

## REPORT REPRINT

# The ball starts rolling for NGD's computational storage

**FEBRUARY 04 2021**

**By Tim Stammers**

The emerging field of computational storage has attracted attention as a way to accelerate machine learning and other data-intensive applications in datacenters and edge locations. Startup NGD is one of the companies pioneering the technology, and says it is beginning to see large deployments of its products.

---

THIS REPORT, LICENSED TO NGD, DEVELOPED AND AS PROVIDED BY 451 RESEARCH, LLC, WAS PUBLISHED AS PART OF OUR SYNDICATED MARKET INSIGHT SUBSCRIPTION SERVICE. IT SHALL BE OWNED IN ITS ENTIRETY BY 451 RESEARCH, LLC. THIS REPORT IS SOLELY INTENDED FOR USE BY THE RECIPIENT AND MAY NOT BE REPRODUCED OR RE-POSTED, IN WHOLE OR IN PART, BY THE RECIPIENT WITHOUT EXPRESS PERMISSION FROM 451 RESEARCH.

451 Research

**S&P Global**

Market Intelligence

### Introduction

NGD is one of a handful of startups that has pioneered the concept of computational storage. Aimed at accelerating processing for data-intensive workloads in both edge and datacenter deployments, the concept has also gained support from large vendors. In 2017 NGD began shipping FPGA-powered computational storage drives, and in 2019 it began shipping the industry's first (and to our knowledge only) ASIC-powered CSDs on the market. That distinction is significant in terms of the addressable market of the devices. Although customer adoption has been slowed by the COVID-19 pandemic, NGD says it has scored multiple deployments, including one large implementation that is in progress now and another that will begin very soon, with others expected to start during 2021.

### 451 TAKE

NGD's potential customers need to evaluate its products in their datacenters or laboratories, and COVID-19 lockdowns have made this impossible in many of NGD's target markets. This will have slowed sales, even though the company attempted to offset the damage by building a remote-access laboratory. Nevertheless, NGD is now beginning to see traction. The company's major distinction compared with its small number of rivals is that its ASIC-based architecture allows it to address a wider market by offering a platform that is easily adapted to handle existing and future workloads. The market will decide how valuable this quality will be, but we think VMware's recent demonstration of machine learning incorporating NGD's devices (see below) underlines its potential market reach, especially for edge computing.

### NGD's product developments

To our knowledge, NGD is the only company currently shipping CSDs that do not use an FPGA to provide computational services, but instead uses an ASIC. Developed in-house by NGD, the ASIC is called the Newport platform. The chip makes NGD's drives easier to program than rival smart SSDs, and therefore suitable for use by a wider range of potential customers, including mainstream enterprises. This may appear counterintuitive – in other situations, FPGAs are more flexible than ASICs. However, defining or modifying the computational tasks to be handled by FPGAs in CSDs requires firmware skills that are generally restricted to hyperscale or other advanced IT organizations. Those skills are not needed for NGD's drives, as the Newport ASIC includes Arm cores that host a Linux OS, which can run conventional software, including existing applications.

As a separate virtue, NGD's NVMe drives offer very high physical densities and low power consumption, alongside high claimed performance. Some customers have deployed NGD drives purely because of these qualities and have not used them as computational devices. In November 2020, NGD launched its first ruler-format drive, which is an EDSFF E1.S (short) device that offers up to 12TB of raw capacity. The E1.S format is in the initial stages of adoption, and very few E1.S drives are currently on the market. As a comparison with NGD's device, Kioxia (formerly Toshiba) recently began sampling an E1.S drive that tops out at just 4TB, according to its public specifications.

Since 451 Research last covered NGD, the company has also qualified its Newport platform to run Amazon's entire Greengrass IoT edge software stack, and has demonstrated its drives running Microsoft's Azure Edge software. NGD says future developments will include support for PCIe4 (or perhaps a jump straight to PCIe5) and possibly drives connecting directly to Ethernet.

## REPORT REPRINT

During VMware's virtual VMworld conference late in 2020, VMware made a detailed demonstration of NGD's drives handling analytics work within a vSAN hyperconverged cluster. VMware says the demo was part of an exploration of ways to support workloads at edge locations requiring high physical densities for machine learning and training. The demonstration involved the Greenplum analytics database and calculation of optimal routes for a ridesharing or taxi service, with processing taking place on NGD's computational drives and on GPUs (the latter using VMware's Bitfusion software). NGD says its devices boosted performance compared with conventional deployments of Greenplum, and highlights what it calls the 'edge-friendliness' of the demonstration. Greenplum involves multiple management nodes, and the NGD drives were able to act as those nodes, reducing the need for server CPUs. VMware says the resulting reduction in hardware requirements was one of the three key points of the demonstration (the other two were the use of Bitfusion software with GPUs and the resilience of vSAN). The demonstration involved two 2U Dell servers, and NGD says they replaced what would otherwise be rack-level mirroring.

