NETLIST

# Flash Storage for Content Delivery Networks (CDN)

Whitepaper


## N1962 NVMe Gen4 SSD

Optimized for CDN Applications and Workloads

Netlist's Enterprise N1962 Solid State Drives (SSDs) are built using PCIe Gen4 protocol and efficient NVMe interface. By combining a feature-rich PCIe Gen4 controller with Enterprise-Grade (TLC) NAND and best-in-class firmware, Netlist's N1962 SSDs scale in capacity (up to 16TB) as well as performance to handle the toughest enterprise application workloads.

## Summary

The Netlist N1962 delivers industry-leading low latency, while simultaneously reducing the latency inconsistency to significantly improve application Quality of Service (QoS). From the CDN's perspective, this level of performance consistency has huge impacts on the ability to meet SLA commitments to their customers. The way to mitigate latency spiking is to over-provision their networks; for a CDN this means running their content delivery hardware at a lower level than its maximum capacity. This translates into more hardware, meaning higher acquisition and operational costs. By reducing latency inconsistency, the content delivery hardware can be run closer to its maximum performance and capacity. The N1962 provides exactly the type of performance required that allows CDNs to meet customer SLAs, while at the same time minimizing hardware and operational costs.

## Benefits

### Extremely Consistent Latency

The N1962 delivers predictable and consistent latency critical to CDN workloads. The minimizing latency inconsistency, balancing average latency against latency spiking is critical to maintaining CDN QoS and SLAs.

### Sustained Throughputs

CDN applications and workloads requires tremendous sustained throughput for maximum efficiency. The N1962 is optimized to accommodate very large throughput & High-IOPs requirements and sustain such performance.

## Background

As their names imply, the business of Content Delivery Networks (CDNs) is to deliver content to end-users. The content being delivered varies from videos, movies, and music to web pages to software executables (both complete packages and updates) to digital downloads of licensed content. Estimates of the amount of content delivered globally by CDNs ranges from 15%-30% of total traffic (GlobalDots, 2021[1]), and some experts believe that CDNs will carry up to 70% of global web traffic by 2022 (especially if you consider Amazon AWS, Microsoft Azure, and Google Cloud as CDNs).

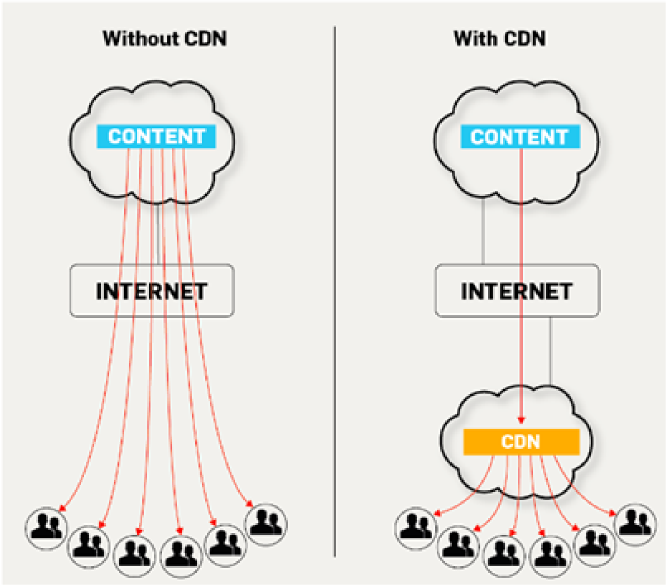


It goes without saying that for CDNs, delivering content rapidly is critical to their success, especially in the case of content that is being viewed or listened to in real time. This is typically measured as a combination of a CDN's total bandwidth or capacity (how much traffic they can carry), and their average latency (how long it takes on average for content to reach a consumer). CDN customers also often have Service Level Agreements (SLAs) in their contracts which specify the performance and capacity that the CDN must deliver.

