

Isilon: Near-Line and Web 2.0 *Broadening Their Horizons*

Date: July 2007

Author: Tony Asaro, Senior Analyst

Abstract: Isilon is expanding their focus leveraging their strengths including ease of management, cost effectiveness and scalability to near-line storage. They are also turning up the heat on their leadership in Web 2.0 with greater density and scalability of up to an impressive 1.6 PB in a single logical system.

Isilon is best known for their high performance computing, multimedia and online market focus. They smartly went after areas in which the leading storage system vendors are not dominant. Isilon has been able to leverage all of their strengths and become the ones to beat in these verticals. They did this with a combination of excellent technology, product and business execution.

It's smart for Isilon to have this level of focus and it has enabled them to build their business. There is always a strong pull to be all things to all people especially when you can see all of the places where your solution can add value. Isilon resisted that temptation and it has paid off. However, at some point you need to broaden your horizons and advance adjacent opportunities. For Isilon, that time has come.

The new Isilon IQ 9000 and EX 9000 are denser systems that support 750 GB SATA drives. Isilon can combine up to 96 of these systems in a single cluster for up to 1.6 PB of capacity. The current talk is that the systems will be available at a list price of \$2,500 to \$4,000 per TB. That's just \$2.50 to \$4.00 per GB (and street price is going to be less than that). Take a step back and think about this. Isilon doesn't just offer stupid, low-end storage. Isilon is one of the most intelligent, sophisticated storage systems, which provides a ton of value. Additionally, you can easily scale the system to outrageous capacity sizes. Since it is a single logical system it is easy to manage regardless of how big it gets. Isilon makes it possible, easy to manage and cost-effective to implement large to massive data stores.

Dynamic and Persistent Data

There are two distinct types of data: Dynamic and Persistent. Dynamic data is new, constantly accessed and often in flux. Persistent data is for the most part fixed. It does not change. Additionally, Persistent data is not accessed as often. Dynamic and Persistent data can be structured, unstructured or semi-structured data. All types of data live dynamically for some period, whether it represents a Word document, a movie, a credit card transaction or an e-mail. It all ends up as fixed digital content or Persistent data.

It is pretty safe to say that the leading storage systems were designed and brought to market well over a decade ago and were built to meet the needs of dynamic data. At that time the cost of storage was high, making it prohibitive to use for anything but the most important data. Therefore, we limited the number of applications and platforms that we stored on expensive storage systems. Additionally, we kept data stored for relatively short periods of time on these costly systems and instead made piles and piles of backup images onto tape. The requirements for these storage systems were that they provide excellent transactional performance that could scale to high levels of concurrency, resilient architectures that mitigated risk of downtime or data loss and data management features that provide data and disaster recovery. All of this came at a price high in cost and complexity.

To this day, this paradigm (are we allowed to use 'paradigm' or is it still forbidden?) still exists. Certainly the cost of these systems has come down and they have become a bit easier to use, relatively speaking. And there is still a fundamental requirement for storage systems that are built for transactional workloads. The

problem is that these storage systems are being used for every kind of data and environment. Utilizing Tier One storage infrastructure while 60%, 80% and even 90% of the data is persistent—well that's just incredibly inefficient.

