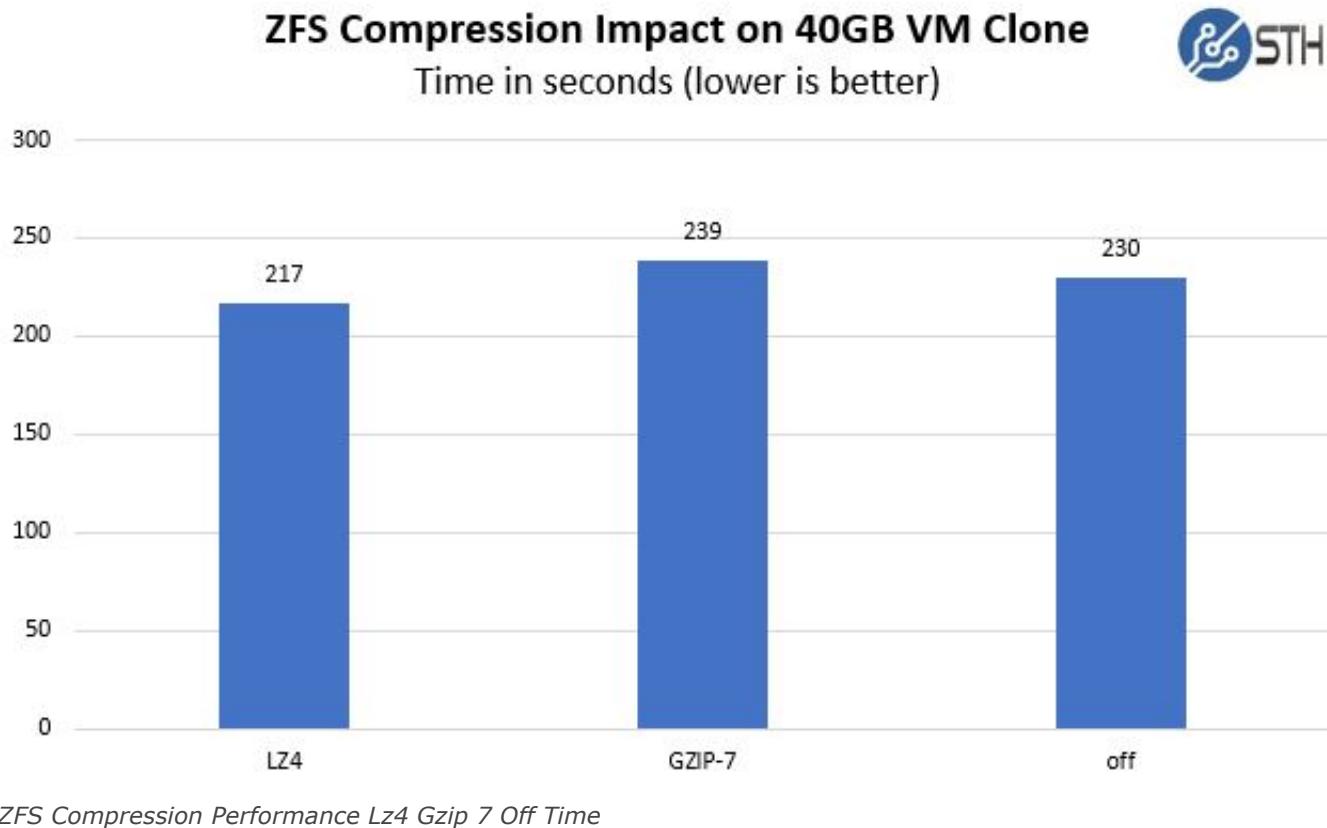


The Case For Using ZFS Compression

By **Patrick Kennedy** - January 2, 2018



An absolutely killer feature of ZFS is the ability to add compression with little hassle. As a resolution: use ZFS compression. Combined with sparse volumes (ZFS thin provisioning) and better disk space utilization. Many workloads work really well with ZFS compression. users overlook when it can be an enormous benefit.

How to find if you have ZFS compression enabled

For many ZFS environments, lz4 compression is the go-to solution. It is fast and gives a tradeoff to get substantial gains. While there are some ZFS environments that default to enabled by default. We looked back on questions we received in 2017, and a common one enabled.

There are many ways you can do this, but the easiest is with "zfs get compression." Here

ZFS Get Compressratio And Compression

Along with “zfs get compression” another useful command is “zfs get compressratio” which uses “compress” not “compression” ratio here. These are two attributes that you will see on zpools (p3600R1/vm-203-disk-1) compression inherited from the zpool (p3600R1.) We can change the compression on general purpose storage and then alter explicitly from there.

