What QoS Means to Computing Applications – Showcasing 5 nines

Introduction to Quality of Service

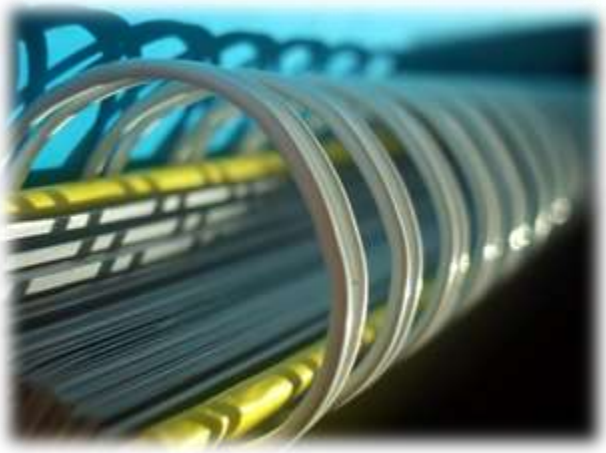

In modern computing architectures it has been shown that being the 'fastest' is not always the best solution for computing and storage problems. While we continue to see the shear IOPS performance of storage increasing with the move from SATA/SAS to NVMe as well as PCIe Gen 3.0 to Gen 4.0, it has been found that maximizing the bandwidth of these interfaces does not necessarily lead to an overall better performing system solution. This is especially true on platforms now available that can support 48 drives or more in a 2U chassis, where we find that individual drive performance specs become irrelevant when the system level interfaces are incapable of aggregating that much raw bandwidth (BW). In effect, the drives go faster, but the system interface BW is still limited – you simply can't force any more data into the pipe. Limited by CPU performance and/or networking interfaces, having that much sheer speed at the device level creates an asymptotic performance cap, and this bottleneck comes at a high cost. A storage device delivering a higher interface bandwidth requires an expensive storage subsystem using a high-power controller with many channels, additional memory and a large die size.

A new trend has been emerging that addresses this very problem. Whether it's called performance optimization or a focus on Quality of Service (QoS), the approach is to have storage, like NVMe drives, that provide a high internal Read bandwidth with an optimized Write performance, tuned to let customers maximize the application use of many drives per system. With this focus on latency, each storage device has the ability to respond to host requests in an extremely predictable manner. A common measure for this consistency is listed as a number of "nines". For a storage device to deliver a 5 nines value, it must complete tasks within a specific target window 99.999% of the time. Another way to describe this is that less than 1 in 10,000 commands completes outside a window of time regardless of the activities within that device.

A Travel-based example

An analogous way to look at this is via the model of boarding an airplane. The most efficient way to ensure success would be to board from the back, window seat first, then middle then aisle, progressively moving forward to the front. This would provide a boarding process that could achieve a 99.999% or 5 nines efficiency. However, today's boarding process is by class, much like user requests to storage, un-aligned and out of order. Where first class fills up the front creating a blockage to other locations, then status members or other priority needs,

randomly dispersed around the plane and then the rest of the passengers or other standard data. While boarding first and possibly even status members can achieve happiness for some passengers or users, and aids in the unload (read) process, once you start getting to the truly random process, your efficiency drops well below even 1 nine of QoS. Boarding passengers, or users looking for data, are frustrated trying to reach seats. The flight attendants or data scientists are doing constant reminders or status updates to move out of the way and so on. Moving to a way to manage data with 5 nines efficiency requires a more planned out method or technology to manage the random data coming in or being sent out. Let's look at a more data driven example.

NGD Systems – Solving Computing Application Qos Needs

To provide this improved level of customer experience, a new approach to data storage is required, a solution that maximizes the Read experience and optimizes the Write performance, power consumption and scalability of the platform. NGD Systems has developed just such a solution with our patented Elastic FTL, Programmable ECC and data placement algorithms integrated into our controller architecture. The core of this platform manages Read vs Write and ensures that house cleaning efforts do not negatively impact host QoS. When plotted, a narrower distribution of response times indicates higher consistency. The figure below shows the 5 nines distribution for the NGD Systems Intelligent Storage platform and that of a competitive solution running identical workloads. For this test, the drives were run against an FIO (Linux test platform) script with a customer workload that mimics their use case of mixed Reads and Writes.

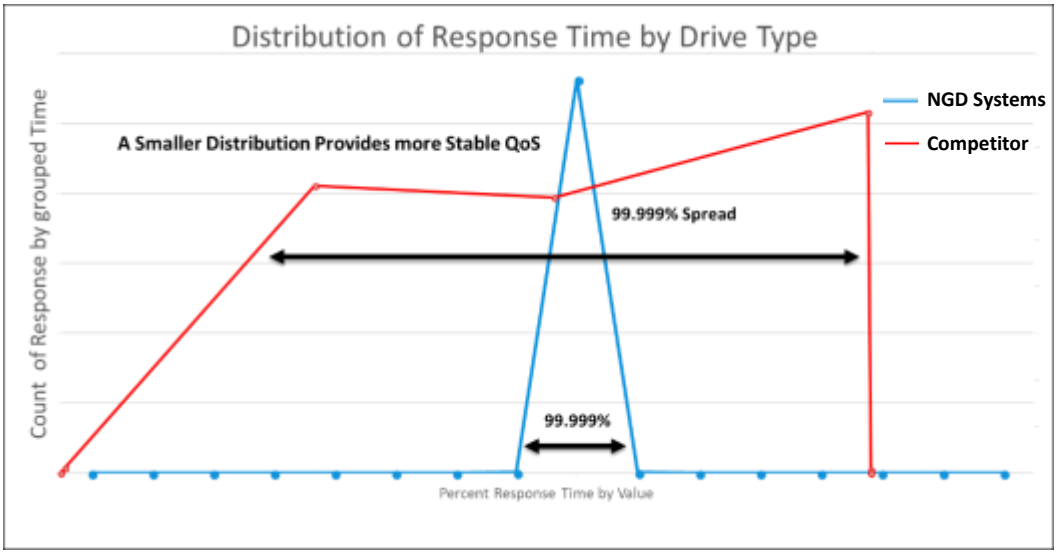

The result: Optimized response time performance that scales across platforms and workloads in an NVMe drive that only consumes 12 watts of power.