Isilon and the Persistent Data Storage Tier

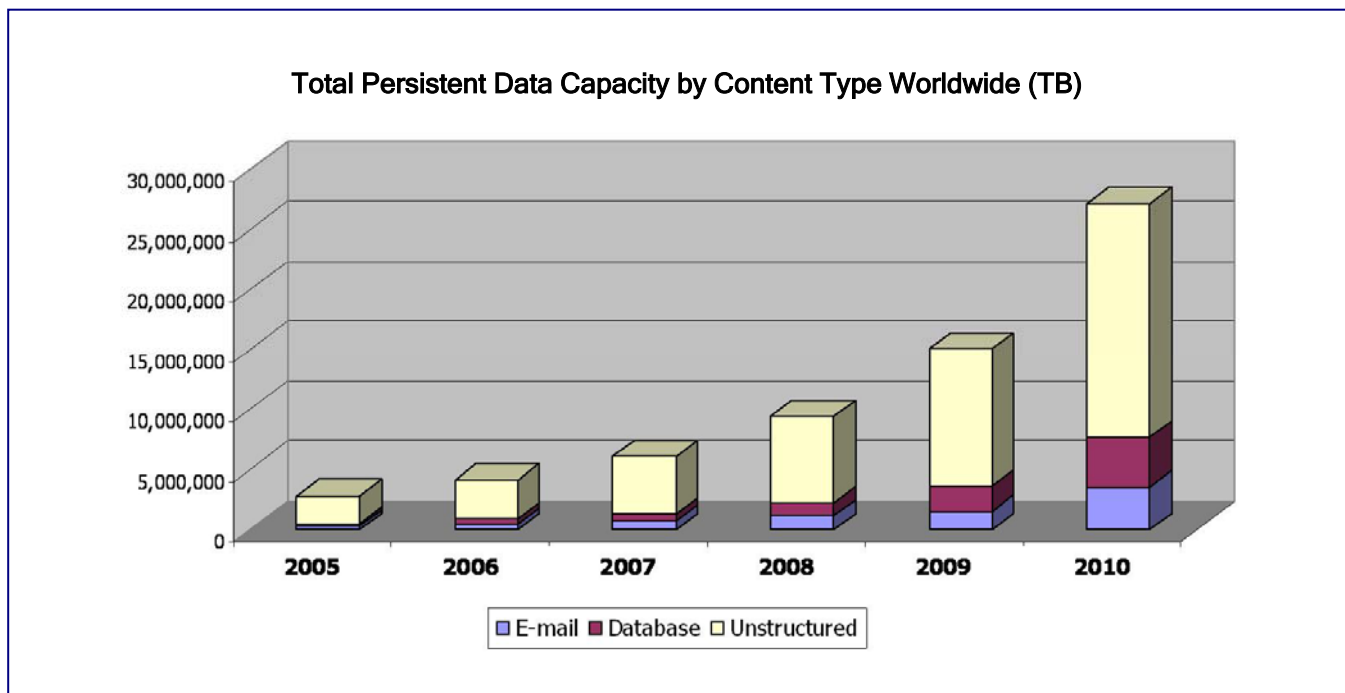
What are the characteristics of the Persistent data storage tier? The challenge is that you can take any one of a number of approaches. However, three elements are requisite: it has to be substantially lower in cost than Tier One, it has to be far less complex and it has to scale to large or even massive amounts of capacity. It's easy to connect the dots on these three elements. It is widely accepted that data growth continues to accelerate, end-users are retaining data for much longer periods of time and the majority of data is persistent—it then stands to reason that this tier of storage needs to support a lot of capacity. As such it needs to be cost-effective, easy to manage and scalable in terms of capacity.

Isilon hits all three of these requisites spot on. Isilon is an extremely easy storage system to manage from the outset and over its entire lifecycle. It is also extremely cost-efficient and can scale up to a whopping 1.6 PB in a single logical system. All of the Isilon value comes into play, including N-way clustering, massive scalability, single file system and cache coherency architecture. All of this works to create one logical system, resulting in a system that is extremely easy to manage, regardless of how large it scales, as well as impressively cost-effective.

Near-Line

ESG estimates that the growth of Persistent data will reach over 26,000,000 TB by 2010 (Figure One). The ratio of Persistent data to Dynamic data is about 10 to 1. However, this ratio will grow to be 100 to 1 in a very short period of time. You just can't efficiently store that much data on traditional Tier One storage systems. It would be impractical, grossly expensive and a management nightmare. Imagine the ramifications of storing 1.6 PB on a traditional Tier One storage system. Compare this to Isilon, a next generation storage system that was designed with these levels of data capacities in mind. A shift has occurred and is on a trajectory. Doing things the same way when the world around you is changing is folly, even though it feels like the safe thing to do. The old adage—when you have a hammer, everything looks like a nail—applies in the way we store data. This approach comes at a huge price in capital and operational costs.

Figure One: ESG Research Forecast of Persistent Data Capacity¹



¹ ESG Research Report: *Digital Archiving: End-User Survey & Market Forecast 2006-2010*, March, 2006.

Near-line, active archive, secondary or tertiary storage are all appropriate terms for how Isilon can be used within this context. If you are looking for an easy to manage, scalable, reliable, low cost solution for any or all of the above, then you should be considering Isilon. Take all of your dormant data—persistent data—and move it from your Tier One storage and to Isilon clustered storage. It will have an excellent home.

Web 2.0

What exactly is Web 2.0 storage? Here is one example: Consider the fact that there will be millions of users accessing a web-based service. They are downloading and uploading images, audio files and video. They may also be leaving messages and e-mails for other people sharing the same or different services. Typically these users are storing data on a service that is free to them. However, the cost of that infrastructure is not free. That is why it is essential that the storage being used for these services is extremely easy to manage, cost-efficient and easily and extensively scalable.

Isilon has already been providing storage to Web 2.0 companies, including KodakGallery, MySpace, and Photobox, for years now. However, their EX 9000 provides a denser and more cost-effective solution. And Isilon is beginning to focus on bringing their solution to brick and mortar companies that are looking to build their own Web 2.0 businesses.

ESG's View

Near-line storage is pervasive and is a horizontal opportunity for Isilon. Additionally, any company that is looking at implementing a Web 2.0 business should consider Isilon for their storage. Isilon has wisely focused on specific markets and applied their competitive advantages and core competencies to become successful. They are now applying their value to broader markets, focusing on specific use models that have wide appeal.

Isilon has always been a great solution for near-line and Web 2.0 applications. In fact, they have an impressive list of customers that are already doing both. The addition of the Isilon IQ and EX 9000 provides even more compelling economics that advance their value proposition in these areas. More importantly, they are setting their sights on horizontal markets. That's the big news and Isilon now has the momentum and resources to make a major impact.