NVMe Gen-4 Thermal Management: Too Hot To Touch

Grace Chen, Phison Electronics

Storage is Getting Hotter

- Problem
  - With the arrival of PCIe Gen-4, SSD storage is getting hotter
  - Heat management concerns move down from the Data Center to Client and Mobile platforms
- Solution
  - Going forward, it is essential to also cool the SSD
  - Passive cooling: Heat Sink
  - Active cooling: Airflow
- But how hot is too hot?
  - NAND likes similar temperature ranges as people
  - 25C = Comfortable
  - 40C = Working hard but still ok
  - 80C = Shutdown

Does bandwidth impact power?

- The PS5012 and PS5016 are ideal for bandwidth vs power
  - Same: Architecture, Process, DRAM, NAND
  - Difference: PCIe Gen-4, Clocks, Bus Speed
- Observations
  - Strong correlation between performance and power
  - Performance increase 47% vs Power increase 44%
  - Process reduction is of limited benefit as it only affects CTRL Core
- Conclusion
  - We can expect full speed Gen-4 x4 to reach M.2 power limit
  - 3.5 GB/s SSD = 3.7W
  - 5 GB/s SSD = 5.4W
  - 7.8 GB/s SSD = 8.25W (M.2 Limit)

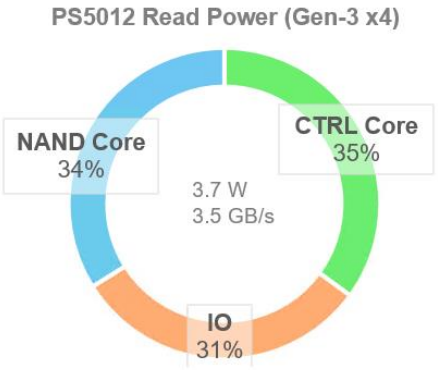

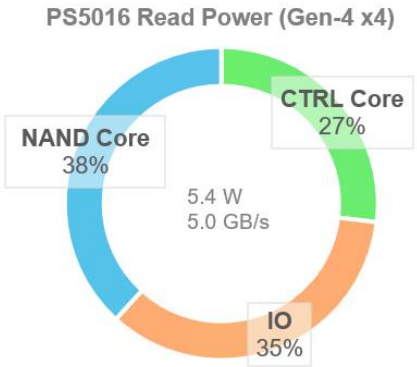

Does room temp affects an SSD?

- **Observations – Room at 0C** (freezing)
  - NAND that is powered is always +15C above ambient/room temperature
  - PS5012 @ 3.5 GB/s IO generates 50C heat in the NAND; faster SSD will generate more heat
  - Heating is not instant; it takes ~200 sec to reach steady state

