Using StoreEngine™ and StorePak™ as a NAS File Server

Abstract

StoreEngine's NAS (Network Attached Storage) capability provides file-level access to its onboard RAID volume to a networked environment. This allows data storage to be accessed and shared, using standard file access protocols such as NFS and CIFS, with other devices that are connected via a standard Ethernet network. StoreEngine provides NAS data access rates of up to 200 MBytes/s. NAS is ideal for systems and applications that require file-level and shared access to data among multiple clients at moderate access rates. StoreEngine (and optional StorePaks) can provide up to 3TB of file storage per blade using either fixed or easily removable SSD storage media.

StoreEngine/StorePak Overview

StoreEngine and StorePak are flexible storage building blocks that are used to implement flexible, scalable data storage systems. **StoreEngine** is an ultra-high performance *Storage Controller* blade that can also host up to 1.5 TB of non-removable on-board SSD storage. The StoreEngine single slot blade can simultaneously serve block data (like a disk drive or RAID system) as well as NAS file sharing (like a NFS/CIFS file server). **StorePak** is a *Storage* expansion blade that can host up to 3 TB of easily removable and hot swappable SSD storage.

StoreEngine, with optional StorePaks, provides unmatched storage capability, ultra high performance and high capacity all within a small size, weight, and power (SWaP) footprint. StoreEngine is ideal for high bandwidth embedded data recording, NAS file serving, and general purpose RAID applications. StoreEngine is easily scalable in capacity and performance by simply adding additional StoreEngine and/or StorePak blades.

Using StoreEngine/StorePak for NAS File Sharing

StoreEngine supports Network Attached Storage (NAS, aka File Server) applications using 1Gb or 10Gb Ethernet connections. Four ports of 1 Gb Ethernet are built in to StoreEngine, while 10Gb Ethernet is available using an optional Rear Transition Module (RTM). StoreEngine also provides options for both fixed and removable SSD storage media. The fixed media version is based on the StoreEngine blade, each hosting three fixed SSD drives. The removable media version uses a combination of StoreEngine VPX blades, along with StorePak VPX removable drive assemblies. Each hot-swappable StorePak driver carrier unit hosts six SSD drives.

