

Test Report of SSD Storage for Antarctic Ground and Balloon-Borne Missions

C. Kulesa

VERSION HISTORY

Version #	Implemented By	Revision Date	Approved By	Approval Date	Description of Changes
1.0	<i>Craig Kulesa</i>	<i>19 April 2013</i>			<i>First complete draft</i>
1.1	<i>Craig Kulesa</i>	<i>15 July 2014</i>			Added results for Intel S3500 based on Supercam testing

Rationale for testing:

The star camera failure on the recent flight of BLAST (2012-13 season) using the Intel model 320 solid state disks has caused a reinvestigation of large-scale storage for future STO and HEAT Antarctic deployments.

Is HEAT or (GUS)STO sensitive to the Intel 320 problem?

No. Both systems boot from a small read-only flash disk for robustness. The data system mounts the large SSDs long after system start and under manual commanding. If the large SSD storage fails, the system is still usable. This is demonstrated by pulling the SSD power on a live system! The worst case scenario is that the SSD fails while mounted and the kernel blocks on the I/O failures, making the system unresponsive. This scenario was not seen in testing but **is** possible. In this case, the data system is simply rebooted and the large SSDs are then not mounted manually. We simply continue to operate – without them.

Drives tested:

From the HEAT and STO-1 projects. OCZ flash 64 GB, Patriot 64 GB flash (flown in STO-1), Samsung 840, Crucial M4, Innodisk 3MP, Intel 320, Intel 710, Intel S3500.

Testing:

The testing operation uses large OS-based parallel writes. Thousands of directories & millions of small files plus one large one were copied, synchronized, CRC checked, then deleted using 8, 32, and 128 parallel processes. Power is not pulled during the first set of runs. When the set repeats, the power is then pulled at random intervals and the system is re-evaluated at restart. Up to 2500 power cycle events are logged and the number of CRC failures recorded.

Results Table:

Drive	Test: Parallel writes to 2 TB	Test: Parallel writes with power failures
OCZ flash	Fails CRC past 1 TB	Fails CRC (stopped after 250 cycles)
Patriot flash	Pass	Fails CRC (stopped after 250 cycles)
ATP flash	Fails CRC past 1 TB	Fails CRC (stopped after 1000 cycles)
Samsung 840	Pass	Fails CRC (stopped after 2000 cycles)
Crucial M4	Pass	Fails CRC (stopped after 1000 cycles)
Innodisk 3MP	Pass	Fails: drive becomes unresponsive
Intel 320	Pass	Fails: Bad context 13x. Drive has 8 MB capacity and requires reformatting.
Intel 710	Pass	Pass
Intel S3500	Pass	Pass

Summary of Results:

A failure mode in which power cycling causes the Intel 320 drive to lose its entire contents and require low level reformatting has been reproduced in the lab at standard temperature and pressure. Other failures in OCZ, Samsung, Crucial, Innodisk drives cause data corruption under power cycling and undervoltage conditions. The Intel 710 drive passed all tests: it has an enterprise-class controller with sufficient charge reserve to safely complete all write operations before the drive powers off completely. Recommend only enterprise-class drives in mission-critical applications. *Update 7/15/2014: The 710's enterprise replacement, the S3500, also appears to pass all write tests and power cycling tests and has been selected for Supercam.*

An additional study showing similar results on drives low low capacity is also available online:

http://lkcl.net/reports/ssd_analysis.html

Lessons:

- a. Use enterprise-oriented disks; you tend to get reliability-selected NAND, much more margin on data lifetime, write-cycles, and a wider operating environment. Ignore consumer disks; ask instead “what drives do data centers use”? The cost penalty is not that much considering the implications for flight.
- b. Size the drives to provide margin on storage using a single drive (500+ GB/drive assuming GUSTO requirements). This allows the use of multiple drives for mirroring and redundancy, not capacity.
- c. Test samples before selection to **near-destruction**. Do gentler **verification testing** on the 'flight units' before deployment.

Final Summary:

Despite the very public problems with the 320 series drive, the more expensive enterprise class Intel drives are the only drives in which the controller properly manages power failure events. The Intel 710 drive and its S3500 replacement should be the drives specified for large format storage for (GUS)STO, HEAT and related missions.