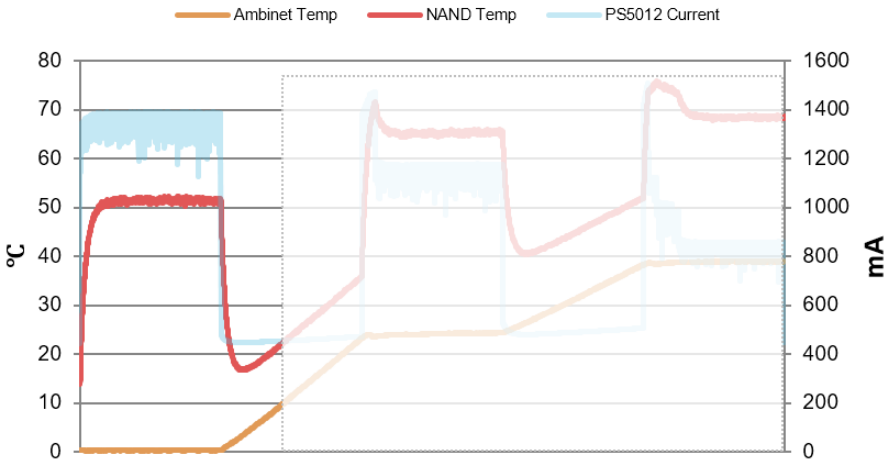

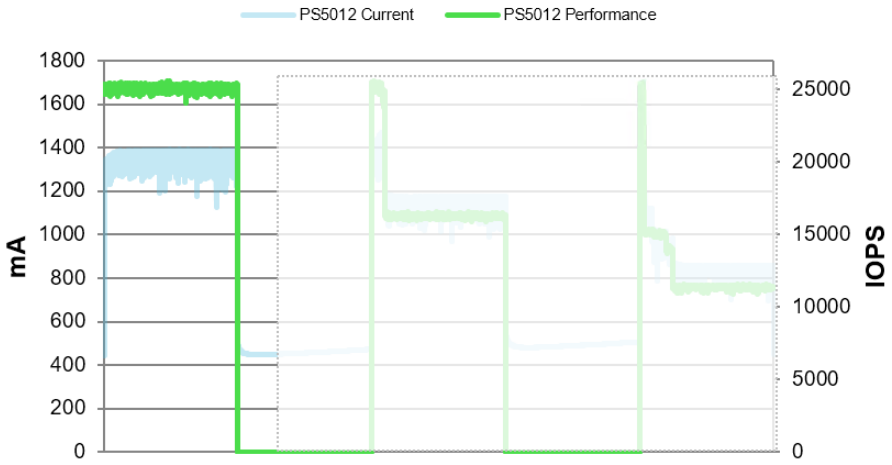

Does room temp affects an SSD?

- **Observations – Room at 25C**
  - Room temperature determines the floor, but the energy used still generates +50C of heat
  - Thermal throttling triggers to keep the NAND below max operating temperature
  - The energy used by NAND is constant, so we reduce the number of operations per second

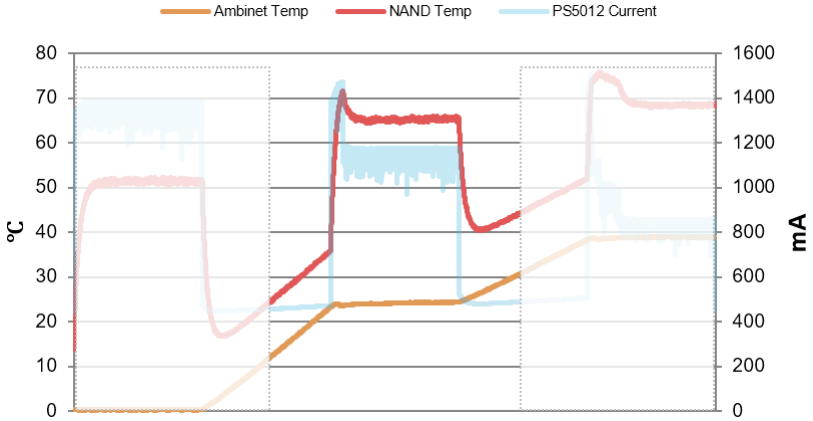

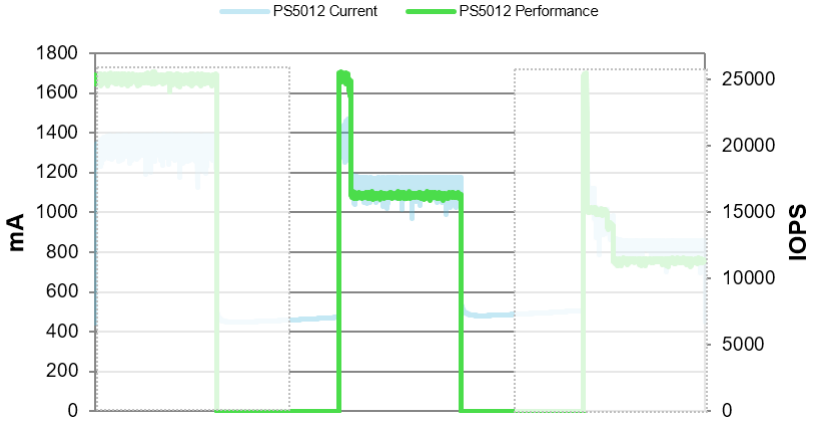

Does room temp affects an SSD?