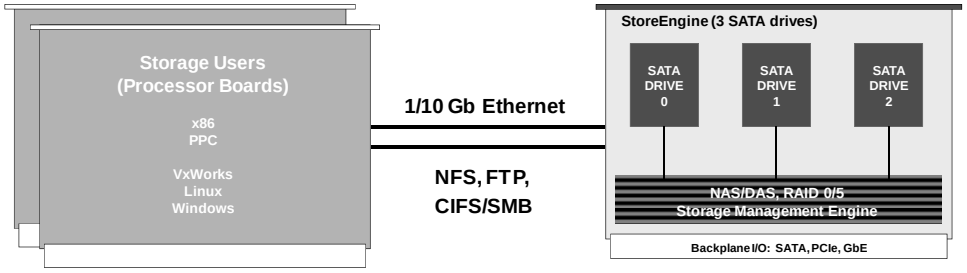

Figure 1. StoreEngine NAS file sharing -- 1.5TB non-removable SSD storage

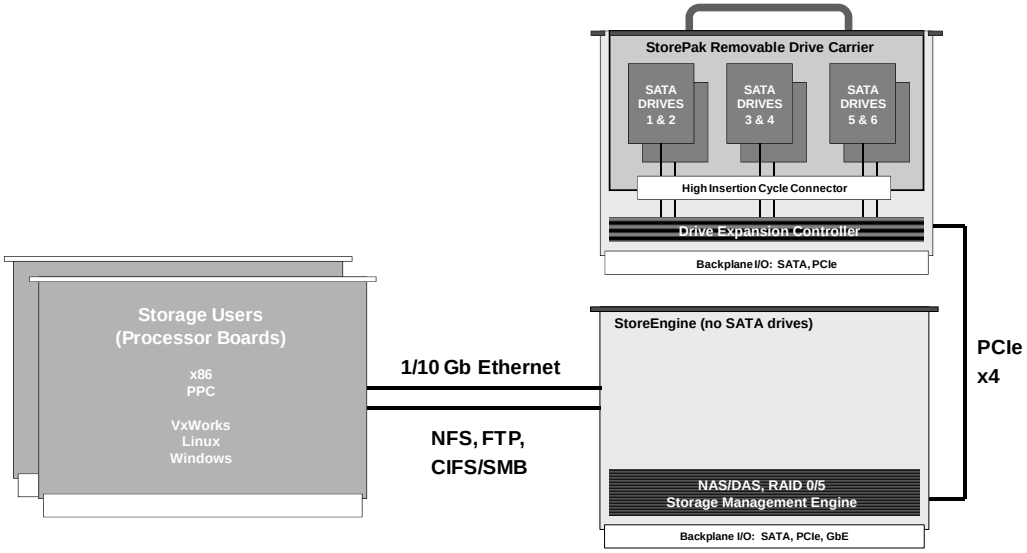

Figure 2. StoreEngine + StorePak NAS file sharing -- 3TB hot-swappable, expandable SSD storage

Data Storage Models

Storage usage models are divided into two main categories: Direct Attached Storage (DAS), which provides ***block level*** storage access (including RAID), and Network Attached Storage (NAS), which provides ***file level*** storage access. The difference between these two lies largely in where the file system is hosted, as illustrated in figure 3 below. For DAS storage systems, the storage client hosts a *local file system*, while for a NAS system, a *server file system* is hosted within the NAS storage device 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, and Ethernet/iSCSI, and Ethernet/FCoE.

