OVERVIEW
ANX 1500 Accelerated Caching Architecture

Executive Summary

The ANX 1500 NFS Throughput Acceleration Appliance is an NFS caching solution geared to increase aggregate data throughput and decrease latency across Network Attached Storage environments. The goal of this short document is to describe, in simple terms, how the ANX 1500 functions on the network and what effects the appliance has on the path of data. While detailed technical specifications of the ANX 1500 are beyond the scope of this document, an overview of the appliance's key features and hardware architecture will serve as a starting point for inquisitive users.

Alacritech, Inc.
1995 North First Street, Suite 200
San Jose, CA 95112
408.287.9997 | www.alacritech.com
©2011 Alacritech Inc., All rights reserved.
Specifications subject to change without notice.

ANX 1500 Caching System

The ANX 1500 is a caching system that specializes in quickly delivering blocks of data to peers on the network. The appliance focuses strictly on client reads and does not take ownership of data – in other words, the appliance relies on the existing infrastructure's file system to store data. Thus, the ANX 1500 exists on the network and is seen as any other NAS server through which clients mount NFS exports as they would with a traditional NAS device, but client writes are satisfied by other NAS devices in the ecosystem.

Architecturally, the ANX 1500 consists of the five components shown in Figure 1:

- DRAM
- Solid State Disks (SSDs)
- 10GbE TCP Offload (TOE) Network Interface Cards (NICs)
- Software and Hardware

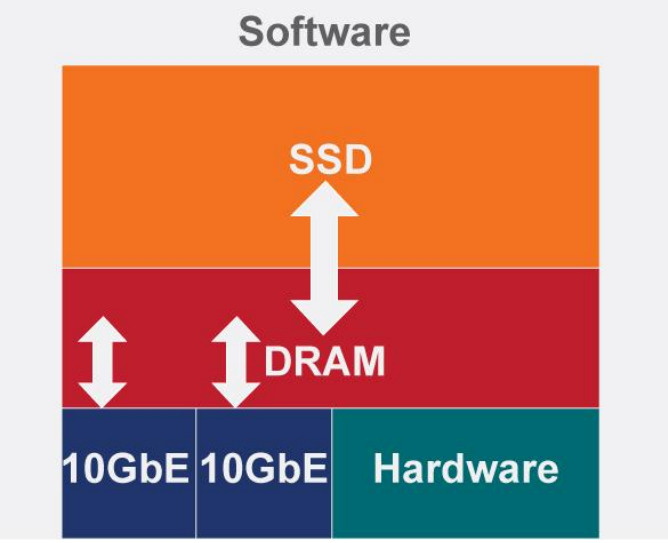

Figure 1

Generally speaking, the ANX 1500 retrieves data from the network and makes that data available to users. Through the use of two Alacritech 10GbE network accelerators incorporating Alacritech's proven TCP Offload Engine technology and NFS acceleration, data is moved from the wire directly into host memory, thus bypassing much of the timely software execution cycles necessary to propagate data through a conventional system. Once the data has been placed into DRAM, that data is then served to its requestor. As data is continually read and placed, the amount of available DRAM diminishes, at which point the data is transferred to SSDs for temporary storage. Data resides on SSD until that block of data is again requested by the client and is then moved back into DRAM for quicker access. This switch-back process continues while clients are accessing given blocks of data; when no further clients are requesting the data, DRAM and SSD resources in the ANX 1500 are recycled to make room for more recent data.

Data Paths

To better understand the deterministic approach by which given blocks of data are transferred from the network to the ANX 1500 and vice versa, the following four cases will be examined:

- Inactive “Cold” Data
- Active “Hot” Data
- Active “Warm” Data
- Committed Data

Case 1: Inactive “Cold” Data

“Cold” data is characterized as data that has not yet passed through the ANX 1500. Cold data can be thought of as inactive data on the network: the data exists, but is currently residing on NAS and is not being accessed. From the vantage point of the ANX, this scenario occurs before the client requests a particular block of data. Once the client requests a file across the network, the ANX 1500 fetches that file from the back-end NAS storage system and moves that data into DRAM. The ANX then satisfies the client read request. The next time this piece of data is requested, a client will read the “hot” data directly from the ANX, which leads us to Case 2.

Case 2: Active “Hot” Data

A client requests a set of data that has recently been read by the ANX. Here, the data is considered “hot”, which means the data resides in the ANX’s DRAM. Any requests for hot data are satisfied directly by the ANX 1500 and will increase performance by reducing latency involved in fetching data from the back-end NAS server.

Case 3: Active “Warm” Data

As more and more data propagates through the ANX, the cache literally becomes full. At this point in time, selected data is moved to and from the ANX’s SSD storage subsystem to ensure that the most active data is able to be satisfied directly from DRAM. Data moved from DRAM to SSD is called “warm” data. If warm data is requested by a client, that data will subsequently be moved to DRAM and satisfied accordingly.

Case 4: Committed Data

As noted above, the ANX 1500 caches and accelerates read data only. In the case that a client desires to save a locally-generated file or commit changes to a file that has been satisfied by the ANX 1500, the appliance acts as a pass-through device, bypassing all caching capabilities and deferring to the back-end NAS server for write operations. It should be noted that under some circumstances, the ANX 1500 will cache write data, but the implications lie outside the scope of this paper.

Conclusion

We have briefly looked at the general architecture of the ANX 1500 NFS Throughput Acceleration Appliance and a few common caching cases of data paths with the ANX present on the network. As designed, the ANX offers a non-intrusive performance addition to Network Attached Storage (NAS) environments by focusing on the most commonly accessed data, while allowing the existing storage environment to facilitate the remainder of data placement functions.