- **Observations – Room at 40C**
  - The same trend continues with respect to +50C thermal energy added to the floor
  - Due to the higher floor, even more performance throttling is required to stabilize the NAND temperature
  - Thermal throttling activates at the same temperature, but it has to iterate through more steps to reach balance
  - The energy output of the NAND is contestant, so we have to reduce ops/sec even more to allow the heat to dissipate out

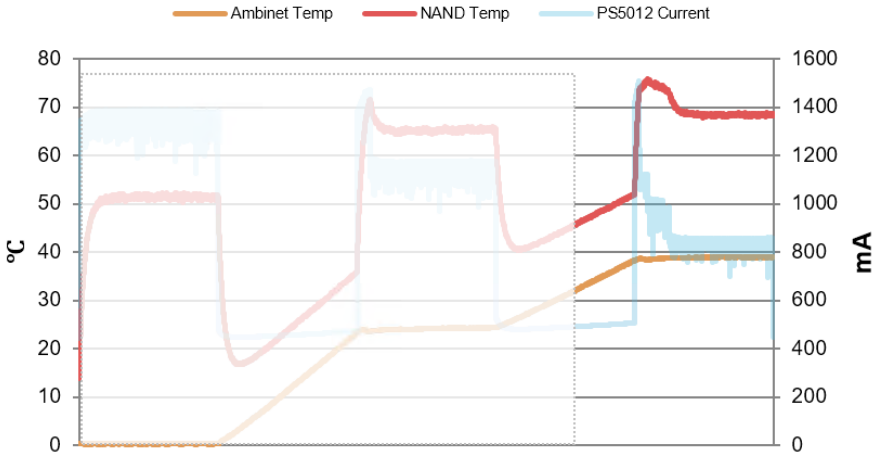

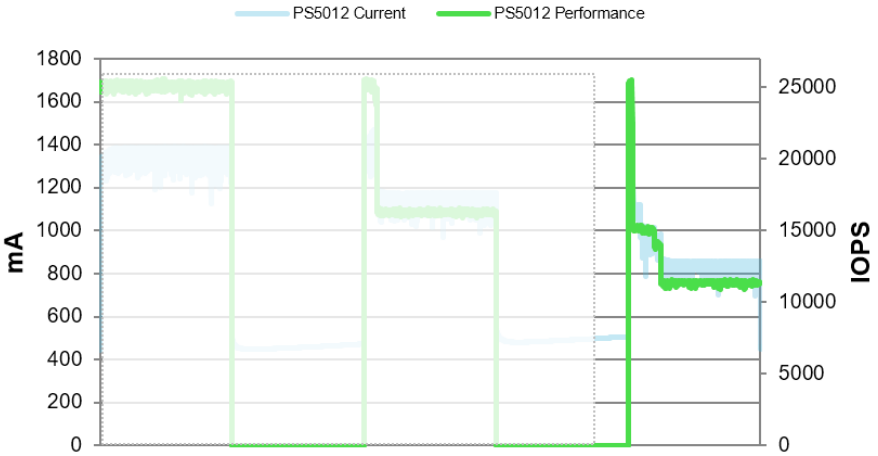

Do heatsinks and airflow affect an SSD?

- **Observations – Adding a heatsink**
  - As with the PS5012, the room temperature sets the floor and determines the headroom before the SSD has to throttle
  - Like the PS5012, the Gen-4 x4 SSD throttles to ~2 GB/s when the room is at 25C
  - ASIC epoxy compound conducts 1 W/mK, static air conducts 0.03 W/mK and aluminum conducts 237 W/mK
  - Moving the heat away from the NAND faster reduces the amount for throttling needed to maintain NAND temperature
  - Adding a heatsink alone isn't enough to allow the SSD to operate at full speed

