

NGD Systems: Closing the Storage-Compute Gap

July 2017

Executive Summary: The advent of high-performance, high-capacity flash storage has changed the dynamics of the storage-compute relationship. Today, a handful of NVMe flash devices can easily saturate the PCIe bus complex of most servers. To address this mismatch, a new paradigm is required that moves computing capabilities closer to the data. This concept, which is known as “In-Situ Processing”, provides storage platforms with significant compute capabilities, reducing the computing demands on servers. For applications with large data stores and significant search, indexing, or pattern matching workloads, In-Situ Processing offers much quicker results than the traditional scenario of moving data into memory and having the CPU scan these large data stores. In-Situ Processing can also enable existing applications to scale to much greater levels than is currently possible with discrete storage and computing resources. Because computation capabilities scale linearly as storage is added into compute nodes, In-Situ Processing can enable new classes of applications for enterprises and cloud service providers.

NGD Systems: Closing the Storage-Compute Gap

Background: The advent of high-performance, high-capacity flash storage has changed the dynamics of the storage-compute relationship. When storage systems were built around hard disk drives, this balance favored the CPU, which often sat idle waiting for the storage system to respond, even if there were hundreds of disk drives all operating in parallel. With flash storage media, the reverse is true. A handful of NVMe flash devices can easily saturate the PCIe bus complex of most servers. The graph in Figure 1 below shows the read time **in minutes** just to read a given amount of data across various speed 32-lane PCI Express® (PCIe™) busses, from PCIe 3.0 to PCIe 5.0 (which will come out in 2-3 years). As you can see on the chart, it takes over nine hours to read 1PB (1024TB) across a PCIe 3.0 bus; even on a PCIe 5.0 bus, this time is over 2 hours and 18 minutes. And if 1PB seems like a LOT of storage, remember that by 2019 it will only take sixty-four (64) 16TB flash drives for 1PB of capacity, which even today could fit within a single storage-class server chassis.

