

## Exanet – Next Generation NAS

**Date:** October, 2007

**Author:** Tony Asaro, Senior Analyst

**Abstract:** Exanet ExaStore is a next generation NAS storage system that supports a wide range of applications and environments. Exanet has just announced its 100<sup>th</sup> customer—a major milestone for storage system vendors.

Exanet is one of the best kept secrets in storage. It is a storage vendor that provides a next generation NAS solution with a lot going for it. Which begs the question: What is a next generation NAS solution? It is a NAS solution—in Exanet's case, the ExaStore—that provides a theoretically infinite file system; a cache coherent clustered architecture; high performance for large sequential reads; and good performance for random I/O and writes. Next generation NAS solutions support valuable features including virtual volumes, snapshots, remote replication and built-in monitoring—and they are focused on both NFS and CIFS environments. In a single cluster, the ExaStore can support a large number of controller nodes—or as few as two—and still support lots of capacity (horizontal and vertical scaling). Combine all of these things together and Exanet offers one of the most versatile NAS solutions in the market with the ability to support mainstream file consolidations as well as high performance computing platforms and a bunch of applications in between.

NAS is an interesting market. There are far fewer players than in the SAN arena and really only a small list of leaders in different segments. Though this leaves room for new players and leaders to emerge, it has yet to happen. However, there is an even bigger opportunity. ESG believes that all storage is going to require some file or content level awareness over time. Having this level of awareness provides such a great amount of value—essentially turning data into information—that at some point it will become pervasive. At the heart of this, we believe that a next generation file system will become a core part of how we manage data and access information in the enterprise.

### Exanet ExaStore – A True Clustered NAS System

The Exanet ExaStore is a true clustered NAS system that can support multiple storage controllers in a cluster while still providing a single logical storage system. This enables applications and users to access any piece of data stored on the ExaStore through any of its storage controllers within the cluster. ExaStore aggregates all its hardware resources into a single system including CPU, memory and bandwidth.

The Exanet ExaStore can start out modestly and then transparently grow to a massively parallel system. This is a cost-effective approach that doesn't require end-users to buy ahead of their needs—which is the case with traditional storage systems that have fixed architectures. Additionally, by adding more storage controllers into the cluster, users can continually raise the performance ceiling. But scalability should not create more complexity. Regardless of how big the cluster gets, it should still remain a single logical system to manage. Essentially, it should be just as easy to manage a clustered storage system with 100 nodes as it is to manage one with two nodes. The system contains more physical storage controllers, but remains a single logical system.

As a true clustered storage system, the ExaStore is a next generation solution that is designed to avoid complexity and remain perpetually optimal. It achieves this by natively providing a single level of access. ExaStore NAS presents a single network drive or mount point that can support a massive amount of capacity and store an enormous number of files therein. Here are some key aspects of the ExaStore NAS solution:

- **It supports all types of I/O.** It is great for large files, sequential data and reads and is also good for small files, random data and writes. The Exanet ExaStore system is versatile and works well with a large number of applications and environments.

- **It supports horizontal and vertical scaling.** In other words you can add lots of storage controllers—hundreds—in a cluster if you need CPU, cache and/or bandwidth. And you can also add more capacity to the cluster independent of the storage controller nodes. An extreme example is that you can have 100 nodes with a small amount of capacity or 6 nodes with 750 TB of capacity—and many variations in between.
- **It offers integrated search and indexing.** If you have thousands, millions, billions or more files on your storage system, how do you access relevant information? The answer is to use a search engine. Exanet ExaSearch provides search capability on ExaStore systems. The name of the game is accessing the information you need quickly and precisely. Otherwise, you can easily be drowned by a sea of content. Or, conversely, you may not find the essential information you need.

A wide variety of applications are creating data separately—there is no commonality or correlation of all of this data. To access information, you need to go to the specific applications that created it. If you are doing research on a broad topic, you may need to go to multiple applications to find all of the required information. Federated search enables leaps in efficiency, since it creates commonality where none existed. It enables you to leverage all of your content independent of the various applications. Exanet ExaSearch provides federated search and indexing using metadata and keywords within content (e.g. documents, spreadsheets, etc), raising it from data to information storage.

- **It enables storage provisioning with virtual volumes.** Customers are already implementing ExaStore for applications that require petabytes of storage. At this level, it is important to carve out separate containers for different data sets—and do it quickly without downtime or touching storage arrays. Exanet provides a virtual volume feature called ExaVolumes that does just that.
- **It supports snapshots and remote replication.** For many end-users, these are requisite features. Exanet supports logical snapshots that are space efficient and don't consume a great deal of system resources. Additionally, you can make as many of them as you want. It is also important to point out that their snapshots are cluster aware—meaning that you can create snapshots across the clustered file system. Exanet ExaSync is a remote replication technology that provides a great deal of capability, supporting directory and volume level replication.
- **It provides built-in performance monitoring.** There are many storage tools out there that provide performance and capacity monitoring. However, they require additional hardware and software to install and maintain, which generates additional cost and adds to the storage management workload. Exanet recently added a built-in performance monitoring and capacity trending capability called ExaMonitor that is stored within the cluster itself and requires no additional components.

For years, ESG has felt that Exanet had great technology, but needed to raise the level of its business execution. The industry has a history of good technology combined with great execution pummeling great technology and not-so-good (or bad) execution. However, Exanet seems to understand this and is changing its behavior. There is a lot at stake here and the race is on. The company just celebrated its 100<sup>th</sup> customer and that is a major milestone for storage system vendors—many end-users feel that a system vendor is valid when it hits this number.

Why do next generation NAS and file systems matter? A file system turns data into information. You can manipulate, access, move, copy, de-duplicate and search data efficiently and intelligently using a file system. A clustered file system removes the limits of scalability. It grows as large as needed, when needed—transparently. A versatile file system works with any application and is not bounded to a particular niche. A file system that works with heterogeneous physical infrastructure can work with high-end storage systems for greater levels of protection, performance and availability; it can work with bulk storage for long term retention of data economically; and everything in between—as a single logical system. A radical idea, but it makes too much sense to ignore.

## **ESG's View**

One of the most compelling things about Exanet ExaStore is that it is so versatile. You can have a massive cluster with relatively little capacity or a small cluster with lots of capacity—and variations in between. You can use it for Windows consolidations; high performance computing, database applications or as a long term near-line archive for bulk storage. Many people may reject this notion—using the same solution for such a wide variety of applications because that is what we've become accustomed. And yet that is backward thinking. We need to manage a small number of logical infrastructural devices that float on top of physical systems—leveraging their CPU, memory, bandwidth and capacity as we need it. We are heading toward a more virtualized data center and ESG believes that intelligent file systems can play a major role in accelerating this because they enable you—through virtualization—to manage, access and utilize your data easily and more effectively.