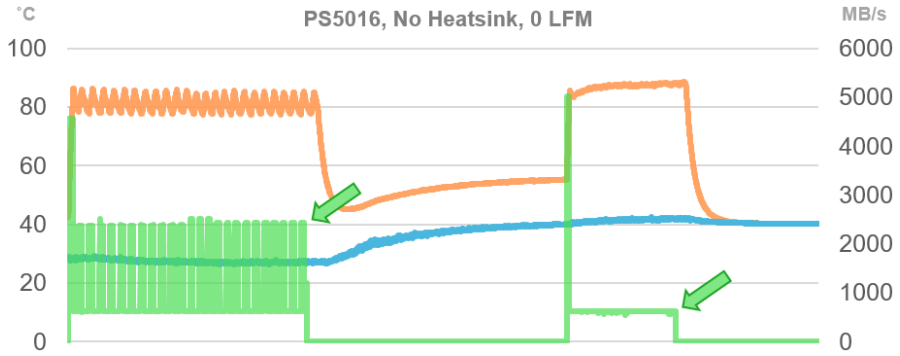

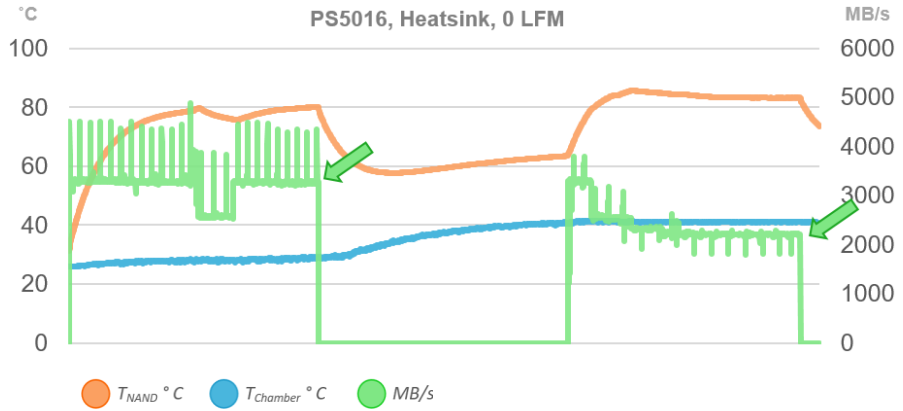

Do heatsinks and airflow affect an SSD?

- **Observations – Adding airflow**
  - Adding moderate airflow has a similar result to adding an aluminum heatsink
  - The test chamber is pre-heating input air to maintain a steady temperature
  - Adding airflow alone isn't enough to allow the SSD to operate at full speed

