Incorporating Proactive Data Rescue into ZFS Disk Recovery for Enhanced Storage Reliability

Zhi Qiao [1] Song Fu [1] Hsing-bung Chen [2] Michael Lang [2]

[1]University of North Texas, [2]Los Alamos National Laboratory

Introduction

ZFS, as the back end of the Lustre, has been used in many HPC systems. It is designed from scratch to mitigate fault management problems and ease the system administration. It also detects and repairs silent data corruption, and rescues data using software RAID when a disk crashes.

However, ***reactive data rescue*** such as post-failure RAID rebuild involves heavy computations. For exascale systems, it introduces non-negligible overhead because of the following facts.

1. Disk failure becomes the norm when tens of thousands of disk drives deployed.
2. Inevitable data loss during multiple disk rebuilds as we move toward RAID 7[1].
3. Service degradation and even unavailability during RAID array rebuilds.

These motivate us to develop ***proactive data rescue*** schemes for ZFS.

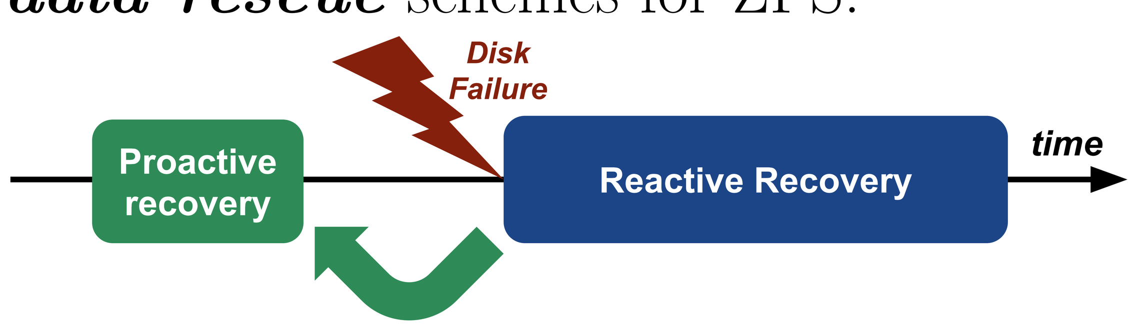

Figure 1: Proactive data rescue uses cloning to copy data from failed drives, thus relieves system from heavy computations, shortens the rescue time, and avoids service unavailability.

ZFS Recovery Performance

To understand the characteristics of ZFS disk recovery, we extensively measure and analyze its performance using following configuration parameters.

RAM Size	Disk Media	Storage Utilization%	RAID Type
8-32 GB	SATA HDD SATA SSD	15%-75%	RAID 5-7 (RAIDZ 1-3)

Following figures shown our findings.

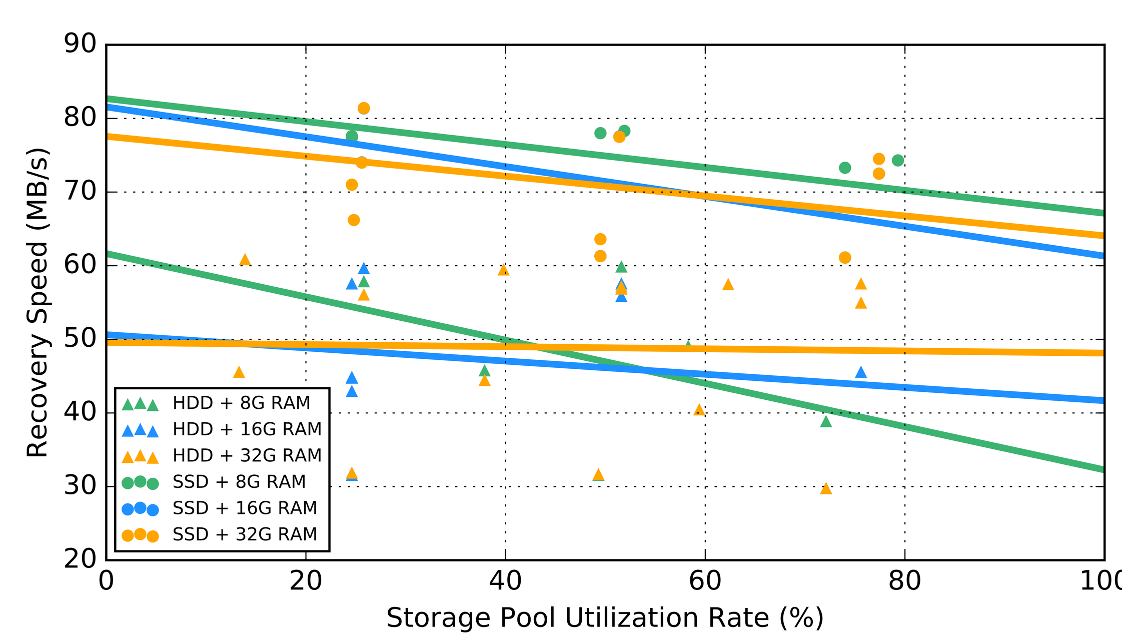

Figure 2: Recovery speed drops up to 60% as storage gets full. Meaning the recovery time will increase exponentially with the storage utilization rate.