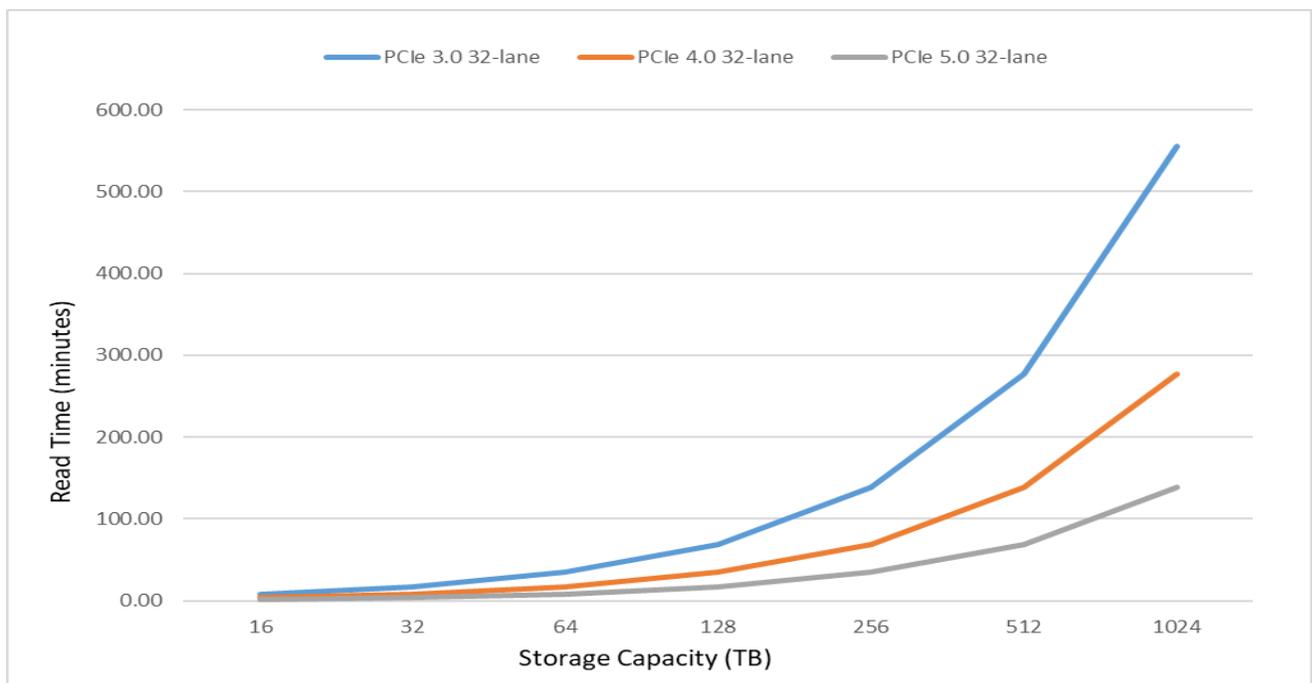
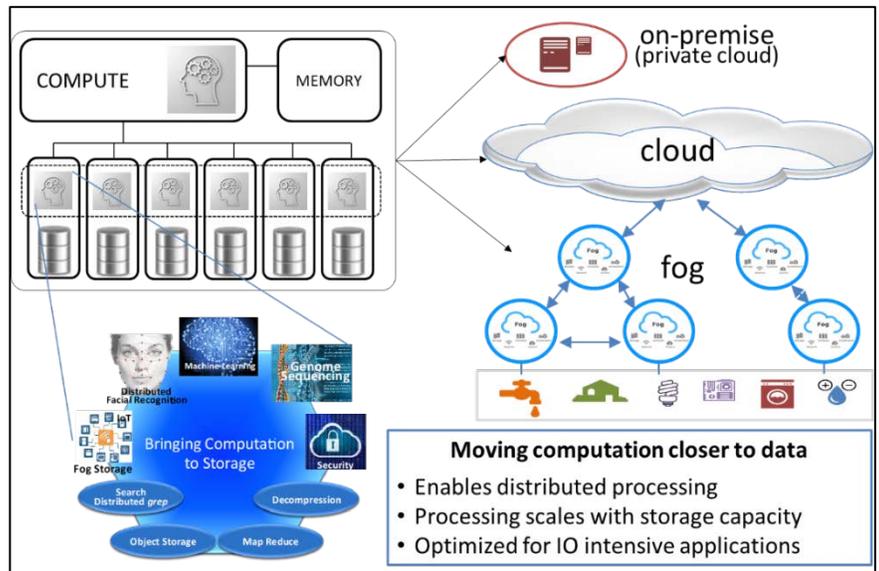


Figure 1: PCIe Read Times by Storage Capacity

The impact of this gap on problems with large dataset is immense. Clearly, this takes many big data problems such as pattern/ facial recognition and searches out of the realm of economical real-time processing. The only option that is available with conventional computing architectures is to split the dataset across a cluster of servers, which significantly increases costs (both CapEx and OpEx), power consumption, and most important, the “footprint” of the system. In many cases, this also makes such problems out of the realm of reality.

The Solution – In-Situ Processing: What if you could avoid moving this data, and instead put the processing capability right in the flash drive? Moreover, if you could read the data at the speed of the SSD’s internal busses, you could reduce load times for each processor by 5X-10X when compared to the PCIe bus. Such an approach would certainly eliminate the “data ingest” problem above. This is the concept of “in-situ” processing, which is at the heart of NGD Systems storage technology solution. In-situ processing is an extension of the “data locality” concept utilized distributed processing systems such as in Hadoop-based distributed big data solutions, where processing work is moved to where the relevant data is, rather than moving the data.



