

# Troubleshooting Faulted Pools after High Temperature Condition

## Problem

After datacenter or computer room where Brickstor Storage System is physically located experiences a climate control system failure, resulting in rise in ambient temperature much above the typical 60 to 70 degree F., problems are observed with storage, in some cases these problems translate to outage of entire pool or pools as a result of too many drives being seen as failed. This problem is really caused by a protection mechanism in the disk drives used in Brickstor Storage Systems.

Disk drives have internal sensors which let them know what the ambient temperature is, whether the device is overheating, etc. When these sensors detect temperature above some threshold, which will vary to some degree with various makes/models, the Operating System is made aware of this through a health and fault status framework. There is a continuous feedback loop which helps to protect the system from various conditions, and allow the system to recover from some things automatically, where possible.

The failure mechanism is quite straight-forward. All disks that are actively used by the system, in other words, any disk in any ONLINE and imported pool, are continuously reporting back to the system any faults that they encounter. One by one, as drives reach the temperature threshold, they begin to report to the system an over-temperature condition. Because of natural sensor precision variance and location in shelf and relative location of shelf in the rack (i.e. warmer air raises, so least elevated shelves remain cooler than most elevated shelves), temperature alerts from drives do not all come at same time. Some drives will reach the threshold sooner than other drives will. Once this alert is reported some number of times, a count is kept internally; the system assumes that the state of the device is critical, and marks it as faulted. This means the system will no longer use this drive. As soon as this happens to first drive, the pool becomes degraded. As this over-temperature warning is reported by more and more drives, more and more of them are faulted in the same manner, and they are effectively no longer usable until administrative intervention occurs.

## Solution

This problem may be resolved in a number of ways, but this is the most practical approach based on our experience.

1. If there are systems trying to access the SAN, failing to do so due to the service outage of pool(s), shut them all down. We need to reboot the SAN, but before we do, we should make sure that anything that the SAN depends on for name resolution and domain membership and such is online and working, if possible.
2. Power off the head unit (controller node). If the OS is not responding to shutdown requests unplug the unit.
3. Power down all the shelves as well, and with the head still powered off, turn on shelves one at a time. Once shelves are online and had opportunity to POST, which should not take more than about 20 to 30 seconds from flip of power switch to being ready, we are to power up the head.
4. Power the head up and wait for it to boot. Depending upon the state of the pool(s) we may or may not have some pools online.
5. If any pools appear to be imported, but faulted, as reported by **zpool list**, which is an unlikely event, power off the head unit, unplug SAS cables from storage, reboot the head and then plug the storage back in.

```
# zpool list
```

6. At this point we need to mark faults as repaired. It is a good idea to look at all the faults, because there may be some reports that address other parts of the system. There is always a chance of hardware failure during climate events like high temperature. Use **fmadm faulty | less** to get a list of faults, in a pageable format and review the faults carefully. We should see a number of faults having to do with disk drives going offline. There are likely to be ZFS related faults as well, which are a direct results of drives becoming unavailable.

```
# fmadm faulty | less
```

7. Once we have a good feel for the things reported as faulted, we should mark them repaired, which will allow for disks to become visible again. This loop should do what we want in one shot.

```
# for fault in `fmadm faulty -s | tail +4 | awk '{print $4}'`; do fmadm repair $fault; done
```

8. Check the state of pool without importing it with **zpool import**. Without any arguments, **zpool import** will only show what pools are available to import. We should now see all pools that the system had online before, and they should hopefully report healthy.

```
# zpool import
```

9. If pools appear to be OK, import them one by one with the following command, replacing poolname with actual name of pool. **zpool import -o cachefile=/etc/.zpools/spa\_info poolname**

```
# zpool import -o cachefile=/etc/.zpools/spa_info poolname
```

10. Reset pool import failure count to 0 in order to allow for pool(s) to be imported automatically next time the system boots. For each pool on system run the following command: **zfs set racktop:poolimport:<pname>:last\_import\_good=0 bp**, replacing **<pname>** with actual pool name, for example assuming we have pool named p01, the command is then **zfs set racktop:poolimport:p01:last\_import\_good=0 bp**.

```
# zfs set racktop:poolimport:<pname>:last_import_good=0 bp
```

11. Once pools import, assuming other parts of the infrastructure are online, check the system log and confirm that datasets are accessible. At this point you should be ready to re-start services in the infrastructure that depend on the SAN.

## Related articles

Page:Brickstor is no longer connected to the domain and is unable to rejoin the domain

Page:BSRADM times out during os operations

Page:vBOX or BrickStor not Booting - Decode Beep Codes and LEDs

Page:GUI won't Connect to BrickStor

Page:BrickStor won't rejoin a Domain it was previously joined to