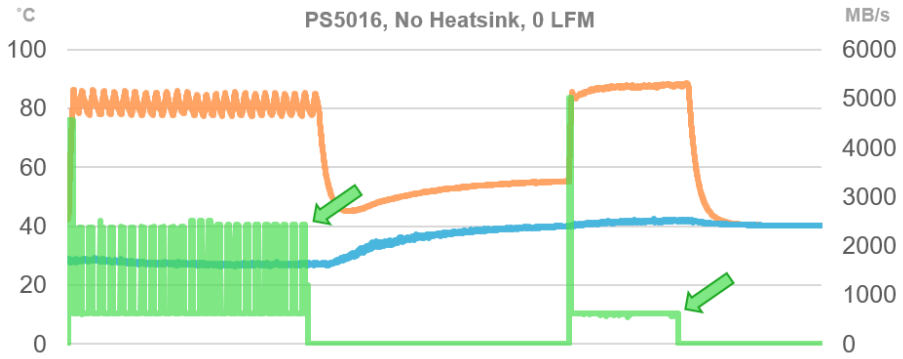

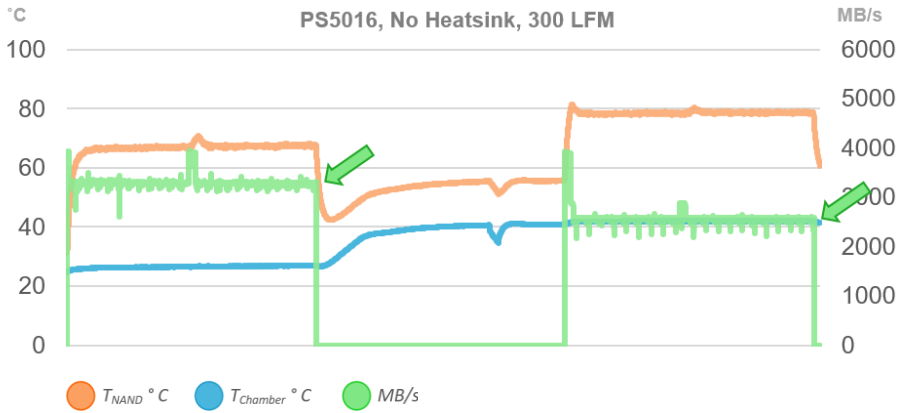

Do heatsinks and airflow affect an SSD?

- **Observations – Adding a heatsink + airflow**
  - The SSD can run at full speed once a heatsink and moderate airflow are added
  - Note that copper has 2x the thermal conductivity of aluminum and could be substituted for airflow
  - The cooling solution can trade off: Air Speed, Input Air Temperature, Heatsink Material and Heatsink Size
  - As SSD's go faster and generate more heat, more consideration must be given to the cooling solution

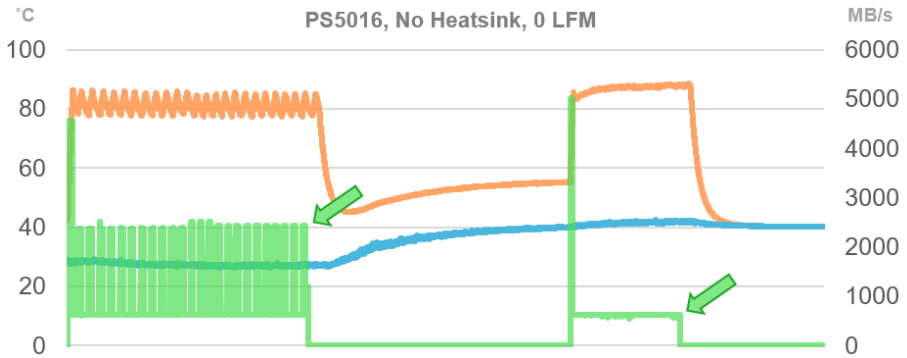

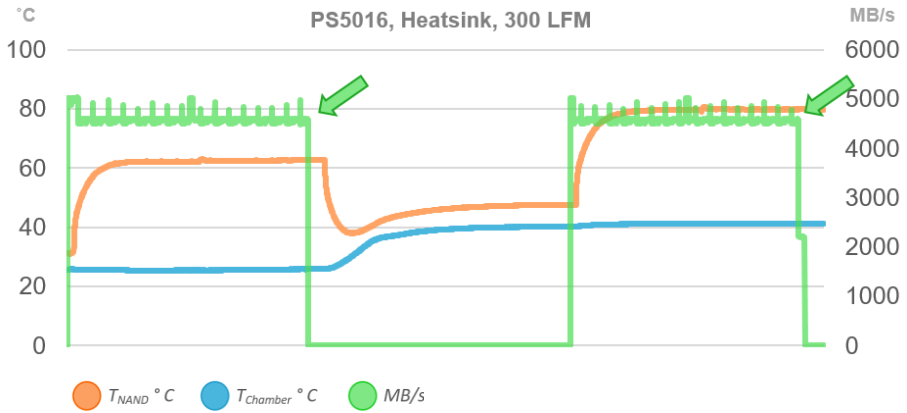

Key Takeaway

1. Power and heat scale with MB/s
2. Foundry process reduction can only reduce CTRL Core power (approx. ~5% overall SSD improvement per tech node)
3. The room (or case) temperature has a direct impact on the SSD thermal floor
4. Heat sinks and airflow can move heat out of the SSD
5. As SSD’s get faster, more consideration has to be given to cooling

PHISON

Knows What You Need

"PUSHING BOUNDARIES"

YOUR PCIe GEN-4 SSD LEADER

VISIT US AT BOOTH #219!