CLABBY ANALYTICS

Research Report

Maximizing Performance by Pairing CGG's Insight Earth with Lenovo's P900 Workstation and SanDisk's Fusion ioFX

Introduction

The name of the game in oil and gas exploration is to increase profits by minimizing uncertainties *before* drilling takes place. If the size and location of a prospect is known before drilling occurs, oil and gas companies can better estimate costs, field size and potential return-on-investment.

To locate oil and gas reservoirs on land, seismic waves are sent deep into the earth and are reflected back using a process known as reflection seismology. Sensor geophones are used to gather this seismic data. To locate oil and gas under the sea, air gun arrays are used to generate pressure waves that travel through water and pass through the sea floor. Towed hydrophones are used to gather the reflected seismic data. In each case, vast volumes of data are captured. Oil and gas exploration companies analyze this data in order to construct a composite picture of the underlying geology and any potential related oil/gas field size. These volumes readily cover hundreds of square miles of surface data as well as potentially tends of thousands of feet underground, and can easily reach many terabytes if not petabytes in size.

To analyze these vast amounts of seismic data, specialized software such as CGG's Insight Earth is used to interpret and graphically display the underlying paleo-depositional environment. Further, advanced high performance workstations are required to process this data. *Workstations with "accelerators" exponentially accelerate this process.*

> ***When selecting a workstation to process large volumes of seismic data, it is important to pay close attention to the workstation system design. The primary selection criteria should be to choose a system that can achieve balanced performance among processors, memory, and the I/O (input/output) subsystem.***
>
> ***Remember that a workstation is only as fast as its slowest component. Finding and eliminating bottlenecks is crucial to maximizing performance. Once a bottleneck is identified, accelerators such as graphics adapters which have fast graphical processing units or Flash memory can be used to eliminate bottlenecks and accelerate the processing of large amounts of data.***

A close look at one accelerated workstation design – a Lenovo ThinkStation P900 with a built-in SanDisk Fusion ioFX accelerator – shows that using SanDisk's Fusion ioFX Flash memory accelerator yields a 300% performance improvement over traditional workstation designs. This translates to a 30% improvement in the time needed to execute a typical seismic analysis workflow. In this *Research Report*, *Clabby Analytics* takes a closer look at this accelerated workstation design – and it is easy to conclude that for time sensitive

seismic analysis projects having fixed deadlines, accelerated workstations using Fusion ioFX yield significantly faster results.

The Age of Accelerators

For decades, improving processor performance involved placing more transistors on an integrated circuit. Intel cofounder, Gordon Moore, observed in the 1960s that processor performance doubled every 18 months to 2 years – and this observation became known as Moore's Law. In the mid-2000s, however, this "law" ran into limitations as it became physically impossible to build microprocessors that could exceed 6 GHz performance levels. This effectively ended Moore's Law and chip manufactures were forced to look for new methods to increase processor performance. With the introduction of multi-core processors (using multiple processors on a single die), the industry has been able to continue to increase CPU performance.

However improving performance through multi-core design is only one step in decreasing run-times. Placing large amounts of data in main memory located close to the processors speeds performance, as does improving I/O bus speed. Other approaches such as using solid state drives instead of mechanical disks, feeding processors over fast I/O channels using flash caches can also lead to faster processing. Field programmable gate arrays (FPGAs) can improve communications performance. Graphical processing units (GPUs) can accelerate processing via inexpensive graphics cards. These new system design enhancements are referred to as "accelerators".

> ***At Clabby Analytics, several reports have been written about these new accelerator designs. Reports on VelociData can be found here; reports on The Now Factory can be found here; reports on using solid state technology as memory cache for in-memory database management systems can be found here. This Lenovo/SanDisk report, however, is the first report Clabby Analytics report written on an overall accelerated workstation design.***

The Lenovo/SanDisk Workstation Design

The ThinkStation P900 is a high-performance workstation that offers plenty of I/O bandwidth as well as large internal storage capacity. The SanDisk Fusion ioFX accelerator is solid state drive (SSD) technology that is used as expanded, low latency memory (available in 410 GB or 1.6 TB configurations). By taking advantage of the expansive bandwidth of the ThinkStation P900 – and by using the Fusion ioFX accelerator as expanded memory – vast amounts of sensory data can be fed to x64 CPU and NVIDIA GPU processors for rapid and efficient processing.

