Unified Approach Paves New Path for Military Data Storage

A new class of data storage products for military embedded systems has recently emerged. This new class of unified, scalable storage enables common storage blades to satisfy diverse storage requirements.

Greg Bolstad, Chief Systems Architect
Critical I/O

Data storage is an increasingly important component of many military embedded systems. In particular, Intelligence, Surveillance and Reconnaissance (ISR) systems generate large volumes of data that often must be recorded in real time for later analysis. These systems require high storage capacity and performance, usually with size, weight and power (SWaP) constraints. Often there are different types of storage requirements such as high bandwidth embedded data recording as well as more generic file serving and general purpose RAID storage needs. Centralizing a system's diverse storage requirements using a common and unified storage platform can dramatically reduce cost, weight/power and greatly simplify storage management. Diverse and separate storage devices are combined into a single, common and scalable storage architecture.

Storage implementations used in conjunction with embedded systems have historically fallen into two categories. One is low capacity, low performance embedded storage boards. The other is higher capacity, higher performance, but physically much larger and heavier external storage boxes or subsystems. However, current flash-based Solid State Drive (SSD) technology—combined with optimized storage controller architectures—has fueled the development of embedded storage blades that provide high levels of consistent performance, reliability and capacity.

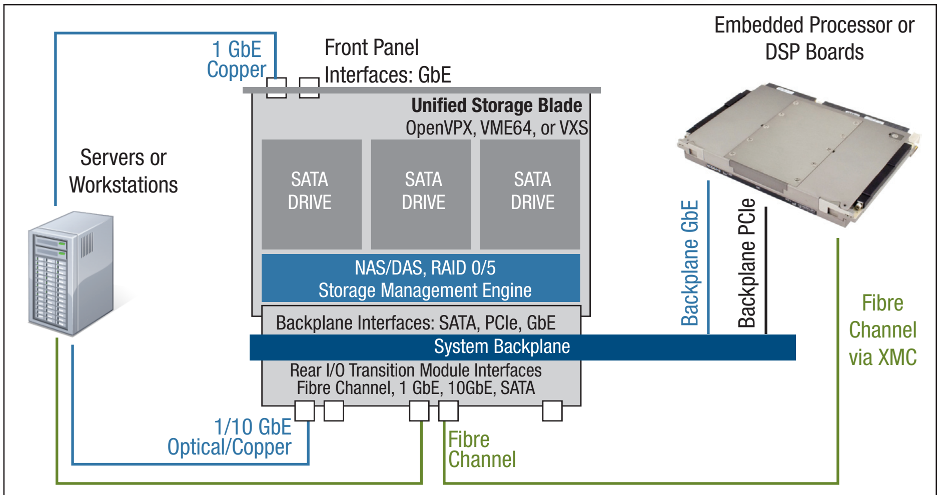

Figure 1

A Unified Storage Blade provides a complete embedded data storage solution that supports a wide variety of storage protocols and interfaces.

Unified Approach

A recent innovation in embedded storage that supports these diverse needs is the concept of unified, scalable embedded storage. The building blocks of this approach are flexible storage blades. Unified storage means that all of a system's storage needs (the various storage models, protocols and interfaces) can be supported by a single storage blade. And since these blades are scalable, multiple instances of the storage blade can be aggregated to provide higher levels of capacity and performance, while still being integrated, managed and used in exactly the same manner as a single blade.

The flexibility of a unified storage

blade architecture allows it to be used for a large variety of embedded storage applications. It is possible to replace large power-hungry external RAID and NAS boxes with a compact, simple, high-performance and high-reliability single blade solution. Some typical applications include Intelligence, Surveillance and Reconnaissance (ISR) systems and Radar/Sonar/Imaging data recording and playback.

Unified Storage Blade Architecture

The architecture and connectivity options of a typical unified storage blade are shown in Figure 1. These blades are generally in OpenVPX or VXS form factors, and contain onboard SATA drives for storage, generally flash-based SSDs. Such blades sport a high-performance storage management engine and provide support for a variety of interfaces to connect storage clients to storage blades.

The storage management engine implements a "storage stack" that controls storage and interface functionality. A unified storage stack supports both file access (Network Attached Storage - NAS) clients by hosting one or more server file systems, which are then exported to NFS, CIFS, or FTP clients. Block access (Direct Attached Storage - DAS, aka RAID/JBOD) clients are supported by providing one or more "virtual disk" or RAID targets, which are then exported to clients using iSCSI, FCoE, PCIe DAS, or Fibre Channel protocols.

RAID offers the advantage of aggregating the performance and capacity of several SATA storage devices, which provides higher capacity and better performance as compared to individual partitions. If RAID 5 is used, it offers the further advantage of parity-based data protection to prevent loss of data in the event of a drive failure.