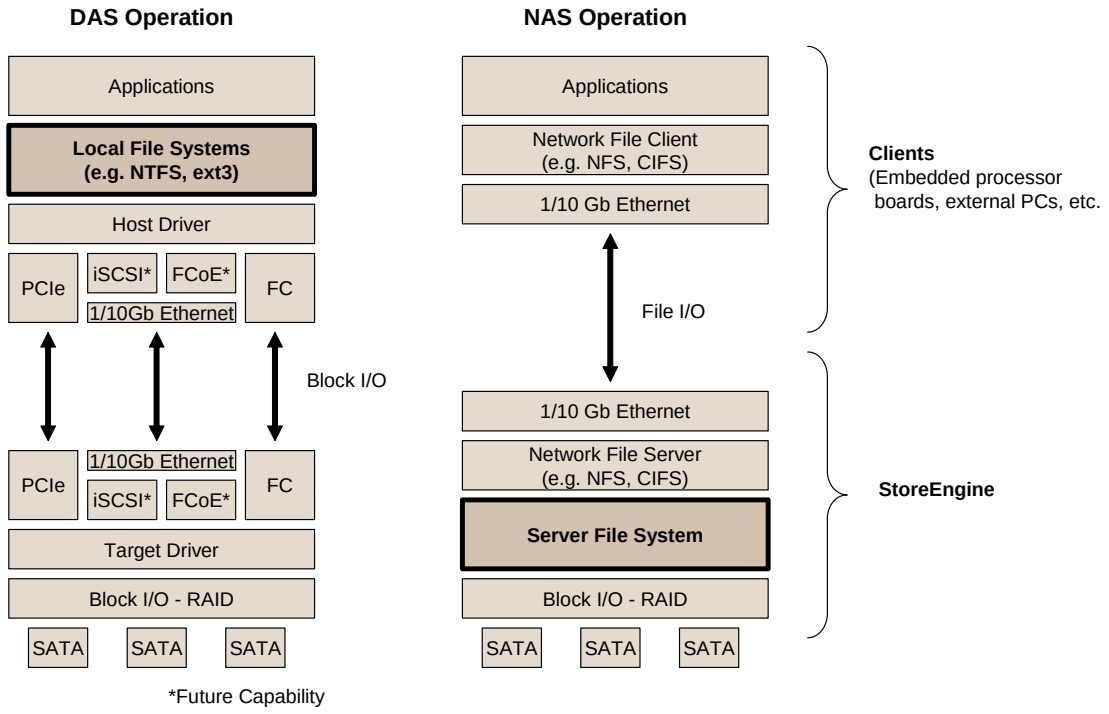

Figure 3. StoreEngine DAS operation contrasted with NAS operation

StoreEngine NAS Application

StoreEngine's NAS (Network Attached Storage) capability provides file-level access to onboard storage volumes in a network environment. This allows data storage to be accessed and shared via standard file access protocols including NFS, CIFS/SMB and FTP by NAS client devices that are connected to StoreEngine via standard Ethernet networks. A unique ultra high performance streaming NAS protocol, running over a 1 or 10 Gb Ethernet network is also available.

Because NAS storage access is always file based, NAS managed data can easily be shared among multiple clients. Data transfer rates for NAS storage are more moderate (as compared to DAS usage), due to the additional complexity of the NAS protocol stacks. StoreEngine typically supports NAS data access rates of up to 200 MBytes/s.

NAS clients are connected to StoreEngine via a 1 or 10 Gb Ethernet network. 1 Gb Ethernet connections may be via front-panel Ethernet, or backplane Ethernet connections, while 10 GbE connections are made using the optional 10 GbE rear transition module (RTM).

NAS Protocols

StoreEngine supports commonly used NAS file sharing protocols including NFS v2/v3, CIFS/SMB, and FTP.

NFS

The Network File System (NFS) allows client computers to access files stored on a StoreEngine file server. The NFS file server on StoreEngine is "mounted" to the client in a manner that allows files to be accessed in the same manner as local files are accessed. NFS may optionally be used with host name based client authentication. NFSv3 can be configured to provide a higher level of performance than NFSv2.

CIFS/SMB

The Common Internet File System (CIFS), also known as Server Message Block (SMB) is the network file system used by Windows. It operates as an application-layer network protocol mainly used to provide shared access to files via a network. CIFS/SMB supports user based client authentication, in a number of flavors.

FTP

File Transfer Protocol (FTP) is a standard network protocol used to copy files from one host to another over a network. FTP is built on a client-server architecture and utilizes separate control and data connections between the client and server applications. FTP is used with user-based password authentication or optionally with anonymous user access.

High Performance Streaming NAS

StoreEngine also supports a unique high performance streaming NAS implementation. Streaming NAS is oriented towards high bandwidth data capture and playback, and allows very high performance NAS storage access (>1 GB/s) using a very simple protocol. A simple, ultra-high performance file system hosted on the StoreEngine manages the SATA drives directly.

Ethernet Connectivity Options

NAS operation requires an Ethernet connection between NAS Clients and StoreEngine. Multiple options for Ethernet connectivity, including front panel, backplane, and rear transition module (RTM) connections. Note that while StoreEngine has six physical Ethernet connections, it has only four logical Ethernet ports which are shared among the six physical connections. The StoreEngine User's Manuals define the specific mapping of logical to physical ports for each StoreEngine model.

Front Panel Ethernet Connections

Front panel 1000Base-T connections (VXS and VME only) are available on some models. These are standard RJ-45 connections which accept normal CAT5e or CAT6 cables.

Backplane Ethernet Connections

Backplane 1000Base-T connections are available on all models. These backplane connections are accessed using an optional Rear Transition module (RTM), which then provides standard RJ-45 connections.

Backplane 1000Base-X connections are also available on all models. These backplane connections may be used in two ways. Some backplanes contain compatible embedded connectivity which may be utilized. Alternatively, these connections may be accessed using an optional RTM which provides optical Ethernet connectivity via standard LC connectors.

Optional Rear Transition Module Ethernet Connections

Optional Rear Transition Modules (RTMs) can be used to provide Ethernet connectivity in systems where front panel or backplane connections are not available or feasible. RTMs provide up to two optical Ethernet LC connections, or up to two 1000 Base-T RJ-45 connections. Other RTMs are available that provide dual 10Gb Ethernet optical connections.