ZFS has a lot of attribute information that you can use “zfs get all” to lookup. Here is an example of the output:

ZFS Get All

If you have zfs compression showing as “on”, and want to see if you are using lz4 already
feature@lz4_compress which should be active if you are using lz4 as the default:

Zpool Get All

Either way, our resolution is to turn on zfs compression if at all possible.

How To Set ZFS Compression

We are going to suggest simply setting ZFS compression at the zpool level. That allows it easy to maintain. To set the compression to lz4, we can use “zfs set compression=lz4”.

ZFS Set Compression

In the first zfs get compression command we see that compression is off by default. We then use zpool (bulksata2) to turn compression on. We then verify that the compression is now set to lz4.

You will notice that the compression ratio is 1.00x which is essentially nothing. That is since we are compressing a new volume.

You can also use different algorithms such as gzip (e.g. gzip-7 that we are using in our new incremental backups on a dedicated backup server, going the gzip route can make a lot of difference in deduplication.

ZFS Compression Impact

We took a 40GB Ubuntu 16.04.3 LTS VM volume (about 32GB of 40GB in-use) used for “source storage”. For us, that means the source was on two Intel Optane 900p 280GB NVMe drives and two 960GB SSDs. These SSDs were configured as ZFS mirrors. Given the size of the VM, we were able to copy the VM to a compressed pool while ensuring the source was many times faster than the destination.

First off, we wanted to see compression ratios. We know that compression=off gives us 1.00x, so this is what we saw with the lz4 compressed pool:

ZFS Get Compressratio Results LZ4

We got a 1.93x compression ratio with lz4 compression. That is good for near transparent use of gzip-7 just to show the compression ratio difference:

ZFS Get Compressratio Results Gzip 7

As you can see, we got a 2.27x compression ratio which is significantly better. We used gzip-7 ratios versus lower levels offered by gzip that are faster.

There was a cost, however.

With lz4 compression, the entire snapshot clone operation (NVMe to SATA) took about 10 minutes which was already running at 52% utilization. Using gzip-7 we saw utilization spike over 60% point, this copy with compression=off pushed utilization up 8-9%. Here are the incremental operations:

ZFS Compression Performance Lz4 Gzip 7 Off Average CPU Utilization

In terms of the actual clone performance, the timings were close but there was a noticeable difference in the time taken to complete the operation.

ZFS Compression Performance Lz4 Gzip 7 Off Time

Not only did Lz4 use less CPU, but it did so over a shorter period of time.

We also were logging iowait while we were doing these operations. Since we were using a system more akin to a running virtualization server, we wanted to see what the impact was on io be minimized.

ZFS Compression Performance Lz4 Gzip 7 Off Max Iowait

If you remember, we are doing this transfer from mirrored Intel Optane SSDs to mirrored you were cloning development VMs to lower-cost storage. We were debugging the impact is why we were watching the iowait number. Through the ten runs we did, each time we | was near 0% on the system otherwise.

To some compression=off may seem like the obvious choice for the highest performance, better compression, lz4 provides “good enough” compression ratios at relatively lower pe recommendation.

Final Words

If you are starting the year and looking for a project, ensure that your ZFS storage is usi compromise between compression ratio and performance is well-known at this point. Usi saving benefit during some operations which makes it a great choice.

Even though lz4 ZFS compression is a well-known solution, over the 2017 holiday seaso from other folks who were not using it on their zpools. Set compression=lz4 at the zpool compression. You will be happy for this new year’s resolution that takes a few seconds ar

Test Configuration Notes

Here is a quick overview of the test configuration for the above. We are becoming ZFS or

- System: Supermicro 2U Ultra
- CPUs: 2x Intel Xeon E5-2698 V4
- RAM: 256GB (8x 32GB) DDR4-2400 RDIMMS
- OS SSDs: ZFS Mirror Intel DC S3610 480GB
- Source SSD: ZFS Mirror Intel Optane 900p 280GB
- Destination SSDs: ZFS Mirror Samsung PM963 960GB
- Proxmox VE 5.1
 - Based on Debian Stretch 9.2
 - Kernel 4.13.3
 - QEMU 2.9.1
 - ZFS 0.7.2

Patrick Kennedy

<http://www.servethehome.com>

Patrick has been running STH since 2009 and covers a wide variety of SME, SMB, and SO industry and has worked with numerous large hardware and storage vendors in the Silicon information about server, storage and networking, building blocks. If you have any helpf