Storage Usage Models

Data storage is divided into two main categories: Direct Attached Storage (DAS), which provides block level storage access (sometimes thought of as RAID or JBOD storage); and Network Attached Storage (NAS), which provides shared file level storage access. The difference between these two is primarily where the file system is hosted. As illustrated in Figure 2, for the DAS block access storage model, the storage client hosts a local file system. For the NAS file access storage model, a server file system is hosted within the NAS storage blade itself. NAS storage is always accessed by clients via an Ethernet network interface, using protocols such as NFS, CIFS/SMB and FTP. DAS storage is accessed using a variety of interfaces/protocols, including PCIe, Fibre Channel, Ethernet/iSCSI and Ethernet/FCoE.

Unified storage blades support both NAS and DAS operation, and typically provide both RAID 0 and RAID 5 storage capabilities. RAID 0 is a performance-centric model, which stripes data across all drives. RAID 5 also stripes data across drives, but adds a parity block for each data stripe, thus providing data protection in the event of a drive failure. Because one "drive's worth" of performance and capacity is essentially consumed for parity in RAID 5 implementations (note that parity is actually distributed among all three drives), the write performance and capacity when using a RAID 5 configuration will be reduced as compared to RAID 0.

Storage and Scalability

Unified storage blades may also provide scalability in storage capacity and performance. Multiple blades can be aggregated, using the high-speed serial paths available in current generation OpenVPX and VXS backplanes to implement PCIe expansion paths between multiple storage blades.

Several different aggregation models

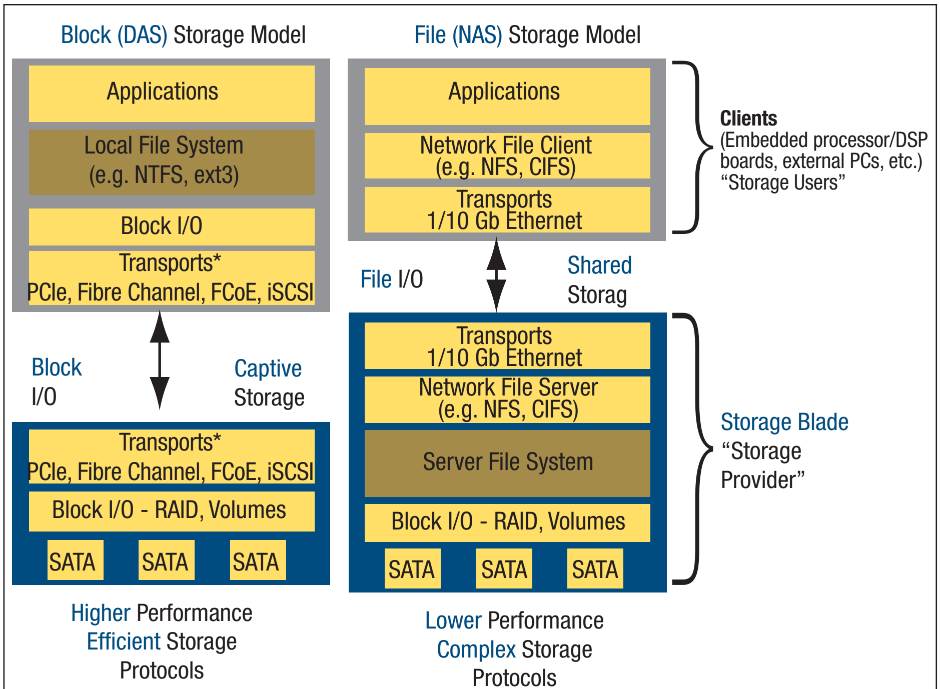

Figure 2

The key differentiator for File vs. Block Storage Access is where the storage file system is hosted—either on the storage client or the storage blade.

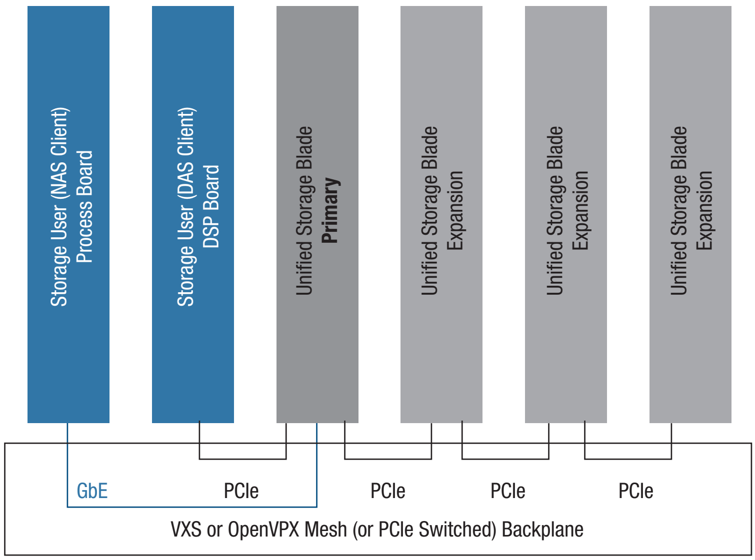

Figure 3

Unified Storage Blades can support aggregation of multiple blades for higher levels of performance and capacity.

can be used, one of which is shown in Figure 3. In this configuration, one storage blade acts as the "primary" blade, with which the storage users (clients) communicate. This example shows one client using file sharing (NAS), while the other client is using block access (DAS). The primary blade then uses backplane PCIe connections to three expansion blades to aggregate the capabilities of all the blades. This makes the set of four blades appear to the storage clients as though they were a single blade.

