

**COTS**

JOURNAL

The Journal of Military Electronics & Computing

# Bonds Between Develop and Deploy Phases Get Stronger

## **JEFF'S PICKS:**

JEFF'S TOP SPACE-QUALIFIED  
ELECTRONICS

PCI EXPRESS PAVES NEW PATH  
FOR SOLID STATE DRIVES

## **DATA SHEET:**

COM EXPRESS BOARDS



Rugged Storage: From RAID to SSD

# High Capacity Storage Systems Look to PCI Express

Defense applications today have an enormous appetite for data storage. With legacy form factors and interfaces play a role, PCIe SSDs are the future.


Steve Gudknecht, Product Manager, Elma Electronic

**M**assive, data-intensive applications are becoming a fixture in today's C4ISR platforms. Fewer boots on the ground means more electronic eyes and ears filling in the gaps and coordinating information. C4ISR systems need to juggle huge amounts of signal and imaging data from incoming intelligence, surveillance and reconnaissance sources, drawn from radar, sonar, high definition video and infrared sensors to name just a few.

Data-intensive applications continuously drive demand for high volume local storage capacity. High bandwidth incoming data streaming into the front end of advanced systems can create a number of challenges with regard to the storage array on the back end. By definition, SWaP (size, weight and power) considerations in deployed equipment increasingly demand that on-board systems economize on all three metrics. Ensuring sufficient storage volumes in the space allowed is a pretty basic idea .

## SSDs Fill Storage Gap

Fortunately, advancements in the density of solid state drive technology are helping to greatly alleviate that space problem for now. Additional concerns, such as read/write speeds, security, redundancy, signal integrity and ease of access, can be addressed by system and board designers who are well versed in storage subsystem



**Figure 1**

A storage subsystem connected via four SATA III backplane ports resulting in an ample 20 Terabyte storage array across just three slots, including the CPU board.

architecture tradeoffs and connectivity options.

The simplest and most common way to architect local storage using today's wide selection of slot card form factors in VPX, VME and cPCI is to take advantage of SATA ports on new single board computers (SBCs) offered by leading board manufacturers. While this is not necessarily new, many things have changed over the past several years that make this connection method even more desirable for building high capacity arrays.

The two most notable are that drive density has grown in leaps and bounds, and the number of SATA ports supported by new CPU silicon has increased—and board manufacturers are making use of them. Leading edge drive density has grown by a factor of 10 over the last five years and four SATA III ports are common on a 6U board, and even some 3U VPX boards. A SATA III drive with its 6 Gbps top end data rate provides the bandwidth needed for higher speed application needs. And these drives are hitting their stride now with numerous suppliers offering a wide range of form factors.



## A Variety of Form Factors

It's no surprise that VPX is the most capable board standard when it comes to taking advantage of these new developments. Central to the discussion is overcoming bandwidth bottlenecks and the "gotchas" along the signal path. As with any OpenVPX system design, it helps to work with a company able to provide the board, backplane and chassis end of the equation, since oftentimes a new backplane design may be needed to support the system requirements.