To understand how this concept changes the storage-compute dynamic, we will compare three configurations based on Open Compute Platform (OCP) servers, each with a 1PB dataset:

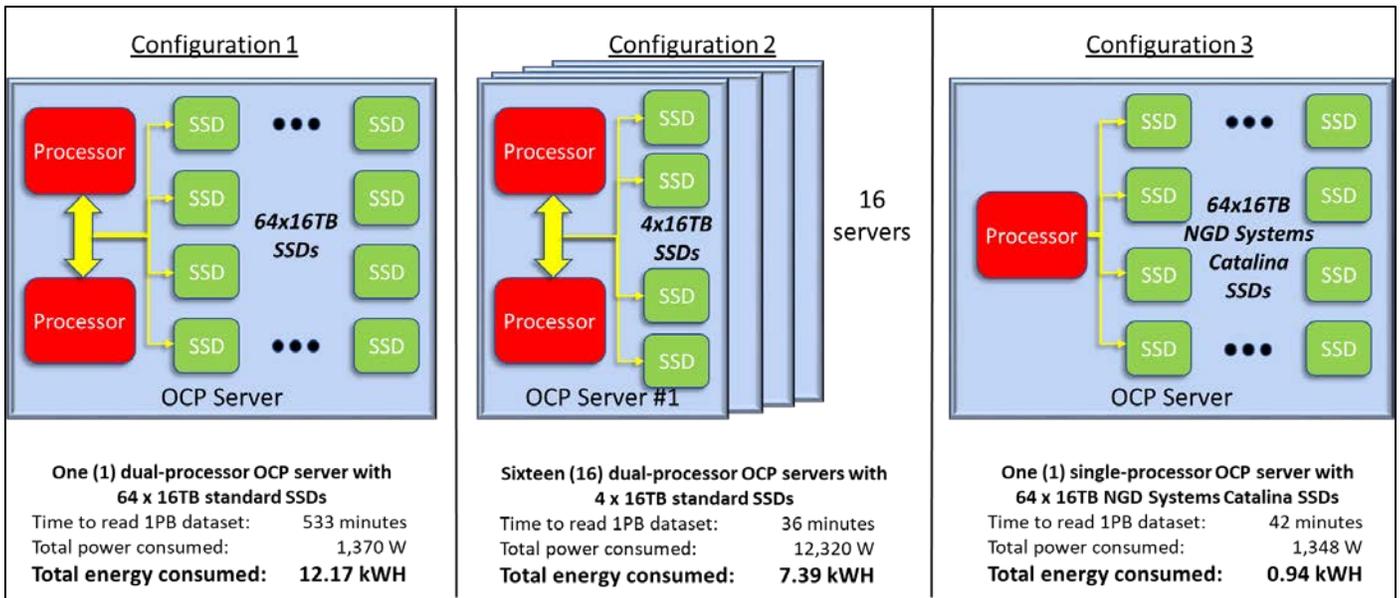
Configuration 1: One dual-processor OCP server with sixty-four standard 16TB U.2 SSDs.

Configuration 2: Sixteen (16) dual-processor OCP servers, each with four 16TB U.2 SSDs.

Configuration 3: One single-processor OCP server with sixty-four (64) U.2 NGD Systems Catalina 16TB drives, each consuming 12W.

These configurations are shown on the following page, along with the time to access the entire dataset (1PB), and the total power and total energy consumed by each configuration. Clearly, the access time and energy consumption improvements from the one dual-processor OCP server with standard SSDs to the one single-processor server with Catalina SSDs is considerable - it is an order of magnitude better on both metrics. While the 16-server OCP cluster configuration has a slightly better access time than the single Catalina-based server, the power consumption is significantly greater - the server cluster would use over 12KW of power, and nearly 7.4 kWh of energy. The server cluster would also take up 34U of space (including 2U for network switches), which is nearly an entire server rack. For many use cases, the cluster configuration is impractical from a power consumption/cooling standpoint, as well as a cost standpoint.