> ***The Lenovo P900 workstation combined with SanDisk's Fusion ioFX solid state memory technology is an accelerated system design optimized for processing data-intensive workloads. This design eliminates slow storage. By placing data in accelerated, low latency solid state memory, seismic data can be presented to processors more quickly, thus yielding a sizeable increase in performance.***

The Lenovo ThinkStation P900

Lenovo's ThinkStation P900 is an "extreme performance" workstation based on Intel's E5-1600 v3 and E5-2600 v3 family of processors (and will support future Intel Xeon EP processors). This workstation has:

- 16 DIMM (dual in-line memory modules) slots for memory;
- Supports 16GB DDR4 2133Mhz ECC RDIMMs (registered dual in-line memory modules) for a total of 256GB of DDR4 memory;
- Features support for a wide variety of NVIDIA graphics cards; and,
- Offers plenty of expansion slots to allow for customization (see Figure 1).

Figure 1 – Lenovo's ThinkStation P900

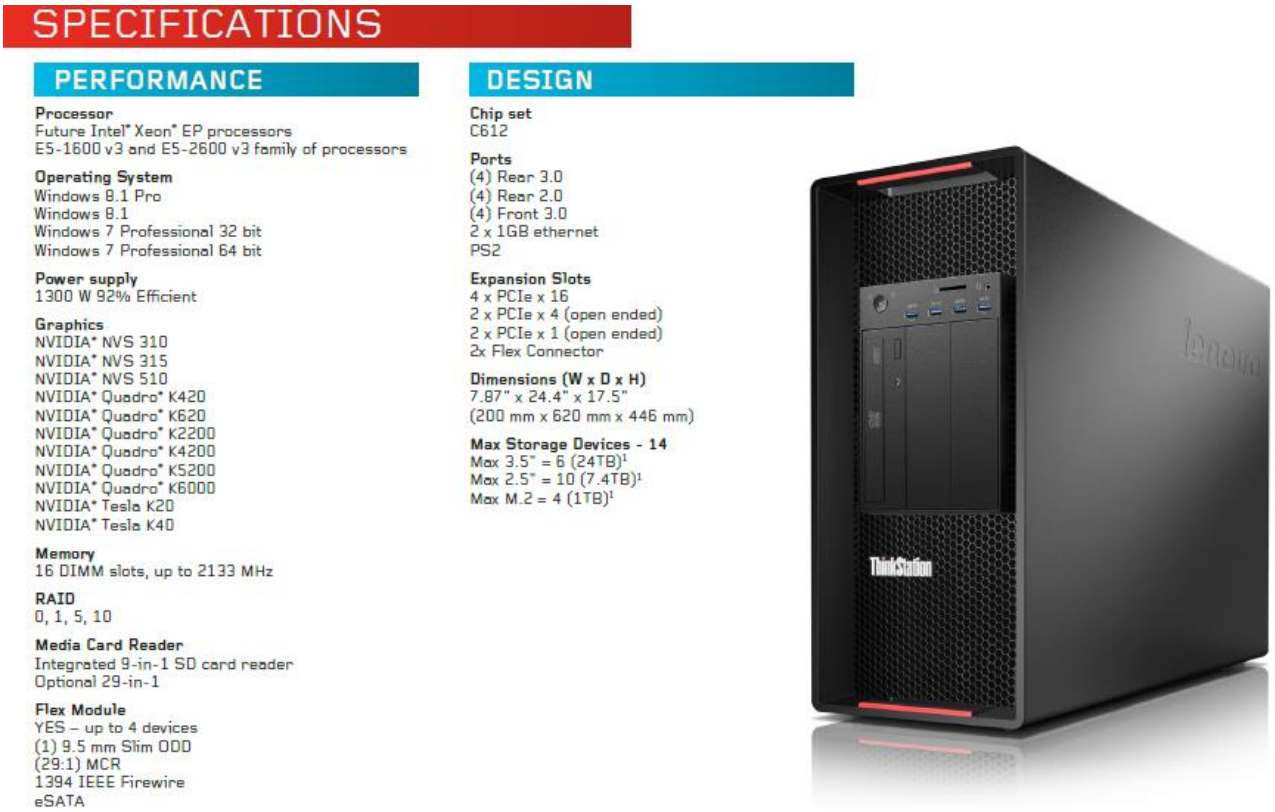

Source: Lenovo – November, 2014

SanDisk's Fusion ioFX

Fusion ioFX uses NAND-based (a type of non-volatile Flash) solid state technology as Flash memory, with its own controller to manage flash data. Furthermore, SanDisk provides a management environment known as "The Fusion ioSphere" that helps information technology (IT) managers manage SanDisk Fusion memory products deployed in a workstation or server environment. ioSphere also monitors performance, and uses live performance metrics and alerts to help ensure optimal performance and operations.

One special feature of the Fusion ioFX environment is its extremely low latency, as compared with pulling data from mechanical drives. This Flash memory can achieve 1.4GB/s read bandwidth with a 42ns average latency time. Depending on the size and type of the mechanical drive being used, Fusion ioFX can be dozens of times faster in terms of latency as compared with mechanical drives. See this formula to calculate the latency on various mechanical hard disk drives.

With extremely low latency and fast throughput, Fusion ioFX solid state drives perform extremely well when processing multithreaded applications such as seismic data analysis and image editing.

Benchmark Comparisons

SanDisk recently examined the throughput and data load times of its Fusion ioFX offering as compared to 7.2 and 10K SATA hard drives, a serial attached SCSI (SAS) drive, and a solid state drive (SSD). These comparisons showed significant performance advantages using SanDisk Fusion ioFX as compared with more traditional drive choices.