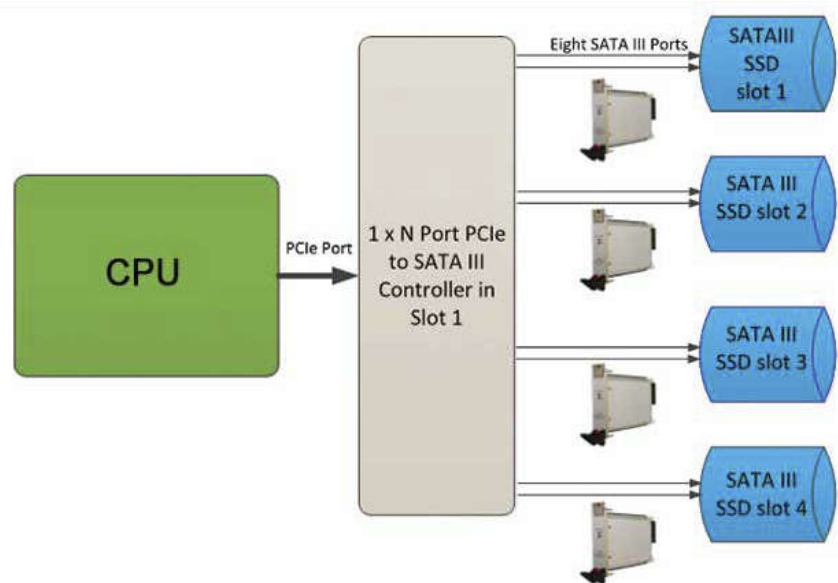
CPU speeds and multi-core processing have advanced to the point of supporting the new data speeds, but getting from the CPU to the drives includes passing through at least one backplane connector pair. Multi-gig connectors in OpenVPX backplanes can handle signal speeds up to 12 Gbps providing acceptable signal integrity – far more than required for SATA III drive connectivity. Systems are becoming more integrated and offering far higher densities in on-board storage arrays (Figure 1).

Here, each storage carrier contains two 5 Gbyte SSDs. An advantage to 3U form factors is, of course, system size, since SWaP is king in defense applications—especially in systems destined for air frames. Conduction and air cooled systems in OpenVPX have an advantage with an option for direct slot-to-slot connections through the backplane with no cables or transition modules.

## Older Form Factors Too

Even older form factors find homes in applications requiring an ever-increasing need for more storage capacity. This basic architecture can be applied to VME or cPCI systems, as well. Although they operate at reduced data speeds, these older systems can be upgraded as missions evolve to support higher storage volumes, if slower data rates are not the paramount concern. It's important to note that backplane architecture and connectors in cPCI and VME systems are limited in terms of signal speed and top out at SATA I (1.5Gbps) and SATA II speeds (3 Gbps) respectively.

Regardless of form factor, all that capacity usually requires some form of data redundancy, and until recently, hardware RAID controllers were the preferred method over software RAID solutions. Software RAID methods slowed down yesterday's



**Figure 2**

One of the multiple benefits of using the PCIe ports showing up on today's SBCs is to free up the SATA ports for other purposes, while making use of the PCIe pipes as intended.

CPUs to the point where other processes suffered.

But higher clock speeds and multi-core processors have almost eliminated this issue in many applications. As a result, high capacity storage arrays can be RAID-protected using Linux-based software RAID, reducing hardware without bogging down the CPU. One-to-N SATA ports internal to the CPU provide sufficient aggregate bandwidth to support the number of SATA III ports offered—up to their theoretical maximum speed. Other factors impact actual data rates, which is what then limits the data rate performance.

## Enter PCI Express

Another way to amplify the amount of storage within an array is to use the PCIe ports now proliferating on SBCs. And OpenVPX is best suited for this methodology. Newer SBC designs can provide four PCIe Gen 3 x8 ports with configuration options to create subsets consisting of multiple x4 or x2 pipes as well.

Using an off-the-shelf controller supporting PCIe Gen 2 x 8 inputs with outputs to eight SATA III ports can, for example, yield aggregate data rates across a PCIe Gen

2 x8 pipe up to a theoretical 32 Gbps. This is 30 percent higher than the theoretical aggregate maximum of 24 Gbps across four SATA III ports coming straight out of the CPU to the SATA III drives. Elma Electronic's 3U VPX Model 5336/6 controller/dual drive SATA carrier card and 5332/3 dual drive SATA carriers are designed for this type of approach (Figure 2).

The bigger benefit is storage arrays up to 8 SATA ports enabled by the available PCIe-to-SATA controllers. In addition, multiple PCIe ports can be used in the same way with each supporting separate arrays. Consider multiple PCIe SBC ports with each driving data to up to eight SATA III drive volumes. In Figure 2, high bandwidth sensor data is pulled into the system front-end via an FPGA for preprocessing.