Use Cases: There are a variety of use cases which lend themselves to in-situ processing. In general, these use cases are read-intensive with large numbers of parallel operations that are performed on the dataset, especially pattern matching, indexing, and searching. We will explore a few here where In-Situ Processing enables new levels of performance and/or capabilities.

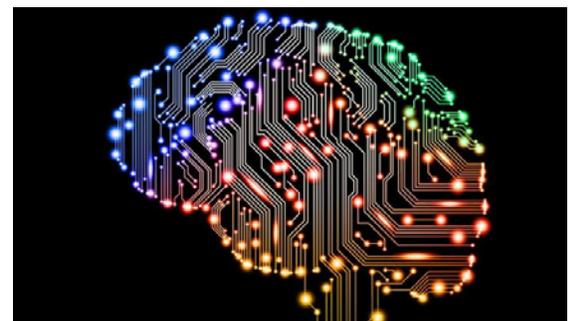


a) Hyperscale Storage: By 2020, roughly 485 hyperscale data centers will contain 57% of all data stored in datacenters worldwide². Moreover, the traffic within these will quintuple by 2020². For a hyperscale storage architect, finding a way to avoid moving data from device to device even in the same rack will become a paramount need, as will reducing the power consumed to process this data. In-Situ Processing offers the best approach today to mitigate both the data movement and power consumption issues. More importantly, In-Situ Processing significantly improves execution times for applications that make use of operations such as indexing, parallel searches, and pattern matching – all of which are becoming more and more common as AI/ML applications proliferate.

By 2020, Hyperscale Data Centers Will House:		Today:
47%	of all data center servers	21%
68%	of all data center processing power	39%
57%	of all data stored in data centers	49%
53%	of all data center traffic	34%

Source: Cisco Global Cloud Index, 2015-2020.

b) Machine Learning/Embedded Artificial Intelligence: Artificial Intelligence (AI) has achieved very promising results in areas such as computer vision, speech recognition, and natural language processing. For example, in a smart city data such as videos or images captured from many distributed cameras need to be automatically processed using video analytics; i.e., object detection, object tracking, facial recognition, image classification, and scene labeling. Vast amounts of cloud data can be used for training on powerful platforms to create generalized yet accurate models. However, there is an opportunity to move the inference, and in some cases the learning, onto “smart” storage. This is where one can take advantage of a distributed system composed of the



computing resources available in a multitude of SSDs to implement shallow learning with weightless neural networks.

c) Edge Computing for IoT (OpenFog): The Internet of Things (IoT) aspires to collect mountains of data by instrumenting nearly everything in our existence. This creates challenges both in storing and processing this data. Gartner Research expects that there will be 20.4 billion IoT devices connected to the internet by 2020, generating 5X the data that we generate today. Cisco expects that IoT devices will generate 403 zettabytes/year by 2018². A single connected aircraft can generate 40TB of data per day², while an autonomous car may generate 2TB of data per hour, according to Intel³. Clearly, this data cannot be simply be “sent home” to datacenters, even to those in the cloud. The key to making constructive use of this IoT data deluge is the ability to process it “in place”, which is exactly what in-situ processing capability on a flash storage device provides for IoT. By processing this data “in place”, IoT nodes can make intelligent updates to the cloud without the burden of the power, space, or cost of a server.



NGD Systems: Leading-Edge Smart Storage: At NGD Systems, we are blazing the trail of in-situ processing for storage devices. Our in-situ storage architecture and Catalina products make the deployment of many large dataset applications possible and practical, whether from an access time, power/cooling, or real estate standpoint. If you would like to find out more on how you can benefit, please contact us for further discussions on how In-Situ Processing can help solve your data center and business issues.

Footnotes:

- 1 – [Gartner Press Release, Feb 7, 2017](#)
- 2 – [Cisco Global Cloud Index: Forecast and Methodology, 2015-2020 \(2016\)](#).
- 3 – [Patrick Nelson, Network World, Dec 7, 2016](#)