Most CDNs operate hundreds or thousands of what are known as Points of Presence (PoPs), with most in close proximity to the metropolitan areas that they serve (in the diagram to the right, the PoP is the cloud with the "CDN" label in it). The purpose of having PoPs is to significantly reduce the latency due to distance experienced by customers, and to provide more redundancy in the case of internet failures. For instance, Akamai (one of the largest commercial CDNs) has over 240,000 servers operating across 4,200 locations in 135 countries, with a total bandwidth of over 300 Tbps. This averages each PoP having roughly 60 servers and 70 Gbps of bandwidth. In many cases, the number of PoPs that a CDN has is a determining factor in their competitiveness. The downside of this is that many of the PoPs in metropolitan areas inhabit very expensive real estate, so PoPs often invest in the latest generation technologies to increase performance and reduce their footprint while maintaining SLAs with their customers.

## The Problem That Storage Performance Presents for CDNs

Obviously, it is a problem for CDNs to under-deliver from a performance perspective for a variety of reasons. This is one of the big reasons that CDNs put a lot of energy into designing their networks to meet their performance objectives. In addition to the contractual requirements embodied in SLAs, poor performance negatively impacts the customer satisfaction of end-users, which also potentially impacts the revenue of the company whose content the CDN is delivering.

Key to the ability of CDNs to meet their performance goals is storage performance. While storage capacity and bandwidth needs are fairly easy to plan for, the predictability of storage latency is another matter. For content delivery networks (CDNs), latency performance inconsistency is a huge problem – it causes buffering of video streams, delayed or lost retail transactions, or interrupted downloads. The figure on page 3 (from StorageReview.com) illustrates this phenomenon for a NoSQL workload.

1: https://www.globaldots.com/resources/blog/content-delivery-network-explained/

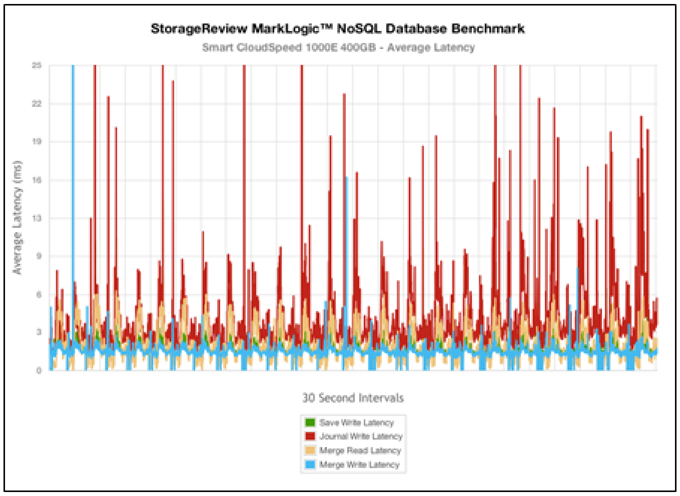


Source:https://www.storagereview.com/review/smart-storage-cloudspeed-1000e-enterprise-ssd-review

While flash-based solid state drives (SSDs) provide significantly better average latency than hard disk drives (HDDs), they can still suffer from latency "tails" (also known as latency "spikes") – seemingly random increases in latency that can be 3x-5x (or more) of the average device latency. These spikes can occur hundreds or thousands of times per second. If a group of SSDs are utilized in a multi-drive array, the frequency will be significantly magnified since each SSD's individual latency will impact the arrays' latency on any given transaction. Moreover, the "size" of the spikes relative to the average latency will be significantly more pronounced, since an SSD array (depending on the array technology being used) will have significantly reduced average latency, while the size of the latency tails will remain consistent.

While there are a number of causes for inconsistency in storage latency, these issues tend to worsen over time. Latency consistency also varies considerably with the type of flash utilized (SLC, MLC, etc.), and by manufacturer. Since different models may have different levels of capacity overprovisioning (the % of storage above the rated device capacity) and different ratios of dynamic random-access memory (DRAM) on an SSD relative to the amount of flash on the SSD. The workloads themselves, in particular the ratio of reads to writes on a device, also impact latency consistency. All of these factors affect write endurance, which requires "firmware intervention" of the SSD controller to move blocks around (for instance), impacting SSD response times.

