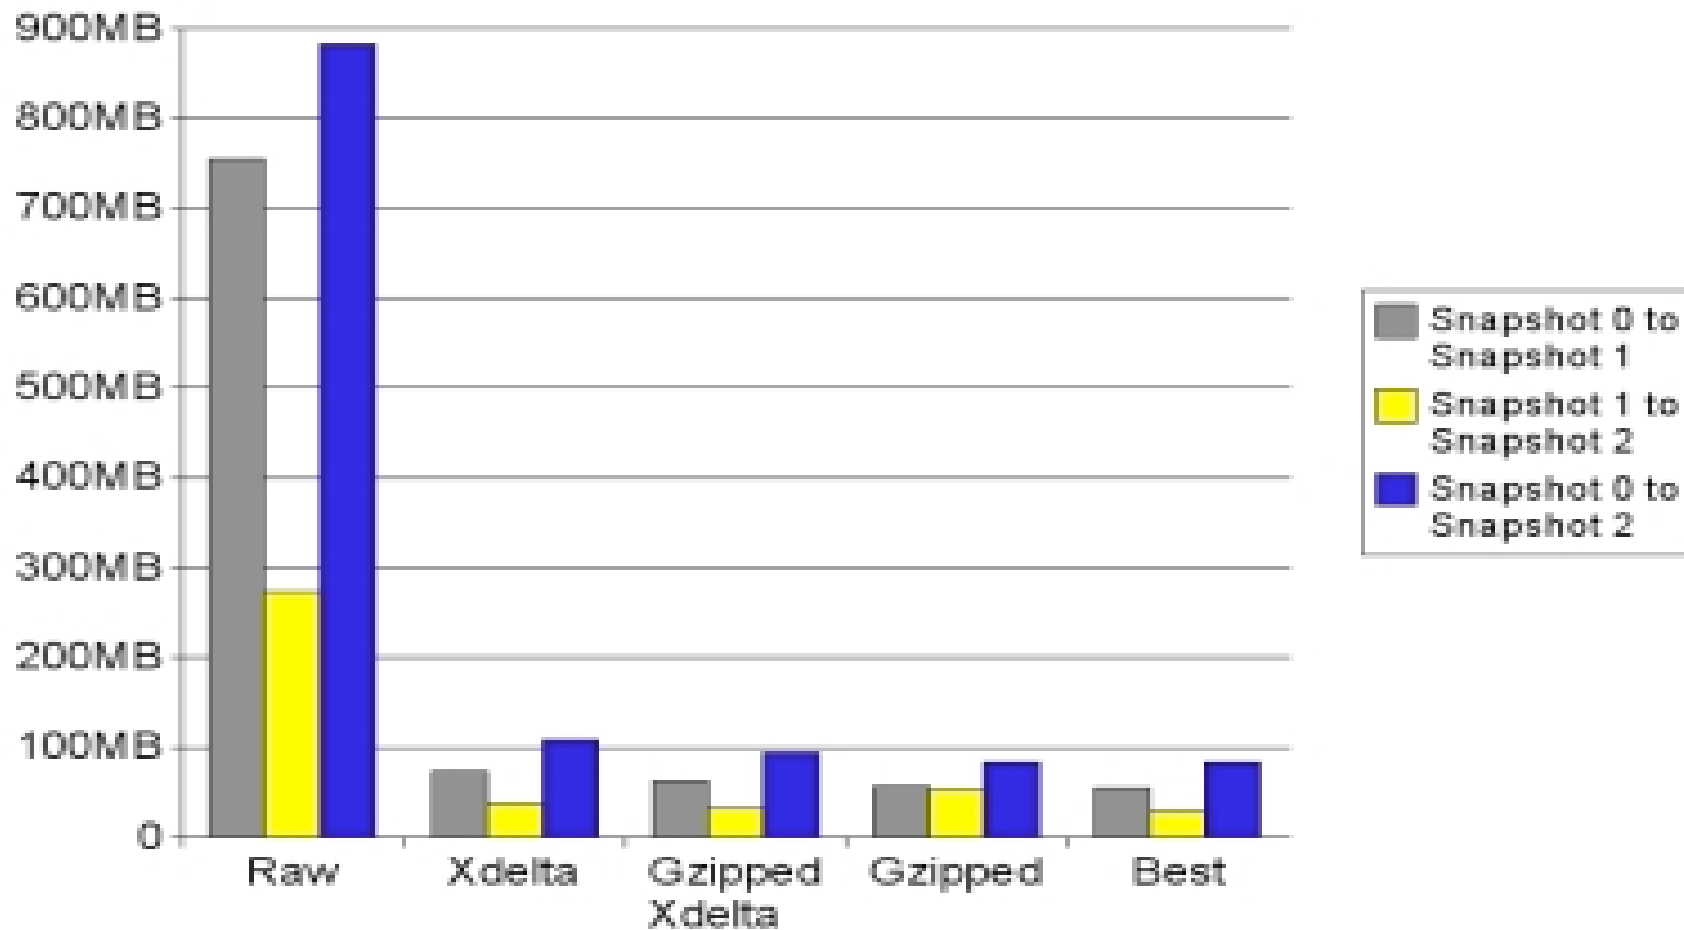
- Origin Read performance is great
- Write performance does not degrade with number of snapshots
- Write performance improves with larger chunk size
- Write performance improves a lot with metadata in NVRAM
- Write performance still pretty bad overall, partly covered up by NFS slowness

Delta compression performance

- Delta size equates to replication time
- Compression is a big payoff for slow links
- Extent oriented, need big chunks to work on but still need to stream
- zlib (gzip) for compression
- xdelta for binary differencing
- Compress or binary difference?
 - Try both and pick the best

Delta compression performance

Delta Size



Zumastor Futures

- It's going to get more features
- It's going to get faster and more robust
- It's going to get bigger

Zumastor Futures

- Give me a graphical front end
- Give it to me over the web
- Give me a real volume manager
- What about online resizing?
- Can I have incremental backup too?
- Faster, yah Faster!
- I don't want to buy NVRAM, can you make it so I don't need it?

Zumastor Linux Storage Server

Zumastor homepage:

<http://zumastor.org>

IRC channel:

[irc.oftc.net #zumastor](irc://irc.oftc.net/#zumastor)