

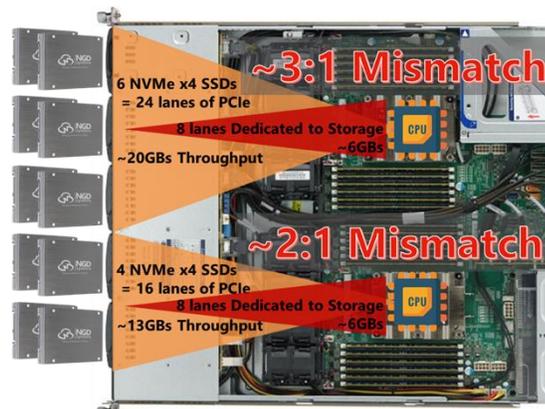
NGD Systems Computational Storage Drives (CSDs) Brings New Life to Elastic.

Ensuring Elasticsearch has scale and growth even with limited available resources to manage the data.

What is Elasticsearch? Elasticsearch is a distributed, open source search and analytics engine for all types of data, including textual, numerical, geospatial, structured, and unstructured. Known for its simple REST APIs, distributed nature, speed, and scalability, Elasticsearch is the central component of the Elastic Stack, a set of open source tools for data ingestion, enrichment, storage, analysis, and visualization.

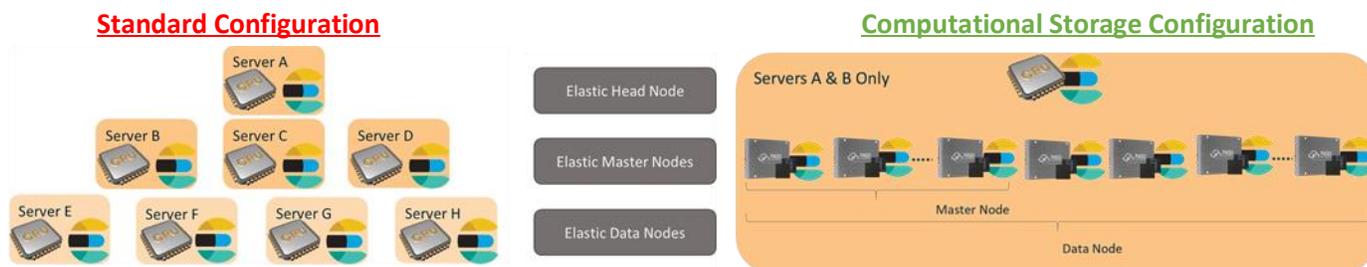
In today's growing marketplace for data analytics the overwhelming need to store and manage data in localized, sometimes resource constrained situations can create havoc for developers and Data Scientists. The goal of this whitepaper is to highlight how taking a distributed Elasticsearch toolset and distributing across NGD Computational Storage Drives can provide the necessary support to increase performance and reduce system complexity to accomplish the goal of proper data management.

Deploying Elasticsearch tools today require access to CPU, DRAM, network, and disk. These resources are always disproportional in today's server and storage architectures where there is always more disk available compared to DRAM and CPU. This requires that data being pulled from disk for analytics must be copied, operated on, discarded, and repeated till the entire data set is complete. Using Computational Storage Drives to offload the initial work required, you can substantially reduce the amount of round robin data movement and time associated with it to get the desired results.

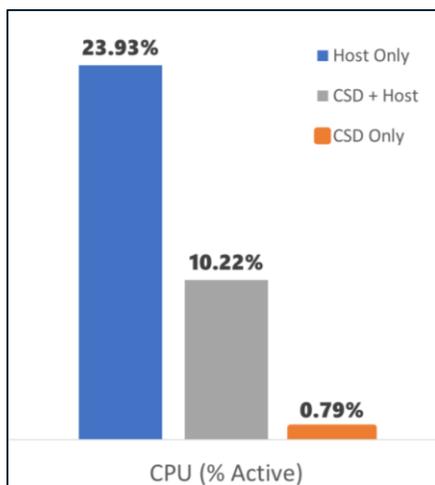


If we take the opportunity to minimize the data transferred from the disk to the DRAM and CPU the efficiency of the system increases while the complexity of the platform is reduced. Since there are no 'new' components required, like GPU or FPGA accelerators, and storage is required anyway, then we can show that the use of CSDs from NGD Systems does not increase system design cost but does provide significant performance improvements by "pre-processing" the dataset while it is on the SSD.

To understand how we can accomplish this task, we need to look at the configuration of the Elasticsearch solution on a given server platform. In this instance, we showcase how many servers may be required to accomplish a deployment of the search tools. By choosing to offload portions of the platform to the Computational Storage SSD's we can provide more value to the overall system performance and Total Cost of Ownership (TCO).



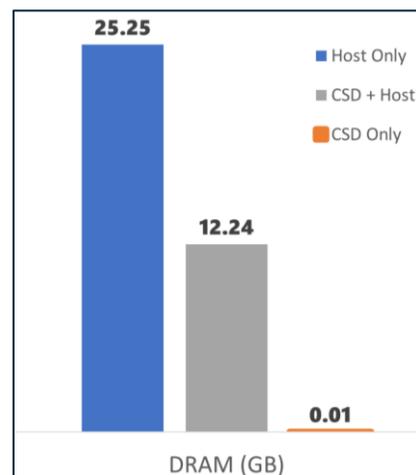
What does this new architecture look like? In our example we reconfigure the platform into only the bare minimum number of servers to guarantee resiliency. We then rely on the CPU resources which exist on the NGD Systems Computational Storage Drives (CSDs) by overlaying the Data and Master onto the SSD's. This allows us to release system resources (like CPUs, GPU's and host memory) back to the system to be used for other tasks. As you can see above, a very simple deployment frees up 6 host system resources for other use cases while still providing better results and lower overall TCO to the Elasticsearch environment.



The ability to all but remove CPU requirements from the search efforts can be dramatic. Our configuration options allow the host to be used for other tasks while the CSDs are doing the search and inference work required of the Elastic platform. If the needs of the system are dedicated fully to just this work, then it can be accelerated by adding the CSD resources and consuming minimal added CPU effort.

When choosing to deploy architectures to accomplish tasks, a key part of the solution is to ensure flexibility of the resources being deployed. By providing a platform to give the user control of where and when host resources are used provides ultimate resource optimization and minimizes the need for additional hardware.

There is also a desire to ensure that such a deployment is not only useful, but deployable with power consumption being considered. We have an active power monitoring system on the primary Elastic host to show how the movement of the search efforts being done at the Computational Storage Drive level provides additional relief to the system power budget which further enhances overall TCO. This allows us to see the amount of power consumed by both the CPU and the drives as well as the aggregate power consumption of the system. All of this continues to show savings at the platform level regardless of the performance of the system.



With proven savings in terms of CPU, system design and system power budget, the last step is to show the amount of value in the improved performance of a given system. As you can see from the graphic, the ability to free up substantial amounts of memory (DRAM) resources required while still improving the net performance shows that offloading and distributing the search functions to the local storage devices provides substantial system level gains. This is without increasing complexity or costs to build the platforms.

To learn more, please reach out to Info@NGDSystems.com