Finally, while latency inconsistency is primarily (but not exclusively) associated with write operations, it also affects read operations as well. Since modern SSD parallelize I/O streams, a read operation could end up behind a write operation inside the pipeline when a latency spike occurs. This will cause the read operation to have latency equal to that of the write operation, which tend (on average) to take 10x the time that a read operation requires. This results in a huge read latency spike when compared to the average read latency. Even read-heavy workloads can have read spikes due to Worst of all, these spikes appear to be non-deterministic when looked at "from the outside" of the SSD, making system-level mitigation difficult at best.

## Performance of Netlist N1962 SSD for CDN Workloads

The Netlist N1962 series of enterprise SSDs encompass the latest technologies and highest performance of any Netlist SSDs to date. The Netlist N1962 series is a U.2 PCIe SSD with a x4 interface, which is compliant with the NVMe version 1.4 specification. It provides 7.0GB/s of throughput and over 1.6M IOPS of performance. It supports features such as self-encryption (AES-256), sanitize, and crypto-erase. It also supports up to 32 different namespaces with quality of service (QoS) settings for each namespace. It utilizes 96-layer 3D TLC NAND, with 176-layer coming in 2022. The improved performance of the PCIe Gen4 Series 1962 SSD versus PCIe Gen3 SSDs is shown in the figure below, where you can see that sequential write performance nearly doubled, and random read IOPS increase to as high as 50% more.

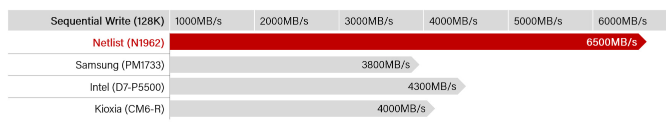


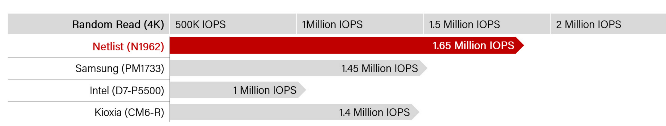

In performance tests run by an independent lab using a simulated SQL Server workloads (15K virtual users and four VMs), the transactions per second (TPS) output of the four drives at their maximum was nearly equal. The Netlist N1962 performed the best at 12,650 TPS, while the Samsung PM1735 performed the worst at 12,625 TPS. However, the average latency of the Netlist N1962 SSD was roughly half the latency of the Kioxia CM6 and CD6 SSDs, and roughly one-fifth the latency of the Samsung PM 1735 SSD.

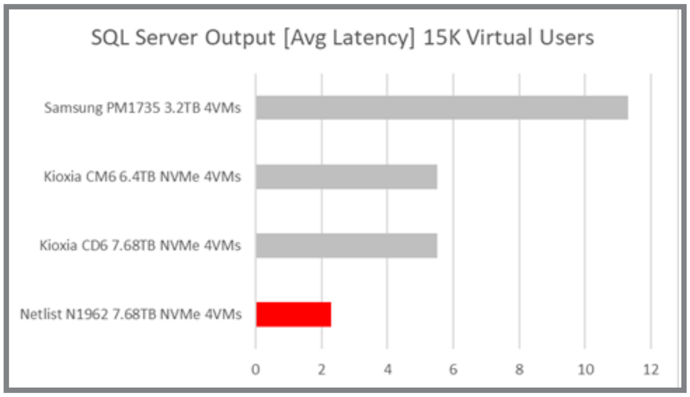


The sequential and random read performance charts below further illustrate that the latency of the Netlist N1962 stays nearly flat and remains the lowest of all the drives tested, even as the throughput and IOPS continue to increase. This is in marked contrast to the competitive drives, where latency increases significantly as throughput and IOPS increase.