The resulting data is then processed through the SBC to a controller/carrier card for break out to eight separate drive volumes across four carrier cards. This approach makes use of known SATA drive technology, while tapping into PCIe connectivity. A PCIe-to-SATA controller approach for establishing a storage array has added flexibility when it comes to RAID solutions, since RAID can either be applied

via software or via built-in RAID options in the controller itself. In terms of current and time-tested technology, straight SATA and PCIe-connected SATA arrays are the norm, and take advantage of a large ecosystem of software and hardware products.

### PCIe SSDs: The Next Frontier

Using a PCIe-to-SATA controller—essentially an interface protocol converter—has its ups and downs. As we've seen, it opens up new and creative ways to build storage arrays, as it lends another use for PCIe ports. It also adds hardware to the system, but there is a data rate penalty to pay almost any time a conversion takes place due to its overhead and management. Also, choices of newer PCIe-to-SATA controllers on the market are currently limited.

Given these various points, and since the embedded industry is always looking for new and better ways to improve performance, the next logical frontier revolves around developing straight PCIe-connected SSDs to take advantage of the quantum data rate leaps they provide and the flexibility offered in the PCIe architecture. This would eliminate the controller between the CPU and the array and replace it with a PCIe switch. The result is an increase in the top-end aggregate data rate as well as higher na-

tive data rates per SSD (Figure 3).

Several SSD manufacturers are now beginning to sell the devices in various sizes featuring native PCIe connectivity. Soligen now offers its Triton 2.5" SSD with a PCIe Gen 2 x 4 interface. Other manufacturers are joining in, with products boasting PCIe Gen 3 capability and, as a result, raw drive data rates can climb to as much as four times the SATA III interface rate of 6 Gbits/s.

### Advantage of Faster Speeds

Consider a leading CPU that is driving a PCIe Gen 3 x 8 pipe through a PCIe switch at a maximum theoretical data rate of 64 Gbits/s. The implication is that for systems where fast data rates are the priority over capacity, fewer drives are needed due to the increase in per drive bandwidth. VPX board designs that support PCIe SSDs are forthcoming and will enable larger and faster extensible storage arrays than have been available in the embedded computing industry. As with the initial storage architectures discussed, RAID strategies for data redundancy and protection are supported via the PCIe fabric.


Beyond RAID protection, nearly all SSD products targeting defense applications now offer 256 AES-compliant data encryption as well as multiple levels of data erasure

methods specific to various DoD agency requirements. Write-protection is a given with custom implementations, as required by end users.

### Rich Set of Choices

With the evolution of connectivity choices and SSD technology comes a rich selection of architectural options to build local storage arrays, and data hungry applications will continue to push the boundaries of storage technology. Architecting a high capacity custom storage array as part of an overall computing platform requires a holistic approach taking into consideration the system level requirements. SATA and PCIe-based strategies will likely co-exist for years to come as each has its place, and balancing the system in terms of bandwidth and performance is key.

Choosing the right connectivity method is crucial, such as the use of cabling for inter- and intra-chassis communication, rear transition module usage or backplane slot-to-slot connectivity vs. front panel connections. VPX systems, by design, rely primarily on backplane slot-to-slot communications. In HPEC systems, designers must work within the given SWaP envelope—often with seemingly mutually exclusive design targets.

Software challenges abound when sorting through driver issues to obtain optimal performance, as do board interoperability challenges. Heat dissipation issues at the chassis level require keen attention to detail and experience with thermal management issues and rugged chassis designs must hold up to high levels of shock and vibration as required in the end application. While selecting the right storage for a system design is key, choosing a competent supplier with proven expertise in both storage and system definition enables faster time to market, streamlined project management and satisfied end users. 

*Elma Electronic  
Fremont, CA  
(510) 656-3400  
www.elma.com*



**Figure 3**

The 533x Series high capacity storage modules enable extensible SATA storage arrays.