Example NAS Configurations

Figure 4 illustrates the different Ethernet configurations that StoreEngine supports for NAS operation. These include support for front panel connection to an external Ethernet network, backplane internal fabric direct client-to-StoreEngine connection, backplane internal fabric connection to an embedded Ethernet switch, and optional RTM based rear I/O connections to an external 1 or 10 Gb Ethernet network.

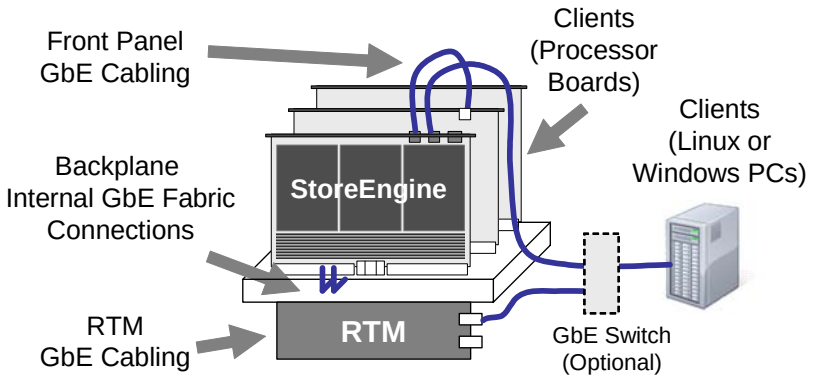

Figure 4. NAS Ethernet Connectivity Options

StoreEngine NAS - Optional Capabilities

RAID Options

StoreEngine storage can be configured to operate as either RAID 0 or RAID 5. Both RAID 0 and RAID 5 modes aggregate the storage of all of the StoreEngine drives into one or more "logical drives". RAID 0 and RAID 5 differ greatly in levels of write performance and data protection.

RAID 0 stripes data across all of the StoreEngine drives, which provides the highest possible read and write performance, as well as the highest usable capacity. RAID 5 also stripes data across drives, but in addition creates a parity block for each data stripe which provides data protection and continued operation (at a lower performance level) in the event of a drive failure. Since one "drives worth" of performance and capacity is reserved for parity in RAID 5, the capacity and write performance when using RAID 5 is approximately 2/3 that of RAID 0 for a standard three drive StoreEngine configuration.

Logical Volumes

StoreEngine supports the optional use of Logical Volumes. Logical volumes provide a method for dividing the available storage into a number of separate volumes. For example, a single RAID 5 array can be created that aggregates all storage. This array can then be divided up into several logical volumes, which may be used for different purposes. (i.e., one volume exported as an NFS server, another volume exported for FTP use). The use of logical volumes offers the advantage of being able to grow or shrink volumes dynamically.

Encryption

StoreEngine provides a option to encrypt all data as it is written to disk. Note that while StoreEngine has hardware accelerated encryption, the use of encryption may still affect performance in some situations.

Secure Erase

Store Engine provides a Secure Erase operation which may be invoked via the Web Management interface. Secure Erase fully erases all data on the selected drive, restoring the drive to an unused condition.

Quick Config

StoreEngine provides two "quick configuration" options. The first is as set of pre-defined configurations that can, with a single mouse click, fully configure StoreEngine to one of six commonly used configurations. The second method allows users to create and store "configuration snapshots". These configuration snapshots can be reapplied with a mouse click to restore a StoreEngine to a known configuration. These configuration snapshots can also be transported to other StoreEngines, allowing users an easy way to "clone" configurations onto additional StoreEngine units.

NAS Performance Optimization

With all data transfer and storage devices, the specifics of how the device is used will affect performance. While there are numerous tunable parameters in NAS applications, there are three basic keys to maximizing performance: 1) maximize the data transfer sizes, 2) use sequential access to files to the maximum extent possible, and 3) enable asynchronous writes.