### Current and expected production deployments

NGD says a deployment involving hundreds of its drives is in progress at a European developer of autonomous vehicle technology, whose clients include major global car and truck makers. That company began taking delivery of NGD drives in Q4 2020 and is continuing to receive deliveries during the current quarter. Initially, the drives were not being used as computational devices, but were instead selected for use within test vehicles as conventional SSDs because of their high physical density, low power consumption and design support. NGD says the drives are now being put to work as CSDs, providing in situ data processing, as well as data storage. The startup says another major deployment is set to begin this quarter at a US government agency, when final qualification testing is completed. This will see NGD drives used as CSDs or computational devices in rugged servers created by Trenton Systems, a supplier of hardware for defense, government, industrial and commercial applications. No further details of this deployment have been made public.

In another government-funded project, the US Space Force says it will use NGD drives within 'constellations' of satellites. This project has entered an official second phase, according to NGD. The Space Force says the NGD drives will work as CSDs handling machine learning, encryption and authentication workloads. 451 Research assumes that, alongside this computational role, the high physical density and low power consumption of the devices are highly attractive at the ultimate edge location of space orbit.

Elsewhere, NGD says PoCs and qualifications by other large organizations have resumed after delays caused by COVID-19. Microsoft is completing final qualification testing for NGD's drives on several platforms, while another 'super seven' hyperscaler is evaluating NGD's EDSFF-format drive described above, according to NGD. Brazilian aircraft maker Embraer has resumed qualification work that will enable usage of NGD drives as CSDs handling analytics tasks. COVID-19 also delayed qualification tests at a large European online commerce operation, but that work has since restarted. During 2020 this customer told us that it planned to deploy NGD's drives because of their power efficiency, but said it might move to CSD usage after deployment.

### Competition

At present, only a small number of vendors are selling computational storage products. Samsung has been developing a CSD for the last few years. Late in 2020, a partnership of Samsung Xilinx and computational storage startup Eideticom launched a CSD called the NoLoad SmartSSD. The device combines flash drive technology from Samsung with an FPGA from Xilinx and control software from Eideticom.

Like NGD's device, the NoLoad SmartSSD drive handles computational works on the drive itself, but does so on an FPGA rather than an ASIC. The argument for using an FPGA is that the device can be programmed to handle a range of tasks (such as database acceleration, erasure coding or data compression), and is potentially faster at these tasks than NGD's ASIC because they do not involve an OS software layer. But as stated above, that OS is the reason why NGD can counter that its device is more flexible and does not require unusual skills to adapt to any workload.

Alongside the NoLoad CSD, Eideticom's technology is also used in NVMe-connected devices that are classified as computational storage processors. These devices do not store data, but offload storage and other tasks from servers and storage systems. Bittware, a maker of NVMe storage devices and an investor in Eideticom, ships such a device, which includes an FPGA and Eideticom's technology. In November 2020, Eideticom announced a deployment of CSPs at the US Los Alamos National Laboratory. Another startup, Scaleflux, has been shipping FPGA-powered CSDs since 2017. In 2019 Alibaba said it planned to deploy Scaleflux's devices to accelerate its PolarDB analytics database.

Other developments in the sector includes Arm's 2020 release of a processor designed for use in computational storage, as well as the development of standards at the Storage Networking Industry Association and the NVMe Organization. SNIA has released a computational storage architecture and programming model, and the NVMe organization is working to extend the NVMe protocol to cover computational storage tasks.

### SWOT Analysis

#### STRENGTHS

NGD is one of very few companies offering computational storage products, and its devices promise to suit a wide range of IT organizations that includes enterprises.

#### WEAKNESSES

Although NGD is reporting sales traction, it has yet to be seen how strong the demand will be for computational storage products in general.

#### OPPORTUNITIES

Demand for mechanisms that boost processing speeds will increase with the ongoing growth in a range of data-intensive applications that include (but are not restricted to) machine learning and real-time analytics.

#### THREATS

Although Samsung is the only large supplier operating in this market at present, others may join.