

# Enabling performant, dense and reliable AI infrastructure



Paperspace



RAIDplus

PLiOPS  
EXTREME DATA PROCESSOR

## Introduction

Paperspace is a high-performance cloud computing and ML development platform for building, training and deploying machine learning models. Tens of thousands of individuals, startups, and enterprises use Paperspace to iterate faster and collaborate on intelligent, real-time prediction engines.

Paperspace currently has several PB of active customer storage in production across various regions which are being served to customers using Linux based NFS servers. Each server is equipped with twenty two 2TB SATA SSDs and multiple fast network interface cards. To protect the customer data, they were using traditional hardware RAID solution from Broadcom.

## Challenge

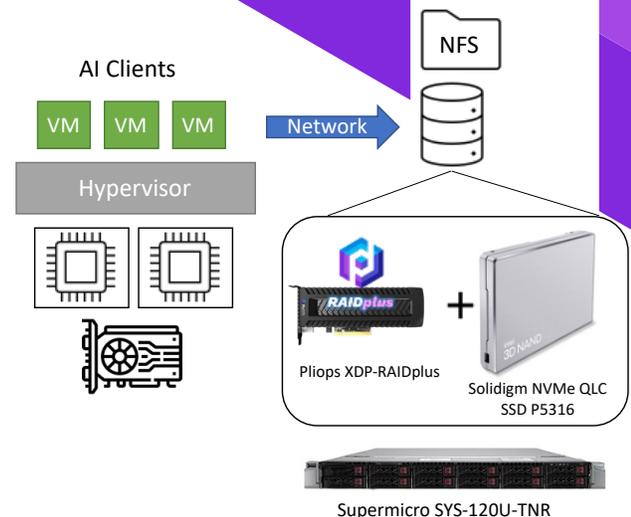
Paperspace had three major challenges with their existing NFS storage infrastructure.

First was growth. Paperspace is growing rapidly, and their customers are adding up to 40TB of data per day. To stay ahead of the demand, Paperspace usually adds multiple storage nodes at a time. Due to supply chain constraints, they were unable to get the Broadcom MegaRAID cards needed to protect their customer data in time.

Second was performance. Paperspace was performance limited by the SATA SSD throughput due to bandwidth constraints. They wanted to move to NVMe drives for better performance but were concerned about storage density and reliability.

Third was storage availability. Paperspace was facing weekly drive failures resulting in extended recovery time. In addition, there was severe performance degradation during the drive rebuild process.

Based on these challenges, Paperspace realized that they needed a storage solution that was denser, more performant with better reliability and recovery times.



# Solution

Paperspace tested Pliops XDP-RAIDplus with Solidigm QLC NVMe SSDs which not only met their requirements, but also exceeded their expectations. The configuration details are as described in the table below.

<b>Accelerator Card</b>	Pliops XDP-RAIDplus PCIe x8 form factor Low Profile HHHL (6.6" x 2.536")
<b>Storage</b>	Solidigm P5316 QLC NVMe SSD 15.36 TB x 8 drives
<b>Server</b>	2U Supermicro SYS-120U-TNR
<b>CPU</b>	Dual Intel® Xeon® Silver 4310 CPU @ 3.30 GHz
<b>Memory</b>	32GB x 8 DDR4 2667 MT/s
<b>OS</b>	Ubuntu 20.04.01
<b>Kernel</b>	5.8.0.53

# Results

5x throughput improvement compared to the older SATA SSD solution.

The storage density per node was also increased by 8x to effectively 128TB per node even with RAID 5 protection due to built in compression.

Drive rebuild efficiency has been improved by 7x with recovery time 2x faster.

Disk Endurance has been enhanced due to all data being written sequentially. Paperspace has faced no drive failure in past 6 months.

Cost per TB of the new NFS server is 33% less than older SATA SSD solution.

The Pliops solution has won over the Paperspace technical storage team.

The Paperspace team has already implemented several servers in production with the Pliops XDP-RAIDplus card solution and are continuing to add more.

To schedule your own XDP-RAIDplus evaluation, please reach out to [demo@pliops.com](mailto:demo@pliops.com)  
Learn more about XDP-RAIDplus at [pliops.com/raidplus](https://pliops.com/raidplus)

# About Pliops

Pliops multiplies the effectiveness of organizations' infrastructure investments by exponentially increasing datacenter performance, reliability, capacity, and efficiency. Founded in 2017 and named as one of the 10 hottest semiconductor startups by CRN in 2020 and 2021. Pliops global investors include NVIDIA, Intel Capital, SoftBank, Western Digital, KDT, and Xilinx. **Learn more at [www.pliops.com](https://www.pliops.com)**

# Key Solution Benefits

- Up to 5x throughput improvement
- 23x throughput improvement during drive rebuild
- 2x improvement in drive rebuild times
- 8x storage density improvement
- Cost per TB reduced by 33%
- SSD Write Amplification Factor reduced to 1
- Benefit from Solidigm QLC NVMe SSD density with the effective endurance of TLC
- RAID 5 like protection with no capacity penalties
- No dedicated hot spare required with Virtual Hot Capacity (VHC) feature
- Protection against 2 sequential drive failures with VHC