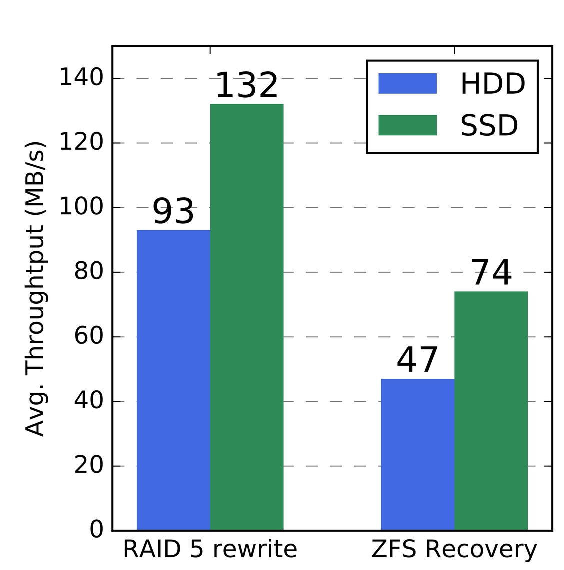

Figure 3: ZFS recovery performance is CPU dominated, so it can only achieve half of the RAID 5 rewrite throughput.

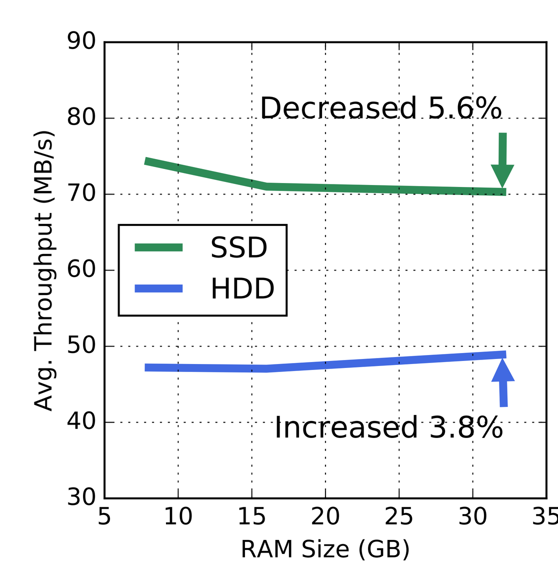

Figure 4: ZFS is known for RAM hunger. But increasing RAM size yields negligible impact on recovery performance.

Proactive Rescue Strategies

Disk failure prediction technologies, like [2][3], continue evolving and are capable of accurately forecasting when disk failures will occur. They enable ZFS to replace failing drives and use spare ones in advance.

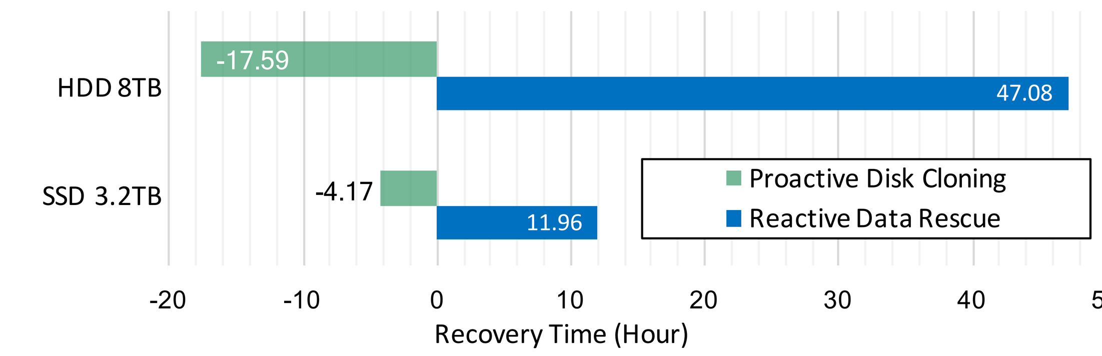

Figure 5: Use *Proactive Disk Cloning* to rescue a full disk of data can avoids up to 2 days of service degradation.

As a Copy-on-Write (COW) filesystem, ZFS uses the *"active"* tag to keep track of data, and repair them if a disk crashes. So full disk clone is not always faster.