Maximize Data Transfer Size

Data transfer sizes should be optimized on two levels. First, applications should perform reads and writes in the largest block sizes possible. For example, a single 32K application level read from a file will be much more efficient that 32 1K reads. And second, the Ethernet MTU (frame size) should be set to the largest value that is supported on the network. For most high performance Ethernet NICs and switches, this size is 9000 bytes. Compared to the standard Ethernet frames of 1500 bytes, these 9000 byte "Jumbo" frames significantly increase NAS performance.

Maximize Sequential Access

Sequential data access is always more efficient that random access. There are two reasons for this. Sequential access allows multiple accesses to be gathered in to fewer, larger accesses using read-ahead and write-behind techniques. And second, solid state drives and especially rotating hard drives both perform better in sequential access scenarios.

Enable Asynchronous Writes

Enabling asynchronous writes causes client writes to be stored in buffer cache on the StoreEngine, from which they will later be asynchronously flushed to SATA storage. This improves write performance because it reduces the time that clients must wait for write operations to complete. Note however that the use of asynchronous writes does increase the risk that data could be lost (without notice to the client) in the event that the system is powered off prior to the buffered writes being flushed to storage.

StoreEngine Web Management Interface

StoreEngine NAS functionality is enabled and managed via a simple web based management interface, which supports monitoring and configuration of all StoreEngine interfaces, operating modes, and storage options. Specific web management capabilities include:

- BIT status (self test, voltages, currents, temperatures)
- Storage status (available/used capacity, status, errors)
- Network statistics
- Interface Status (link status, errors)
- Enable/Disable protocols & options
- Manage NAS and DAS exports
- Security and permissions
- Secure Erase

An example of a StoreEngine management page (this specific one is used to create an NFS export) is shown in Figure 5. The other main management categories are shown in the menu on the left side of the page.

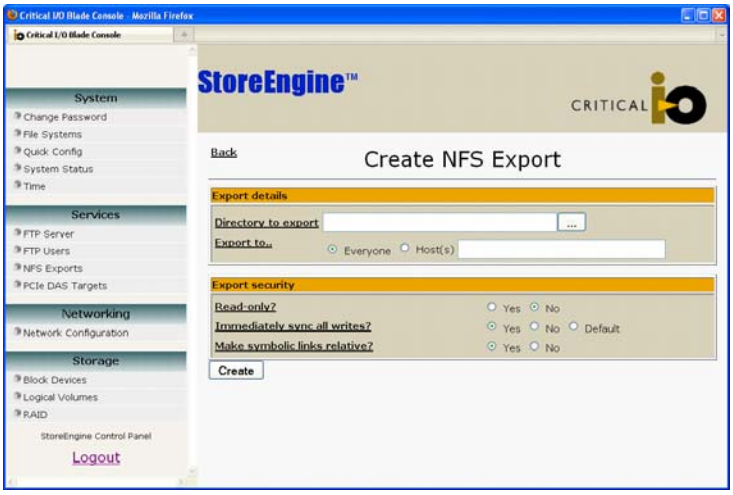

Figure 5. StoreEngine web management interface example page

StoreEngine Hardware Architecture

The StoreEngine hardware architecture is shown in Figure 6. The key components are:

- Storage Management Processor – The highly power efficient PowerPC based storage management processor provides overall control of the StoreEngine, implementing storage management functionality as well as implementing the NAS file sharing protocols.

- Hardware Accelerated RAID – The RAID hardware function provides acceleration of RAID-5 operations.

- PCIe Switch – An 8 port PCIe Gen 2 switch provides internal and external connectivity for StoreEngine. It connects all of the internal resources, and provides up to four backplane PCIe ports (depending on the specific StoreEngine model) for external PCIe backplane or RTM connections.

- Gb Ethernet NICs– StoreEngine has four on-board Gb Ethernet interfaces. Two interfaces provide TCP Assist Hardware capabilities.

- SATA Controller – The SATA controller manages the three internal SATA drives. It operates under the control of the Storage Management Processor.
-
- SATA Drives – Three internal SATA drives provide data storage. Up to four additional external SATA drives may also be connected using the RTM SATA connection. The SATA drives are used only for user data; no StoreEngine management software or configuration information is stored on the SATA drives.