Physical Interfaces

Storage clients can connect to storage blades using backplane, front panel, or rear transition module (RTM) connections. A variety of interfaces are typically supported.

Ethernet (NAS) – NAS operation requires an Ethernet connection between NAS clients and storage blades. Multiple options for Ethernet connectivity include front panel, backplane and rear transition module (RTM) connections. Many standard VPX and VXS systems provide 1 Gbit Ethernet connections integrated into backplanes, and these can be used to provide NAS connectivity between processor boards and storage blades.

PCIe (DAS) – PCIe connections are generally made using embedded backplane high-speed fabrics, though in some situations the connections may also be made using RTM PCIe access. Storage clients run a lightweight PCIe DAS storage driver.

Fibre Channel (DAS) – Storage clients may use standard Fibre Channel PMCs/XMCs along with the use of standard Fibre Channel initiator drivers.

Storage Protocols

NAS (Network Attached Storage) capability provides file-level access to data. This allows data storage to be accessed and shared via standard file access protocols. These include NFS, CIFS/SMB, FTP and High Performance Streaming NAS. DAS clients use block level storage access protocols to store or access data. There are several commonly used DAS block level protocols, including PCIe Direct Connect, Fibre Channel, Ethernet/iSCSI and Ethernet/FCoE. PCIe DAS is an extremely efficient method of adding high-performance RAID storage to PCIe enabled embedded processor boards. Storage clients use a lightweight PCIe client (host) driver that runs on the client board. Fibre Channel, meanwhile, is a multi-gigabit storage networking technology, which is the standard for storage area networks (SAN). Clients use a Fibre Channel interface and driver to connect to storage blades.

Performance Considerations

Storage blade performance is determined largely by three factors: the storage access protocol used, the storage access pattern and the choice of the physical storage media. For DAS (block access) applications, the client processor hosts a file system that is provided by the client's operating system. The client thus utilizes storage in a low level block mode. The allocation and use of these low level storage blocks is controlled completely by the client file system, thus (unlike NAS storage) DAS stored data cannot typically be shared between multiple clients. Because the DAS transport protocols are extremely efficient, data transfer rates for DAS storage can be very high, over 700 Mbytes/s.

Because NAS (file sharing) storage access is file based, NAS managed data can easily be shared among multiple clients. But data transfer rates for NAS storage are more moderate (as compared to DAS operation), due to the additional complexity of the NAS protocol stacks. A typical data transfer rate for NAS operation is 100 Mbytes/s.

Example System

An example of a unified storage blade is Critical I/O's new StoreEngine (Figure 4) storage blade. It provides simultaneous block (DAS/RAID) and file sharing (NAS) access to up to 3 Terabytes of onboard SATA-based SSD storage at rates of over 700 Mbytes/s for DAS access, and over 150 Mbytes/s for NAS access. Multiple units may be aggregated for even higher performance and larger capacities. Hardware-based RAID0 and RAID5 capabilities are built in, along with a simple but comprehensive web-based configuration and management interface. Available form factors include 6U OpenVPX, VXS and VME64.

StoreEngine provides protocol support for all major storage access protocols, including file sharing (acting as a file server) using NFS, CIFS/SMB, or FTP, as well as block level storage access (acting as a RAID or virtual disk) using PCIe or Fibre Channel. StoreEngine supports simultaneous usage of all of these protocols, so users may "carve" up the unit's SATA SSD storage such that one portion is used for a NFS file server with RAID5, another portion used for a CIFS file server, and yet another portion exported for block level storage via PCIe or Fibre Channel. Special high-speed recording modes are also supported by StoreEngine, allowing aggregation of the units to record data at rates of over 2 Gbytes/s with aggregated recording capacities of up to 12 Terabytes using four StoreEngines.

The flexibility and performance available from the new generation of unified embedded storage blades allows them to be used for a large variety of embedded system applications, replacing large, power-hungry external RAID or file server boxes with a compact, simple, high-performance and high-reliability single blade storage solution.

Critical I/O
Irvine, CA.
(949) 553-2200.
[www.criticalio.com].

Figure 4

The StoreEngine unified storage blade provides simultaneous block access (DAS/RAID), file access (NAS) and high-speed recording access to onboard RAID 0/5 solid state storage.