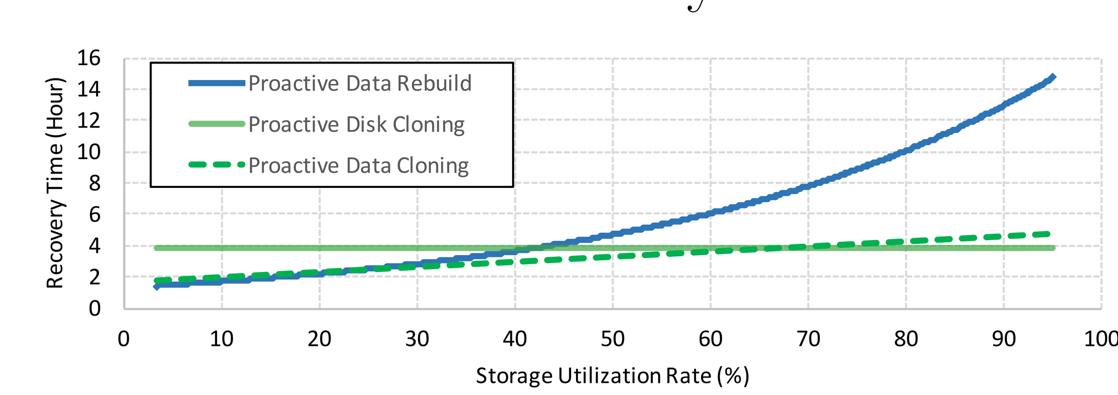

We propose following strategies for proactively rescues data.

Strategies	Descriptions	Operations	Speed
Proactive Disk Cloning	Full disk cloning	R/W	Fast
Proactive Data Cloning	Cloning data with *"active"* tags	Seek R/W	Slower
Proactive Data Rebuild	Rebuild data with *"active"* tags	Seek R/W XOR	Slowest

Strategy Optimization

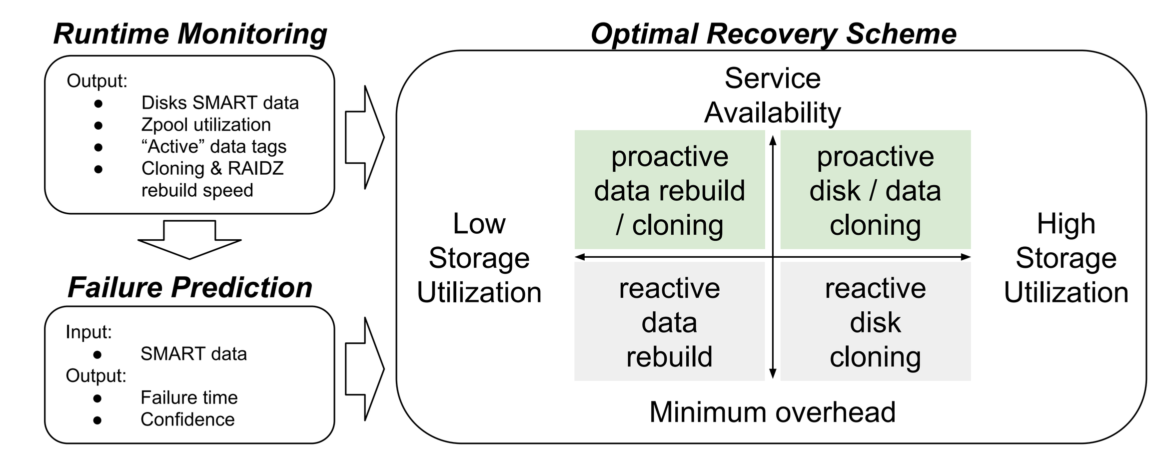

Conclusion

1. Discovered that storage utilization seriously affected ZFS recovery speed.
2. Developed proactive data rescue strategies to avoid costly post-failure disk rebuild.
3. The optimal strategy need to consider runtime system status.

References

[1] Adam Leventhal.
Triple-parity raid and beyond.
Communications of the ACM, 53(1):58–63, 2010.

[2] Song Huang, Song Fu, Quan Zhang, and Weisong Shi.
Characterizing disk failures with quantified disk degradation signatures: An early experience.
In *IEEE International Symposium on Workload Characterization (IISWC)*, pages 150–159, 2015.

[3] Mirela Madalina Botezatu, Ioana Giurgiu, Jasmina Bogojeska, and Dorothea Wiesmann.
Predicting disk replacement towards reliable data centers.
In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining(KDD)*, pages 39–48. ACM, 2016.

[4] Jeff Bonwick, Matt Ahrens, Val Henson, Mark Maybee, and Mark Shellenbaum.
The zettabyte file system.
In *Proceedings of the 2nd Usenix Conference on File and Storage Technologies (FAST)*, volume 215, 2003.

LA-UR-17-27639