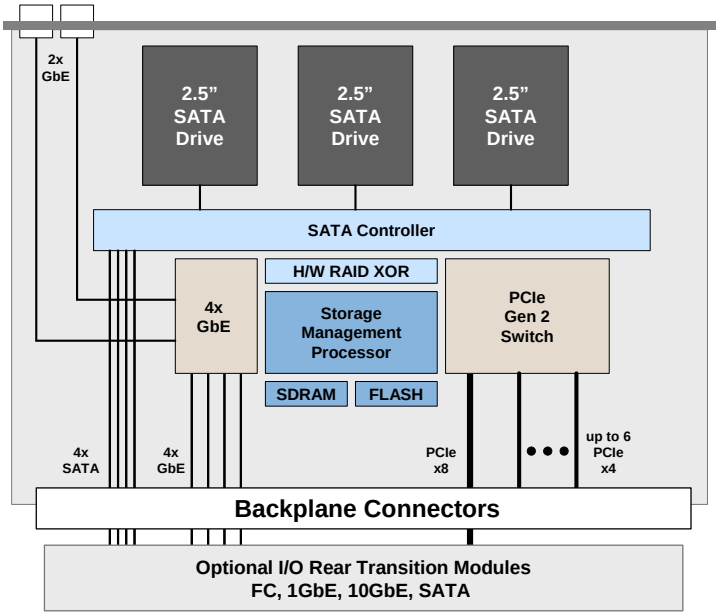

Figure 6. StoreEngine hardware architecture

StoreEngine Software Architecture

The StoreEngine internal software architecture is shown in Figure 7. The key layers in the architecture are 1) the Transport layer, which provides the implementation of the various storage access protocols, 2) the Storage Management Layer, which manages the file and block level storage resources, and 3) the Storage Layer, which provides the interface to the physical on-board SATA storage SSD or HDD devices.

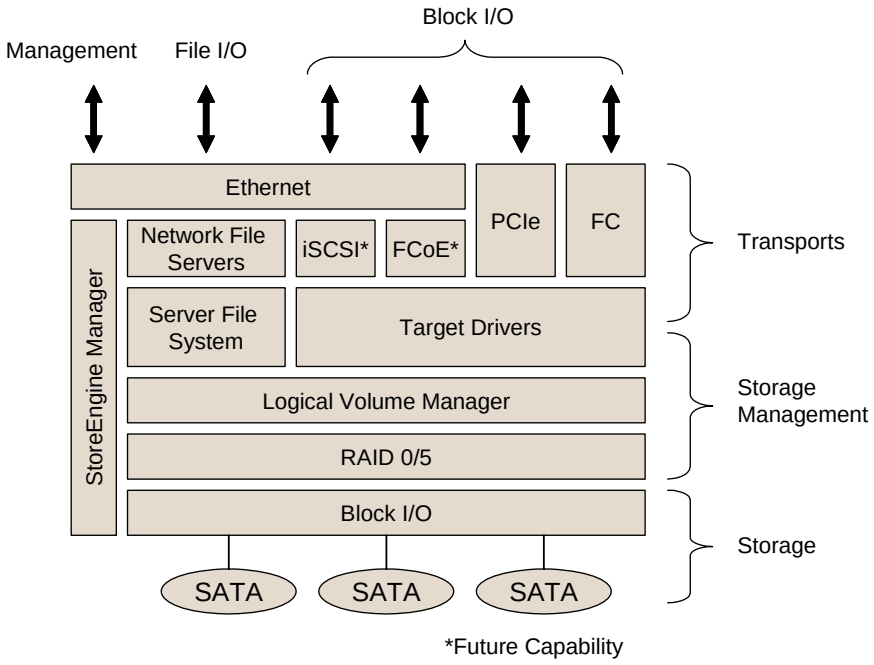

Figure 7. StoreEngine Software Architecture