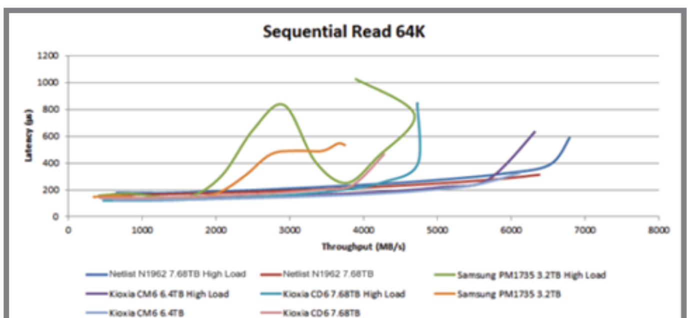


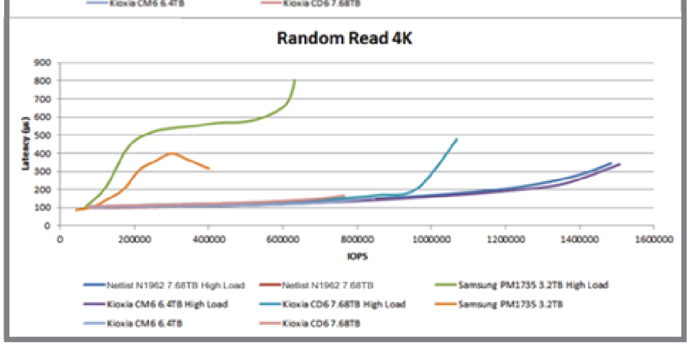


This improvement in SSD latency can be critical to a variety of applications, including databases, web services, cloud service providers, and content delivery networks (CDNs). In these applications, the additional latency translates into reduced performance, which can only be compensated for by adding more hardware (and more capital and operational costs).

## Performance Netlist N1962 for CDN Workloads

A leading content delivery network (CDN) was looking for a Gen-4 PCIe SSD that they could utilize in their operations. Critical to the CDN was the minimization of latency inconsistency – balancing average latency against latency spiking is very important to maintain their QoS. They tested the Netlist N1962 Series SSD across a number of scenarios that are representative of their typical workloads, including real operational testing. These included read, write, and mixed I/O scenarios with a variety of random I/O sizes consistent with the CDN's workload. Testing was performed over extended periods of between 6 hours to 24 hours to quantify the effect of "aging" on device latency and latency consistency. An example of the performance of the N1962 SSD for 99th percentile read latency for an actual workload is shown below. You can notice that the read latency was nearly flat, with extremely high latency consistency.

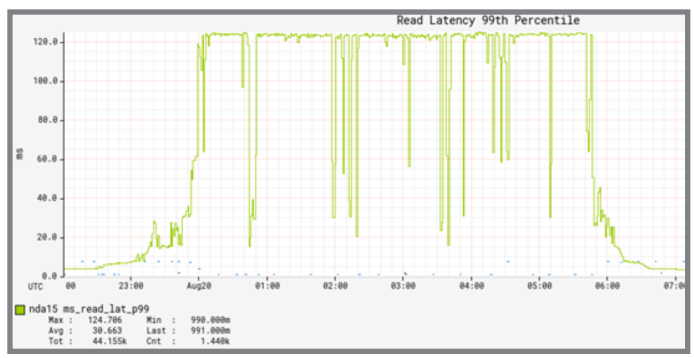


From the CDN's perspective, this performance consistency has huge impacts on the ability to meet service level agreements (SLAs) with their customers. The only means to mitigate latency spiking is to over-provision their networks; for a CDN this means running their content delivery hardware at a lower level than its maximum capacity.

This translates into more hardware, which also means more acquisition and operational costs. By reducing latency inconsistency, the content delivery hardware can be run closer to its maximum performance and capacity.


## Summary

The Netlist N1962 delivers industry-leading low latency, while simultaneously reducing latency inconsistency. From the CDN's perspective, this level of performance consistency has huge impacts on the ability to meet SLA commitments to their customers. The only means to mitigate latency spiking is to over-provision their networks; for a CDN this means running their content delivery hardware at a lower level than its maximum capacity. This translates into more hardware, which also means more acquisition and operational costs. By reducing latency inconsistency, the content delivery hardware can be run closer to its maximum performance and capacity. The N1962 provides exactly the type of performance required that allows CDNs to meet customer SLAs, while at the same time minimizing hardware and operational costs.