Throughput comparison data showed SanDisk's Fusion ioFX with an almost 7:1 advantage over a 10K SATA hard drive – and it was over twice as fast as a solid state drive.

	I/O Testing	
Device	**Burst (MB/s)**	**Sustained (MB/s)**
10K SATA	299.4	117.6
7.2K SATA	318.6	119.4
10K SAS	315.3	177.1
SSD	374.8	276.4
SanDisk Fusion ioFX	794.3	681.4

SanDisk's data load time comparison also showed strong advantages for its Fusion ioFX offering over hard SAS and SDD drives. Note that the only drive technology that outperformed Fusion ioFX was a SDD in the 11GB sequential test, and that difference was statistically negligible.

ISE Project Data Handling/Flow (load times in seconds)				
Device\Workflow	**2GB Mixed (MB/s)**	**11GB Seq. (MB/s)**	**18GB Mixed MB/s**	**32BG Seq. (MB/s)**
10K SATA	18.6	111.9	318.74	279.34
7.2K SATA	17.43	105.4	270.84	247.18
10K SAS	13.18	86.99	243.93	198.1
SSD	9.66	60.23	192.53	130.74
SanDisk Fusion ioFX	6.63	60.62	174.88	82.67

The configuration tested included CGG Insight Earth running on Microsoft's Windows 7 64-bit operating system. The hardware used included a Lenovo ThinkStation P900 with 128 GB of memory, dual Xeon E5 CPUs, NVIDIA K5200 graphics and dual Telsa K40 compute graphical processing units, SanDisk's 1.6 TB card; and a SAS drive used to mount the operating environment.

Maximizing Performance by Pairing CGG's Insight Earth with Lenovo's P900 Workstation and SanDisk's Fusion ioFX

Summary Observations

As individual processor performance has reached its physical limitation in the 6 GHz range, vendors have had to develop new approaches in order to increase systems performance. We are seeing new "accelerated systems" designs which effectively eliminate system-level bottlenecks. The Lenovo ThinkStation P900/Fusion ioFX design discussed in this report is one such design.

The performance bottleneck addressed by the Fusion ioFX solid state memory card is storage. By offering improved throughput and faster load times, this card significantly outperforms traditional storage – delivering data faster to Intel and NVIDIA processors for faster results. In the case of data-intensive applications such as the analysis of seismic data, Fusion ioFX improved overall performance by almost one-third. For oil and gas professionals who are under pressure to produce better results faster, the combination of Lenovo's P900 workstation and SanDisk's Fusion ioFX represents a means to significantly accelerate seismic interpretation.

Clabby Analytics
http://www.clabbyanalytics.com
Telephone: 001 (207) 846-6662